

Document Version

Final published version

Licence

CC BY

Citation (APA)

Kuo, N., Sergejuk, A., Chen, V., & Izadi, M. (2026). Developer Interaction Patterns with Proactive AI: A Five-Day Field Study. In T. Kuflik, S. Kleanthous, L. Chen, G. Jaccuci, & A. Renner (Eds.), *IUI 2026 - Proceedings of the 2026 Conference on Intelligent User Interfaces* (pp. 349-362). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3742413.3789148>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Developer Interaction Patterns with Proactive AI: A Five-Day Field Study

Nadine Kuo*
 JetBrains
 Amsterdam, Netherlands
 nadine.kuo@jetbrains.com

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

Valerie Chen
 Carnegie Mellon University
 Pittsburgh, Pennsylvania, USA
 valeriechen@cmu.edu

Maliheh Izadi
 Delft University of Technology
 Delft, Netherlands
 m.izadi@tudelft.nl

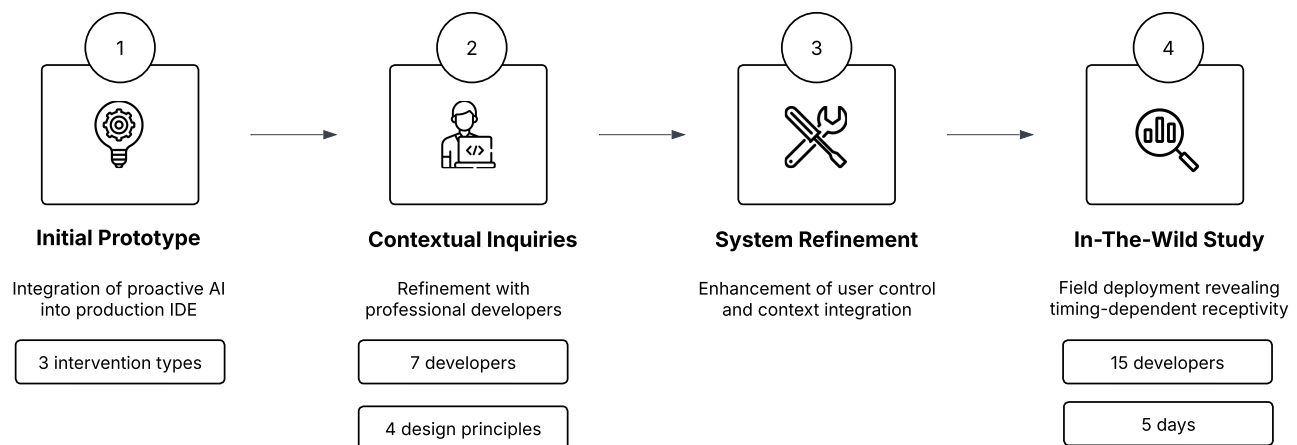


Figure 1: Study Overview. We developed and iteratively refined *ProAIDE*, a proactive feature for code quality of an in-IDE AI assistant (Phases 1-3), then conducted a 5-day field study with 15 professional developers (Phase 4). Analysis of 229 interventions across 5,732 interaction points revealed that timing critically impacts receptivity: post-commit suggestions were more readily accepted than mid-task interventions.

Abstract

Current in-IDE AI coding tools typically rely on time-consuming manual prompting and context management, whereas proactive alternatives that anticipate developer needs without explicit invocation remain underexplored. Understanding when humans are receptive to such proactive AI assistance during their daily work remains an open question in human-AI interaction research. We address this gap through a field study of proactive AI assistance in professional developer workflows. We present a five-day in-the-wild study with 15 developers who interacted with a *proactive* feature of an AI assistant integrated into a production-grade IDE that

offers code quality suggestions based on *in-IDE* developer activity. We examined 229 AI interventions across 5,732 interaction points to understand how proactive suggestions are received across workflow stages, how developers experience them, and their perceived impact. Our findings reveal systematic patterns in human receptivity to proactive suggestions: interventions at workflow boundaries (e.g., post-commit) achieved 52% engagement rates, while mid-task interventions (e.g., on declined edit) were dismissed 62% of the time. Notably, well-timed proactive suggestions required significantly less interpretation time than reactive suggestions (45.4s versus 101.4s, $W = 109.00$, $r = 0.533$, $p = 0.0016$), indicating enhanced cognitive alignment. This study provides actionable implications for designing proactive coding assistants, including how to time interventions, align them with developer context, and strike a balance between AI agency and user control in production IDEs.

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *IUI '26, Paphos, Cyprus*
 © 2026 Copyright held by the owner/author(s).
 ACM ISBN 979-8-4007-1984-4/26/03
<https://doi.org/10.1145/3742413.3789148>

CCS Concepts

• **Software and its engineering** → **Integrated and visual development environments**; • **Human-centered computing** → **Empirical studies in HCI**; • **Computing methodologies** → **Artificial intelligence**.

Keywords

Human-AI Interaction, Proactive Assistance, Intelligent User Interfaces, Software Development, Large Language Models, Field Study

ACM Reference Format:

Nadine Kuo, Agnia Sergejuk, Valerie Chen, and Maliheh Izadi. 2026. Developer Interaction Patterns with Proactive AI: A Five-Day Field Study. In *31st International Conference on Intelligent User Interfaces (IUI '26)*, March 23–26, 2026, Paphos, Cyprus. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3742413.3789148>

1 Introduction

Programming tools powered by large language models (LLMs) have changed how developers interact with code in IDEs [19, 29, 32, 60, 67]. These tools evolved from simple code completions [43, 61] to collaborative assistants that support various software development tasks [21, 26, 31]. Despite these advancements, existing LLM-based tools often face practical limitations when adopted in real-world developer workflows [7, 47, 59]. Most tools require explicit prompts from developers, creating cognitive overhead through the need to craft precise queries, supply context, interpret responses, and integrate suggestions [32, 45, 67]. These costs diminish productivity gains, especially in complex or urgent tasks, where frequent context switching between coding, prompting, and interpreting outputs interrupts developers' flow and increases mental fatigue.

To reduce such cognitive overhead, both researchers and practitioners have begun exploring proactive AI programming assistance, capable of anticipating developer needs and offering suggestions without explicit requests [13, 17, 21, 42, 52, 58, 68]. Yet, a key challenge persists: when prompts are poorly timed or perceived as intrusive, they risk undermining the user experience. Existing results highlight a critical open problem: how to accurately identify suitable moments for intervention that support, rather than disrupt, the user's workflow. Moreover, existing evaluations of proactive assistants have been confined to *controlled*, *web-based* environments with *predefined* tasks and *monolingual* settings [13, 52]. These limitations restrict our understanding of how proactive assistants perform under real-world conditions, where timing, contextual relevance, and code maintainability are critical concerns.

In this work, we investigate how proactive capabilities can be integrated into in-IDE assistants, advancing toward practical, context-aware developer support. We focus on providing proactive assistance to improve *code quality*, a particularly relevant developer concern that often goes unnoticed during development [71] despite its direct impact on code review efficiency, production defect rates, and deployment velocity [4, 48, 63]. While developers have shown strong interest in AI tools that help with code quality, including bug detection, refactoring, and adherence to best practices [12, 13, 62], many existing tools addressing code quality are offline or reactive [9, 30, 65]. Reactive tools provide limited usefulness in fast-paced development settings where timing and context

are critical [69, 70]. We aim to develop proactive assistance that is timely and non-obtrusive to help developers resolve problems before they might escalate.

To this end, we design, implement, and evaluate through a field study *ProAIDE*, a proactive feature of an AI assistant, that offers suggestions for code quality improvements. *ProAIDE* is integrated into *Fleet*¹, a production-grade, polyglot IDE with built-in AI capabilities used by professional developers, and owned by JetBrains², a major software vendor specializing in the creation of intelligent development tools. The development of *ProAIDE* followed an iterative, human-centered design process comprising four phases (Figure 1): in Phase 1, we built an initial prototype with three trigger types for proactive interventions based on established guidelines [1, 13, 24, 52]; in Phase 2, we conducted contextual inquiries with seven professional developers to refine the system's timing heuristics and interface design, extracting four design principles; in Phase 3, we refined the system based on these insights, resulting in the final *ProAIDE* implementation; and, in Phase 4, we deployed *ProAIDE* in a five-day in-the-wild study with 15 professional developers working on authentic development tasks spanning 12 programming languages. This integration enabled an in situ evaluation of proactive support in realistic workflows. Our analysis takes a mixed-methods approach, combining telemetry logs with structured daily and post-study surveys. This setup allowed us to analyze not only how developers interact with in-IDE proactive AI suggestions, but also how developers experience these interactions and the impact of proactive assistance on their work.

Our results offer a deeper understanding of how, when, and why proactive suggestions can support adoption and usefulness in real-world developer workflows. Across 5,732 interaction points, we found that suggestions triggered at natural workflow boundaries, particularly after commits, were significantly more likely to be accepted. On a System Usability Scale (SUS) [36], developers rated *ProAIDE* as easy to use (SUS = 72.8 out of 100, 95% CI [64.1, 81.5]) and appreciated its integration into the IDE, but varied in their preferences for timing, control, and granularity of suggestions. While alignment with task intent shaped trust and receptiveness, utility declined when the AI failed to capture lower-level contextual understanding of the developer's source code, including technical design choices or domain-specific patterns. Compared to reactive use of AI, proactive suggestions were interpreted significantly faster ($W = 109.00$, $r = 0.533$, $p = .002$), revealing the potential for enhanced workflow efficiency. Based on our empirically grounded findings, we present design implications for integrating proactive AI assistance with real-world IDEs to guide future systems toward more developer-aligned support.

Our contributions are as follows.

- **Empirical characterization of developer interaction patterns** with proactive AI suggestions across workflow contexts, revealing timing-dependent receptivity patterns in authentic professional environments.
- **Quantitative evidence of proactive assistance efficiency**, showing significantly reduced interpretation time compared to reactive AI use through deployment in production IDE.

¹Fleet: <https://www.jetbrains.com/fleet/>

²JetBrains: <https://www.jetbrains.com/>

- **Empirically grounded implications for building** proactive assistants: when to surface suggestions, what makes them effective, how trust and context influence adoption.
- We share our study material via a **replication package**³ [34].

2 Background

2.1 Developer-AI Interaction and Proactive Assistance

Understanding how developers interact with *proactive* AI assistants in the real-world IDE is lacking in current human-AI interaction research, compared to developer interactions with *reactive* systems. For example, Ross et al. [53] demonstrated the value of extended conversational interactions with AI systems for complex tasks, showing that multi-turn dialogues can enhance productivity compared to single-invocation approaches. Relatedly, Barke et al. [7] showed through a grounded theory that user interactions with AI can be classified into two modes akin to the two systems of thought in dual-process theories of cognition: acceleration and exploration. For interactions with reactive assistants, Mozannar et al. [46] suggested categorization of 12 programming states, including the most prominent time-wise verifying suggestions, prompt crafting, and editing code. In the recent work of Sarkar and Drosos [54], it was found that vibe coding with AI follows cycles: set a goal, prompt the AI, test results, refine, and sometimes edit manually. As these works primarily investigate *reactive* scenarios where humans explicitly seek AI assistance, they place the cognitive burden of recognizing assistance opportunities entirely on the developer. This interaction model introduces cognitive load and often limits real-time utility [14, 47, 49].

Proactive tools take a different approach. These systems anticipate developer needs and offer assistance without the need for explicit prompting. However, a fundamental challenge persists: determining when AI systems should *proactively* intervene without disrupting developers' workflow and cognitive state [1, 24]. The timing challenge is particularly critical in real-world professional developers' workflow because poorly timed interventions can disrupt focused work and reduce trust, while well-timed assistance can enhance workflow efficiency without cognitive overhead [5, 27].

While research on more proactive assistants has been studied in forms ranging from physical robots [6, 50] to virtual chat assistants [37, 38], the actual capabilities of prior deployed proactive assistants do not always meet user expectations [39]. Moreover, the effectiveness of proactive AI in professional software development contexts remains underexplored. Only some research investigated how to improve the timing and scope of proactive code completions [18, 44] and several works explored proactive chat-based agents for professional coding [13, 74]. As many of the previous systems relied on simple models or preset messaging, with the advent of modern LLMs and their growing usage in agents for software engineering [72], there is significant potential to design more capable proactive coding assistants.

Furthermore, a critical limitation of existing proactive AI research in software development is the reliance on controlled evaluation environments. Most of these systems were evaluated in toy

editors, controlled settings, or Python-only environments [13, 52]. These design considerations limit what we know about how proactive assistance behaves in full-featured IDEs with realistic workflows, diverse programming languages, and authentic time pressures. Building on prior work, we integrate proactive support into a production-ready in-IDE AI assistant and conduct a five-day in-the-wild deployment with professional developers. Through this deployment, we capture developer interactions with proactive suggestions triggered at three AI-driven workflow-grounded points [10, 11]: ambiguous prompt detection, declined AI edit follow-up, and post-commit review.

2.2 Code Quality as a Domain for Proactive Assistance

Developers increasingly expect AI tools to support not just their productivity (e.g., how quickly they can write code), but also make meaningful edits to existing code to improve the code quality [8, 57]. Poor code quality has direct business consequences: it increases time spent in code review [4], leads to higher defect rates in production [48], and reduces deployment frequency [63]. Despite these impacts, code quality issues frequently escape detection during initial development [71], making them ideal candidates for proactive AI intervention. While recent work shows that LLMs can help with refactoring [51], bug fixing [9, 65], review [15], optimization [20], and smell detection [69], the aforementioned tools are largely reactive. Developers must spot a problem, switch tasks, and phrase a suitable prompt, which are multiple barriers that can limit adoption.

Proactive support can reduce the friction of reactive tools by offering timely, relevant improvements without requiring initiation. Moreover, code quality suggestions often involve changes that do not affect program behavior, such as style improvements or small refactorings. These suggestions can still be helpful even when they are not directly related to the developer's current task. Developers also appreciate support that enhances clarity and maintainability, as long as it is not disruptive [13, 52]. Our work builds on these observations by exploring how proactive AI can surface code quality improvements in real workflows. Note that we do not focus on building new quality assessment tools and instead study how such improvements can be integrated into natural developer workflows.

3 Proactive System Design and Iterative Refinement

To investigate how professional developers interact with proactive AI assistance across workflow contexts in situ, we created a prototype of a feature that provided proactive suggestions for code quality improvements inside a production in-IDE AI assistant. We developed ProAIDE through an iterative, human-centered design process spanning three phases (Figure 1). First, we created an initial prototype with three workflow-grounded intervention triggers based on prior work (Section 3.1). Second, we conducted contextual inquiries with seven professional developers to refine the system and extract four design principles (Section 3.2). Third, we refined the prototype based on these insights, resulting in the final ProAIDE implementation (Section 3.3).

³We share recruitment texts, survey instruments, and contextual inquiry protocol. Note that the collected user data cannot be made public due to user privacy considerations.

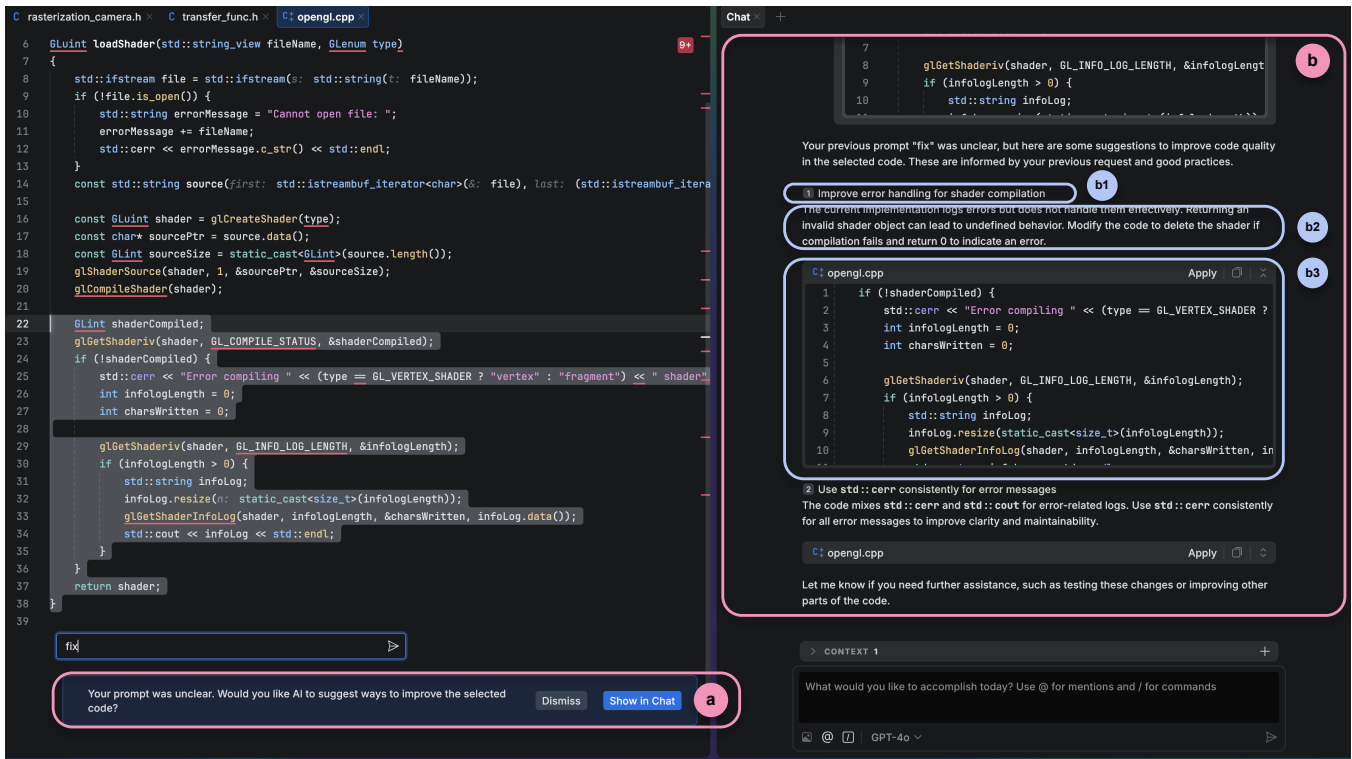


Figure 2: User interface for proactive intervention at an *ambiguous prompt*. (a) Lightweight in-editor cue indicating AI activity. (b) Opened chat panel with: (b1) suggestion title, (b2) rationale, and (b3) code snippet. The UI for *declined AI edit* followed the same structure. For *post-commit*, the interaction interface remained, but the visual cue (a) appeared as a popup in the upper-right corner after a successful commit.

3.1 Initial Prototype

For the initial user interface, the in-IDE’s AI chat served as the central communication channel (see Figure 2), specifically designed to minimize context switching while preserving developer focus and workflow continuity based on guidelines from prior work [13, 52, 64, 66]. When the system detected intervention opportunities, which are described below, it displayed a minimal inline banner providing options to dismiss or engage with the suggestion. Upon engagement, the chat panel presented suggestions with associated code patches. For automatically integrating suggestions, diffs were directly displayed in the editor for review. Developers could subsequently accept or reject changes. The initial interface remained virtually unchanged in the final system, with the exception of detailed specifications and interaction flows described in Section 3.3.

Unlike prior work, our in-IDE assistant is present throughout the entire developer workflow. In particular, the system triggers proactive suggestions across three conceptual stages of the code improvement process, depicted in Figure 3. These trigger points were informed by prior studies on developer workflow patterns [10, 11] and were refined through contextual inquiries:

- **Formulating Need:** When developers submitted a prompt through the IDE’s built-in inline AI action “Edit Code”, the system evaluated this request by making LLM-as-judge assessments against predefined criteria to identify *ambiguous*

prompts, such as lack of context (e.g., requesting a fix without providing error details) or inconsistent references (e.g., using undefined variables). Upon detecting such ambiguity, the system intervened to offer additional guidance and code quality targeted suggestions to enhance their code further. This heuristic is motivated by prior work showing that accurately formulating specifications is non-trivial, imposes substantial cognitive overhead, and has a significant impact on both the efficiency and quality of generated code [49, 73].

- **Executing Idea:** Developers *declining AI code edits* was treated as a heuristic signal that the proposed solution may not have fully aligned with their intent. To explore whether follow-up assistance could be valuable, ProAIDE proactively offered suggestions in such cases with focus on code quality. These could include the reasoning behind the original solution or alternative implementation approaches, aiming to reduce the need for developers to reformulate requests after an unsuccessful attempt.
- **Getting Feedback on Written Code:** the system also proactively provided suggestions after developers *committed their changes*, representing a natural checkpoint [28, 52]. Committing code generally indicates that a developer has finished iterating on a logical unit of work, potentially creating a

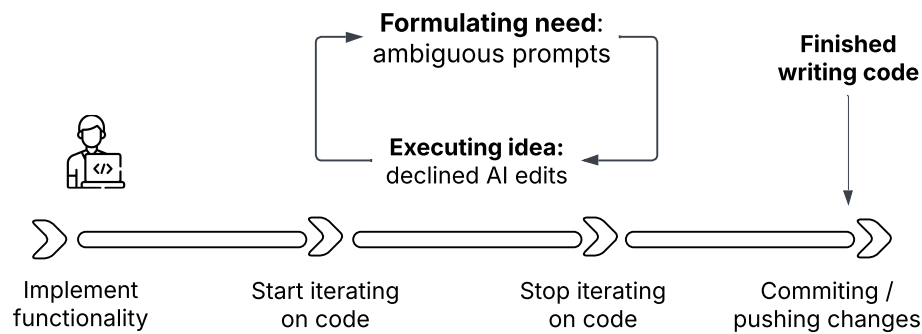


Figure 3: Main stages of the code iteration workflow and corresponding points where ProAIDE triggers proactive suggestions.

receptive moment for review. Hence, upon successful commits, developers were proactively presented with an option to “Review Changes with AI” which invoked the AI chat where they received feedback on best practices and code quality in terms of criteria, including modularity, efficiency, and robustness.

Due to proprietary constraints, complete system prompts could not be disclosed. However, relevant prompt segments are provided in supplementary materials [34].

3.2 Contextual Inquiries

We conducted contextual inquiries [23] with seven professional developers at JetBrains to refine approach described above.

3.2.1 Procedure. Participants were chosen to be professional developers (more than 1 year of experience), regular (at least once per week) users of the IDE under study, who frequently (daily) used in-IDE AI features. Each remote session lasted approximately one hour and was both recorded and transcribed. Participants interacted with the early prototype described in Section 3.1, addressing code quality issues in any familiar project of their choosing. The setting characterized their daily work, with participants elaborating on their thought processes aloud. We concluded each session with a semi-structured interview to gather feedback about participants’ impressions, preferences, and suggestions for improvement. The script and transcripts of contextual inquiries are included in the supplementary materials [34]. Recordings were transcribed, and through collaborative discussion, the first two authors extracted prominent and common themes to inform the refinement of the early prototype.

3.2.2 Analysis of Themes. Combining these insights with existing literature, we identify four Design Principles (DPs) that guided the refinement for our resulting implementation of ProAIDE.

DP1. Timely, proactive AI assistance emphasizes anticipating developer needs without causing disruptions. In prior work, periodic suggestions had been shown to interrupt workflow and reduce focus [13, 74]. To address this, our initial prototype triggered suggestions at two code iteration stages (Figure 3): formulating needs and executing their ideas. However, participants of the contextual inquiries recommended including triggers at task boundaries later

in their workflow, such as after commits. One participant noted: “I separate writing code and optimizing it. When I optimize, I want to do it later. Reviewing should happen after I’ve written something and feel it’s ready for review.” [Pilot2]. We returned to this distinction between task stages and boundary-triggered interventions in our analysis of how developers engaged with proactive suggestions across workflow phases (see Section 5).

DP2. Contextually relevant suggestions are important for ensuring that the assistant supports, rather than interrupts, existing in-IDE workflows. Instead of general-purpose proactive suggestions [13, 74], we specifically targeted code-quality improvements. These improvements spanned optimizations based on best practices and general code quality criteria, tailored to the developer’s context. Whereas the LLM underlying the AI assistant by default had access to tools to dynamically fetch code artifacts across the repository, we augmented its context with potential problems detected by the IDE (e.g. language servers) and the relevant portions of code the user was actively working on. In case of interventions on commits, the corresponding git diffs were provided. Participants of the contextual inquiries appreciated that ProAIDE inferred their intent without requiring explicit prompts or manual context augmentation. However, they occasionally perceived suggestions as redundant, overly generic or not aligned with their initial intent.

DP3. Explainability and transparency of proactive suggestions are crucial to build developer trust, particularly in industrial settings. Earlier work had highlighted how over-reliance and subsequent limited understanding of proactive AI suggestions could lead to long-term maintenance and extensibility issues [52]. Addressing this, our design included clear explanations of suggestions and rationales behind AI intervention. Additionally, an in-chat context panel provided transparency into the AI’s working context (e.g. relevant files, commits or other manual attachments from the user). Participants valued the ability to preview changes in the editor when receiving patches, but requested inline diffs within snippets themselves, without them being integrated directly. Suggestions felt more trustworthy when users could inspect edits before applying them: “Maybe smaller suggestions or steps would help. Right now it feels too risky to accept all at once.” [Pilot3]. Hence, offering multiple small patches was preferred over single large edits.

DP4. Preserving user control balances proactive assistance with the developer’s autonomy in decision-making. From the contextual inquiries, it became evident that automatic invocation of the AI chat upon intervention points was seen as intrusive. Developers preferred lightweight, in-editor pop-ups to confirm chat invocation, noting that suggestion rejection often stemmed from task switching rather than disapproval. We found that signaling AI activity close to the user’s current interaction scope within the IDE was necessary to minimize distraction. One participant commented: “I would expect some confirmation before running a new request to the chat, otherwise it would be quite unexpected.” [Pilot5].

3.3 ProAIDE

Based on findings from contextual inquiries with the early prototype, we refined ProAIDE to address the identified issues as described above (see Figure 2 for the resulting user interface):

- When intervening on *ambiguous prompts* and *declined AI edits*, we additionally augmented the assistant’s context with the corresponding user’s request, to enhance alignment with their initial intent. We also refined the system prompt to avoid suggestions that may be repetitive given previous attempts (as inferred from the user request or code context).
- We implemented lightweight visual cues (see Figure 2) with optional expansion directly in the IDE near the user’s current cursor position (a), providing users with the option to invoke the chat panel (b), where the AI suggestions would be displayed. Each suggestion contained a concise title (b1), a brief explanation of the issue and the rationale behind the improvement (b2), and a relevant code snippet (b3).
- We added an explicit option to apply suggested patches from within the AI chat, rather than directly displaying inline diffs within the editor. This on-demand model gave developers full control over the integration of suggestions while avoiding unnecessary computational cost for LLM requests in the background to apply these patches.

Our iterative design process resulted in an enhanced version of ProAIDE to be used in our systematic in-the-wild study of developer interactions with proactive AI assistance across authentic workflow contexts, which we describe next.

4 In-the-wild Study Design

To investigate developer interaction patterns with proactive AI assistance in real-world professional contexts, we conducted an in-the-wild study that employed a mixed-methods approach, combining quantitative telemetry data with qualitative survey responses. The study was conducted in line with institutional ethical standards, adhering to the values and guidelines outlined in the ICC/ESOMAR International Code [25]. The research questions we aim to address are as follows:

- **RQ1.** How do developers *interact* with in-IDE proactive AI suggestions across workflow stages?
- **RQ2.** How do developers *experience* the interaction with in-IDE proactive AI suggestions?
- **RQ3.** How do developers *perceive* the impact of in-IDE proactive AI suggestions on their work?

Table 1: Participant Demographics

ID	Role	Languages	Exp.	AI Use Freq.
P1	Dev, DevOps	C#, Groovy, Py	6–10y	Multi/day
P2	Dev	Kt, TS	11–16y	1/week
P3	Dev, Researcher	Dart, Py, Shell	3–5y	Multi/day
P4	Dev	C#	16+y	Multi/day
P5	Dev, DevOps	Dart, Kt, Py	3–5y	Few/week
P6	Dev	Java, Kt, SQL	11–16y	Multi/day
P7	Dev, Analyst	Py, SQL, TS	6–10y	Multi/day
P8	Dev	Go, JS	11–16y	Multi/day
P9	Architect	Go, Kt, Py	16+y	Multi/day
P10	Dev	HTML, JS, Kt	1–2y	1/day
P11	Dev	C#, TS	3–5y	Multi/day
P12	Dev	Dart, Rust, Swift	6–10y	Multi/day
P13	Dev, ML Eng	Py	6–10y	Multi/day
P14	Dev, Architect	Go, PHP, TS	16+y	Multi/day
P15	Dev, Architect	Kt, Py	11–16y	Multi/day

4.1 Participants

We recruited participants through JetBrains’ research panel using a comprehensive screening survey, requiring active use of the Fleet IDE and the JetBrains AI assistant ⁴ under study, regular AI tool usage (at least weekly), and professional development experience. The screening process yielded 29 qualified participants who met all selection criteria. According to the daily and post-study surveys, 15 participants were actively engaging with the tool. The number of participants is in line with prior studies of proactive AI [52] and in-the-wild studies of software developers [35, 40, 41]. All participants indicated familiarity with JetBrains’ products *and* in-IDE AI tools, minimizing potential confounding variables. Participants were required to have regular experience with AI tools in day-to-day development work. Out of the participants, 80% indicated using AI tools multiple times daily, with 66.7% using GitHub Copilot and 46.7% using JetBrains AI Assistant. The participant pool consisted of 93.3% identifying as Developer/ Programmer/ Software Engineer roles, and 33.3% as DevOps/ Architects. Regarding professional experience, 46.7% indicated having over 11 years of coding experience and 26.7% having 6-10 years of experience. In terms of primary programming languages used by the participants, Python (46.7%), Kotlin (40%), and TypeScript (26.7%) were the most common. Full demographic details are in Table 1.

4.2 Data collection procedure

The study followed a structured timeline spanning four weeks from 7 May 2025 to 3 June 2025. Participants were recruited on a rolling basis and enrolled in batches. Eligible participants, identified through a screening survey, received an invitation to start the study in the following week. The core study period required five active days of Fleet IDE use. This duration balanced capturing interactions with ProAIDE beyond initial novelty [33, 40] with the practical constraints of professional developers. To further accommodate participant schedules, we allowed a ten-day window for

⁴JetBrains 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.

completing these five active days. Participants received a reminder midway through the study and a final survey prompt at the end of the participation window. This design provided flexibility for participants while maintaining structured, consistent data collection. Supplementary materials [34] provide detailed texts used for that.

4.2.1 Capturing in-the-wild usage and daily surveys. During the study, participants engaged in self-directed programming tasks representative of their routine professional activities (*i.e.* “worked as usual”). There were no constraints on programming languages, frameworks, or project types for this reason. With participants’ informed consent, we collected anonymous IDE usage statistics as telemetry data. Importantly, we did not collect personal data, source code, file names, or full inputs and outputs exchanged between the IDE and LLMs. To prevent fraudulent participation, in-IDE notifications prompted participants to complete short daily surveys. These surveys appeared five minutes after interactions with any proactive feature, limited to once per IDE session. Surveys captured the specific proactive feature used, the relevant development task, and perceptions of feature behavior, along with an open-ended question for additional feedback. Daily surveys also helped track participation, as telemetry data were anonymous and thus not linked to individual participants.

4.2.2 Post-study survey. After completing at least five active days and corresponding daily surveys, participants filled out a post-study survey. This survey included Likert-scale items from the System Usability Scale (SUS) [36]. Additional questions evaluated participants’ perceptions of the assistant’s reliability, intent alignment, intrusiveness, and impact on efficiency. Two open-ended questions encouraged participants to describe specific effects on code quality and suggest improvements to better support their workflow.

4.3 Data Analysis

We applied both quantitative and qualitative analysis methods to address our research questions. To examine developer **interactions** with the proactive feature of AI assistant across workflow stages 3.1 (**RQ1**), we analyzed 5,732 data points and 229 interventions in telemetry logs. Specifically, we focused on three interaction types for each AI intervention: engage (user opted for viewing the suggestions in chat), dismiss (user actively closed the notification), or ignore (no user response). We processed raw telemetry events by filtering for proactive intervention triggers, then categorized subsequent user actions following each intervention. Descriptive statistics, including engagement patterns and temporal patterns across the five-day study period, were computed to characterize interaction behaviors for each intervention type (ambiguous prompt, declined AI edit, commit changes). To evaluate developer **perceptions** of the assistant (**RQ2**) and its perceived **impact** (**RQ3**), we mainly analyzed post-study survey responses. This included quantitative data summarized through frequency distributions for Likert-scale items, and qualitative insights from open-ended questions. Moreover, we employed a Wilcoxon signed-rank test [56] to compare interpretation times (from the moment a suggested snippet appeared in the chat interface until the user’s first response to it) between proactive and reactive chat sessions. Given the moderate sample sizes and the observed distributional

skewness inconsistent with normality assumptions in both conditions, the Wilcoxon signed-rank test provided robust inference without requiring the parametric assumptions necessary for t-tests. We combine the telemetry and survey data to derive contextualized insights into the usability and perceived usefulness of ProAIDE.⁵

5 Results

Our analysis combines in-IDE telemetry (N=18 machines, with active usage declining to 8 machines by day five) and a post-study questionnaire with open-ended feedback (N=15).

5.1 RQ1: Interaction Patterns Across Workflow

To answer the **RQ1**, we analyzed 229 interactions (Figure 4) across three workflow-triggered intervention types: (1) *Ambiguous Prompt*, (2) *Declined AI Edit*, and (3) *Commit Changes* over the study period of five days. Each intervention provided developers with options to either engage with the AI for code improvement suggestions, dismiss the notification, or ignore it entirely.

- *Ambiguous Prompt* interventions (N=35) saw moderate engagement (46%), with 23% dismissals and 31% ignores. This balanced distribution may reflect a mix of situations: in some cases, developers appeared to be in an exploratory mode, open to proactive suggestions. In others, they seemed to have a clear intent in mind but preferred efficiency over detail, being surprised that “*simplify was insufficient context for the AI to act upon*” [P1].
- *Declined AI Edit* interventions (N=39) achieved the lowest engagement rates at 31% and highest dismissal rates at 62%. Despite signaling unmet needs, these interventions appear to feel “*like advertisement*” [P3], indicating unnecessary intrusion when developers have already formulated specific task intentions. Moreover, context-switching between different AI interaction modalities appears to be causing cognitive overload, with participants expressing confusion about “*when to use inline AI and chat*” [P15].
- *Commit Changes* interventions (N=155) achieved highest engagement rates (52%). This timing is based on natural workflow boundary points when developers complete logical units of work and are “*still in the mindset of wanting to submit good code*” [P5]. This suggests that developers, having finished writing code and entered an evaluative mindset, experience reduced cognitive switching costs and thus are more receptive to proactive AI.

An analysis of interaction patterns over the study period revealed stability across all intervention types (See Fig. 4). For *Commit Changes* interventions, the consistent engagement rates throughout the study period indicate that the utility of boundary-point interventions is not diminished by repeated exposure. The daily survey feedback confirmed that having AI review changes after commits remains helpful, with users expressing a desire to “*continue using it once the feature becomes available*” [P5]. In contrast, engagement rates for *Declined AI Edit* interventions remained consistently low, suggesting that the poor reception was not necessarily due to an

⁵Survey data supporting the findings of this study may be shared upon reasonable request. Due to privacy and confidentiality concerns, telemetry data collected from in-IDE usage cannot be made publicly available.




Figure 4: Rates of interaction with proactive AI interventions across various development stages: formulating needs (Ambiguous Prompt), executing the idea (Declined Edit) and finishing writing code (Commit Changes). Users had the option to Engage, Dismiss, or Ignore AI interventions. Active participation decreased from 18 to 8 users between Day 1 and Day 5, yet interaction patterns remained stable across intervention types. Commit Changes interventions achieve the highest engagement (52%), while Declined Edit interventions show consistently low engagement (31%). In total, we recorded 229 proactive AI interventions over the 5-day study period.

initial learning curve, but rather reflects a more fundamental disruption to their workflow. Moreover, 8 out of 18 participants chose to continue using ProAIDE beyond the mandatory 5-day study period, suggesting that its value extended beyond initial curiosity. This sustained engagement indicates that the proactive feature was not merely perceived as a novel experimental tool, but had begun to integrate meaningfully into participants’ regular workflows.

Summary of RQ1. *Proactive AI’s utility depends on developers’ cognitive context and thus readiness to receive assistance. The most accepted interventions occur at natural workflow boundaries when developers are in the mindset of reviewing submitted code, while interventions during focused implementation tasks often feel disruptive and intrusive.*

5.2 RQ2: User Experience

Having identified when proactive AI interventions are most interacted with, we now examine how developers experienced these interventions (RQ2). The findings below are primarily based on analysis of qualitative post-study survey responses (Figure 5).

5.2.1 ProAIDE was Easy to Use and Learn. Participants generally rated ProAIDE positively in terms of usability, reflected by an above-average SUS score of 72.8 out of 100 (95% CI: [64.1, 81.5]). This indicates good usability, though with room for improvement, suggesting users hold favorable but not strongly promotional views [16, 55]. The majority found the system quick to learn (93%), likely due to its seamless integration within the already familiar IDE environment. Additionally, participants considered the proactive feature easy to use (87%). Explicit confirmation mechanisms were particularly appreciated, with one participant noting the system was “*even better than some other IDEs [...] offering custom UI in form of buttons when asking for confirmation*” [P12]. Furthermore, the transparency in AI

presence contributed to helping developers feel like they “*always knew what the suggestions were talking about*” [P3].

5.2.2 Developers Had Varying Needs and Preferences. In terms of how proactive interventions were triggered, many expressed a desire for control, as [P5] noted: “*I prefer to use hotkeys where possible or a command that can be used*”. Conversely, some preferred greater AI autonomy, with another participant stating “[*They*] *would like to immediately move context to ‘Chat’ without even asking me, since [they] didn’t accomplish what [they] tried to do*” [P15]. Moreover, opinions regarding the degree of proactiveness and granularity of suggestions differed. Some participants indicated that “*it would be good if the ‘Review Commit’ option could be triggered automatically on any commit*” [P5], or anchored to “*any changes made in git*” [P10]. Others found intervention frequency excessive, stating that suggestions were “*not always necessary [...] dismissing the dialog is annoying*” [P15]. Regarding suggestion granularity, some participants sought detailed, “*fine-grained comments in code and documentation*” [P3], while others preferred concise suggestions, stating “*no need for explanation text*” [P15]. Also, participants also wished to select suggestion types explicitly, such as those related to “*performance, readability, security*” [P4].

5.2.3 Alignment Between AI and Mental Model Was Key. An important factor influencing developers’ receptiveness to proactive suggestions was how well the AI’s behavior aligned with their expectations and mental models. Survey responses indicated that 66% of participants easily understood how the suggestions related to their intended tasks. One participant noted, “*In most cases, suggestions were exactly what I wanted*” [P6]. Interestingly, one participant emphasized the role of developer context, stating, “*I did not feel that any suggestion was misleading [...] so perhaps the coder’s perception and friction with the code they are writing is an important factor here.*” This suggests that perceived alignment may depend not only on the AI but also on the developer’s experience with the task at hand.




Figure 5: System Usability Scale and user experience responses from 15 participants who completed the entire study period. Results show above-average usability (SUS = 72.8).

5.2.4 *Perceived Utility was Shaped by the AI’s Contextual Understanding.* The usefulness of proactive AI suggestions was strongly influenced by the AI’s ability to understand the developer’s code context. However, only 27% of participants rated AI chat suggestions as reliable, and just 47% felt comfortable accepting them. In fact, the perceived reliability diminished when intent behind design choices was not understood; in some cases “repetitiveness was necessary and unavoidable for the specific functionality intended to achieve” [P3]. Others experienced frustration when the “AI failed to account for domain-specific syntax and patterns” [P13]. We note that LLM training data here plays a substantial role, as P3 observed: “at least half of the times this happened [...] other LLM chats [...] would get it wrong in their first try as well.” Another described the effort required to manually adjust inaccurate suggestions: “The suggestions were sort of helpful, but messed up the code file. I manually copied suggestions, undid changes, then pasted suggestions into the correct place” [P4], disrupting their workflow.

5.2.5 *Specific Friction Points.* Beyond contextual understanding, session length also affected reliability. Participant [P7] observed that “if a particular chat session gets long, then the model would start glitching out”, entering loops where it would “list out the entire plan as is and not actually implement the plan”. This degradation over extended interactions forced developers to start fresh sessions, disrupting workflow continuity. Finally, some participants noted that commit review suggestions sometimes restated changes rather than critiquing them. Participant [P13] provided feedback that

“instead of reviewing my commit, it summarizes the changes I made and presents them as its ‘suggestions’ ”.

Summary of RQ2. Developers rated the proactive feature as easy to use (SUS = 72.8), expressing ease of learning and good integration. However, preferences varied widely regarding user control, degree of proactivity, and granularity of suggestions. Perceived usefulness was influenced by how well the AI’s suggestions aligned with developers’ mental models and workspace context.

5.3 RQ3: Impact of Proactive Suggestions

While RQ2 focused on the experience of using ProAIDE, we now turn to RQ3, which investigates how its suggestions influenced developers’ perception of impact on their work. The study reveals that proactive AI assistance generates overall positive effects on perceived productivity. We find that the majority of participants indicated that using the proactive suggestions helped them complete tasks more quickly (66%) and easily (60%), as seen in Figure 5. Interestingly, from qualitative responses, we find that developers engaged with ProAIDE for various underlying purposes.

5.3.1 *The Proactive AI Served Exploratory Purposes.* Our ProAIDE served as a code quality gate, with participants appreciating its role as a second pair of eyes for identifying overlooked issues or general enhancements. P3 noted: “There were a few suggestions around safety ... an aspect that actually is really important in order




Figure 6: Interpretation time comparison between proactively triggered chat sessions ($\mu = 45.4s$, $N = 33$) and reactively triggered ones ($\mu = 101.4s$, $N = 50$). Proactive suggestions demonstrate significantly faster interpretation times ($p = 0.0016$, Wilcoxon signed-rank test), suggesting improved efficiency when AI assistance is tailored to the user’s workflow context.

to catch senseless errors when speed coding and did improve my code significantly.” Similarly, P5 mentioned how “[They] didn’t catch that loading the asset might throw an error as [they were] following an implementation guide from the library docs.” Participants also seemed to get an intuition for the AI’s strengths over time: “the AI is helpful for identifying repetitive code and extracting it to improve modularity. It also handles typing well.” [P13].

5.3.2 The Proactive AI Served Accelerative Purposes. Beyond exploratory purposes, we observed that ProAIDE served accelerative purposes by alleviating the manual burden of context management and switching between different AI modalities, especially on *Declined AI Edit* interventions. While developers generally rejected proactive suggestions after declining AI edits, some participants appreciated that the system automatically escalated to the AI chat with relevant context (e.g. when inline AI failed to produce the desired result), “offering proactive code editing while [they were] writing code so that [they wouldn’t] have to open a chat window or go somewhere else to fill in what’s missing” [P12].

Interestingly, users spent significantly less time interpreting suggestions in proactively-triggered chat sessions ($\mu = 45.4s$, $N = 33$) compared to reactively-triggered sessions ($\mu = 101.4s$, $N = 50$), as shown in Figure 6.

We compared interpretation times (*i.e.*, the time between the moment a suggested patch or snippet was presented in the chat until the user responded using either “Accept” or “Copy”) between proactive and reactive chat sessions using data from the same set of 18 participants. Given the paired nature of the data and the observed distributional skewness (see Figure 6), we employed the Wilcoxon signed-rank test as a non-parametric alternative to the paired *t*-test. To address the unequal group sizes (50 reactive vs. 33 proactive sessions), we created balanced subsets by randomly sampling 33 observations from each condition. The test revealed a statistically significant difference ($W = 109.00$, $p = 0.0016$) with

a medium-to-large effect size ($r = 0.533$). These findings suggest that proactively delivered suggestions, whether accepted as patches or adapted as snippets, were substantially easier and faster for developers to interpret than reactively requested ones.

Examining the types of suggestions reveals differences in user engagement. In the reactive condition, users encountered relatively more exemplary code snippets (21/50; 42%) than in the proactive condition (5/33; 15%), where patches dominated. This discrepancy suggests that reactive use was more information-seeking, while proactive suggestions provided more actionable code. This distinction likely contributed to the reduced interpretation time observed for proactive suggestions, as users were presented with ready-to-apply patches that generated diffs directly in the editor. In contrast, exemplary snippets were displayed only as in-chat code fragments, likely requiring more effort to interpret and adapt.

Summary of RQ3. Developers engage with proactive AI for both exploratory and accelerative purposes, achieving high perceived productivity and workflow efficiency. In fact, we observe significantly lower interpretation times for proactive suggestions compared to reactive AI use.

6 Discussion

6.1 Design Implications for in-IDE Proactive AI

Our study provides empirical evidence for how developers interact with proactive AI assistance in authentic professional workflows, revealing timing-dependent patterns that inform the design of context-aware coding assistants. While prior work explored proactivity in controlled environments and limited contexts, our in-the-wild study reveals how real-world developers engage with and evaluate proactive interventions over time, spanning various programming languages and frameworks. We also revisit findings from earlier studies, contrasting them with our observations based on real-world behavior. Building on these insights, we offer design implications and recommendations grounded in the four design principles introduced in Section 3.

DP1. Timely, proactive AI assistance. Our findings reveal that timing is a crucial aspect impacting developer engagement with proactive AI. Unlike static recommendation systems, LLM-based proactive assistants must dynamically align their understanding with developer mental models through contextual reasoning. Workflow boundary interventions (e.g., after commits) were consistently well-received, due to natural blending into developer workflows. This reinforces prior findings from Pu et al. [52], now ecologically validated in a production environment. On the other hand, mid-task interventions (e.g., after declined edits) were frequently dismissed or ignored. This pattern aligns with Iqbal and Bailey’s Index of Opportunity framework [27], which demonstrates that mental workload decreases at task boundaries, creating natural windows for interruption. Post-commit moments represent such boundaries: developers have externalized their working memory into the commit, freeing attentional resources for evaluating new information. Mid-task interventions, by contrast, compete for cognitive resources engaged in active implementation. This theoretical lens suggests that receptivity to proactive AI depends on the underlying cognitive state that different workflow moments reflect.

Implication and Recommendation: Developers' workflow context and cognitive state must guide when interventions are triggered to ensure effectiveness. As such, anchor proactive suggestions to clear task boundaries (e.g., failed builds or test runs). For mid-task interventions, we recommend tuning down the AI's prominence in the IDE or providing configuration options to respect individual workflows.⁶ Overall, we highlight the need for more accurate intent modeling to adapt triggers dynamically based on developer behavior and context. Reinforcement learning (RL), for instance, can be used to learn optimal interruption moments to balance relevance and flow disruption [22]. Our study contributes the empirical foundation such approaches require: the engagement and dismissal outcomes we observed provide ground truth labels for developer receptivity, which could serve as reward signals in future RL formulations. On the other hand, in-IDE telemetry events offer candidate state representations. However, learning such policies requires larger-scale deployments with appropriate privacy safeguards, which we leave to future work.

DP2. Contextually relevant suggestions. Developers valued suggestions that reflected their current task, responding positively when suggestions felt well-aligned. Moreover, our findings support Chen et al. [13]'s framing of proactive AI as a "second pair of eyes" to catch overlooked issues. Nonetheless, we identify that the perceived value of suggestions diminished when they were perceived as irrelevant, for instance, due to unawareness of domain-specific patterns, which can be attributed to both context integration and LLM training data.

Implication and Recommendation: Suggestion quality is largely influenced by the AI's understanding of developer intent and their workspace, highlighting the importance of enhanced integration with workspace context. As such, adopt more effective heuristics for integrating in-IDE context sources, prioritizing signals that reflect the developer's immediate focus i.e. "possibly using nearby open tabs or recently edited files" [P3], ideally grounding these signals in a broader semantic representation of the codebase through repository indexing. Notably, recent standardized protocols for connecting LLMs to external data sources such as Anthropic's MCP [3] exemplify a shift toward structured, tool-mediated context retrieval enabling real-time personalized augmentation of model context.

DP3. Explainability and transparency. Transparency in the AI's in-IDE presence was essential for establishing reliability. Developers needed to understand not just what the AI suggested, but why it triggered at that moment and how it interpreted the code context. This differs from traditional recommendation systems, where the underlying logic is more predictable. Developers appreciated concise rationales, previewable changes, and minimal in-editor cues that respected their workflow. Notably, while Pu et al. [52] found that flexible interaction scopes improved transparency, our findings show they can cause distraction due to context switching. This highlights the need for careful coordination of concurrent AI features in real-world IDEs to reduce cognitive load. To further enhance trust, participants indicated that it "would be helpful if each suggestion came with a confidence level".

⁶Examples of such customization mechanisms include custom slash commands and hooks in the CLI-based AI coding tool Claude Code [2].

Implication and Recommendation: Proactive suggestions must be understandable at a glance, predictable in behavior, and verifiable before application. As such, enhance developer decision-making by increasing transparency, e.g., showing clear inline diffs when presenting patches (i.e. before application onto the editor) or confidence indicators with suggestions.

DP4. Preserving user control. Control over when and how to engage with AI suggestions was a recurring theme. In fact, preferences varied across individuals; some preferred more autonomy, while others favored greater manual control via hotkeys or menu actions. Beyond interaction modalities, users expressed a desire to have control over the type and granularity of suggestions.

Implication and Recommendation: To accommodate diverse user workflows and preferences, adaptive human-AI interaction strategies are needed. As such, combine user-configurable settings for intervention timing, frequency, and suggestion types with adaptive assistance that learns from both explicit feedback and implicit behavioral signals to provide personalized support.

6.2 Threats to Validity

While our findings provide strong real-world insights into proactive AI use, several limitations must be acknowledged.

External Validity. Participants were recruited from JetBrains' research panel and were familiar with its IDEs and AI tools, which allowed us to study ecologically valid use cases. However, the findings may not generalize to developers unfamiliar with these tools or those with limited AI experience. While the findings likely generalize to similar in-IDE AI assistance tools, future studies should broaden the participant pool to assess representativeness across diverse contexts. Due to the limited sample size, we did not stratify participants by role, experience level, or language. Some findings in our study may be shaped by language-specific workflows or by domain-specific factors. For languages with limited tool support or different compilation models (e.g., C++), tailored designs or fallback strategies may be necessary to achieve similar outcomes. Future studies should examine whether timing-dependent receptivity patterns vary systematically across developer categories.

Construct Validity. Participants interacted with a full-featured AI assistant, of which our proactive feature was one component. As such, post-study perceptions may have been shaped by the assistant as a whole. We see this as a strength of ecological integration, yet it introduces ambiguity in interpreting usability scores. Still, qualitative data confirms that the proactive feature contributed to the observed user experience, making these responses relevant. Controlled A/B testing may complement in-the-wild studies by isolating specific feature effects. Our assessment of proactive suggestions' impact relies on self-report rather than objective metrics. Preserving participant privacy in enterprise settings precluded collecting source code or LLM outputs, preventing direct measurement of code quality improvements. Future work could address this through controlled studies with open-source repositories where code artifacts and peer review outcomes can be analyzed, or by integrating privacy-preserving quality signals into telemetry with participant consent.

Internal Validity. Participants were exposed to all intervention types, but engagement varied by individual workflow, IDE usage intensity, and day-to-day conditions. As telemetry data were collected anonymized and separately from surveys, we could not link behavioral data to subjective perceptions at the individual level. While both datasets are independently reliable, they cannot be precisely mapped to the same individuals. Moreover, telemetry includes data from all participants who engaged with the system, even if they did not complete the full five-day protocol (N=18 machines with active usage declining to 8 machines by day five). On the contrary, survey analyses were restricted to participants who met the five-day filled daily survey protocol criteria (N=15). This preserves privacy and maximizes available data but may have introduced noise.

Statistical Conclusion Validity. The sample size, variability in daily engagement, and non-uniform exposure to intervention types limit the power of inferential statistical tests. Where applicable, we use non-parametric tests to account for data specifics. While our findings point to meaningful patterns, results should be interpreted with caution due to the limited power.

7 Conclusion

We investigated how professional developers interact with proactive AI assistance in authentic workflow contexts through a five-day in-the-wild study with 15 developers within a production IDE. We demonstrated that *timing* is key: suggestions presented at natural workflow boundaries (e.g., immediately post-commit) garnered the highest engagement rates, whereas mid-task interventions were often dismissed. Importantly, we observed sustained engagement patterns over the course of the study, indicating a potential for long-term integration of proactive assistance into everyday development workflows. Participants consistently valued ProAIDE's ability to surface overlooked issues, particularly when it functioned as a lightweight code quality gate. These interventions also supported perceived productivity and efficiency, given the observation that proactive suggestions were interpreted and applied significantly faster than reactive ones. Our findings bridge the gap between conceptual AI prototypes and production-ready IDE integration. We offer design insights and empirical evidence for embedding proactive AI into daily software engineering workflows, emphasizing adaptive support and user-configurable behavior to lay the groundwork for future research. While our five-day study revealed adoption patterns, longer-term research is needed to understand retention, evolving perceptions, and lasting impact.

8 GenAI Usage Disclosure

Portions of this manuscript were revised using OpenAI's ChatGPT for grammar and clarity improvements. The authors reviewed and edited all suggestions and take full responsibility for the content.

Acknowledgments

This work was supported by JetBrains as part of the AI for Software Engineering (AI4SE) collaboration with Delft University of Technology.

References

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [2] Anthropic. 2024. Claude. <https://code.claude.com/docs>.
- [3] Anthropic. 2024. Model Context and the Case for an Open Protocol. <https://www.anthropic.com/news/model-context-protocol> Accessed: 2025-07-19.
- [4] Alberto Bacchelli and Christian Bird. 2013. Expectations, outcomes, and challenges of modern code review. In *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 712–721.
- [5] Brian P Bailey, Joseph A Konstan, and John V Carlis. 2000. Measuring the effects of interruptions on task performance in the user interface. In *Smc 2000 conference proceedings. 2000 ieee international conference on systems, man and cybernetics. cybernetics evolving to systems, humans, organizations, and their complex interactions* (cat. no. 0, Vol. 2. IEEE, 757–762.
- [6] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh Rao, and Minoru Asada. 2016. Initiative in robot assistance during collaborative task execution. In *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 67–74.
- [7] Shraddha Barke, Michael B James, and Nadia Polikarpova. 2023. Grounded copilot: How programmers interact with code-generating models. *Proceedings of the ACM on Programming Languages* 7, OOPSLA1 (2023), 85–111.
- [8] Christian Bird, Denae Ford, Thomas Zimmermann, Nicole Forsgren, Eirini Kalliamvakou, Travis Lowdermilk, and Idan Gazit. 2023. Taking flight with copilot. *Commun. ACM* 66, 6 (2023), 56–62.
- [9] Islem Bouzenia, Premkumar Devanbu, and Michael Pradel. 2025. RepairAgent: An Autonomous, LLM-Based Agent for Program Repair. In *Proceedings of the 47th International Conference on Software Engineering (ICSE)*.
- [10] Nick C Bradley, Thomas Fritz, and Reid Holmes. 2022. Sources of software development task friction. *Empirical Software Engineering* 27, 7 (2022), 175.
- [11] Souti Chattopadhyay, Nicholas Nelson, Yenifer Ramirez Gonzalez, Annel Amelia Leon, Rahul Pandita, and Anita Sarma. 2019. Latent patterns in activities: A field study of how developers manage context. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 373–383.
- [12] Tianyi Chen. 2024. The Impact of AI-Pair Programmers on Code Quality and Developer Satisfaction: Evidence from TiMi studio. In *Proceedings of the 2024 International Conference on Generative Artificial Intelligence and Information Security (Kuala Lumpur, Malaysia) (GAIIIS '24)*. Association for Computing Machinery, New York, NY, USA, 201–205. doi:10.1145/3665348.3665383
- [13] Valerie Chen, Alan Zhu, Sebastian Zhao, Hussein Mozannar, David Sontag, and Ameet Talwalkar. 2025. Need Help? Designing Proactive AI Assistants for Programming. arXiv:2410.04596 [cs.HC] <https://arxiv.org/abs/2410.04596>
- [14] Bhavya Chopra, Ananya Singha, Anna Fariha, Sumit Gulwani, Chris Parnin, Ashish Tiwari, and Austin Z. Henley. 2023. Conversational Challenges in AI-Powered Data Science: Obstacles, Needs, and Design Opportunities. arXiv:2310.16164 [cs.HC] <https://arxiv.org/abs/2310.16164>
- [15] Umüt Cihan, Arda İçöz, Vahid Haratian, and Eray Tüzün. 2025. Evaluating Large Language Models for Code Review. arXiv:2505.20206 [cs.SE] <https://arxiv.org/abs/2505.20206>
- [16] Nicholas Clark, Matthew Dabkowski, Patrick J Driscoll, Dereck Kennedy, Ian Kloo, and Heidy Shi. 2021. Empirical decision rules for improving the uncertainty reporting of small sample system usability scale scores. *International Journal of Human-Computer Interaction* 37, 13 (2021), 1191–1206.
- [17] Cursor. 2023. Cursor: The AI-first Code Editor. <https://cursor.sh>.
- [18] Aral de Moor, Arie van Deursen, and Maliheh Izadi. 2024. A Transformer-Based Approach for Smart Invocation of Automatic Code Completion. In *Proceedings of the 1st ACM International Conference on AI-Powered Software (AIware '24)*. ACM, 28–37. doi:10.1145/3664646.3664760
- [19] 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.
- [20] Shuzheng Gao, Cuiyun Gao, Wenchao Gu, and Michael Lyu. 2024. Search-Based LLMs for Code Optimization. arXiv:2408.12159 [cs.SE] <https://arxiv.org/abs/2408.12159>
- [21] GitHub. 2021. GitHub Copilot. <https://github.com/features/copilot>.
- [22] Bo-Jhang Ho, Bharathan Balaji, Mehmet Koseoglu, Sandeep Sandha, Siyou Pei, and Mani Srivastava. 2020. Quick Question: Interrupting users for microtasks with reinforcement learning. (2020). <https://www.amazon.science/publications/quick-question-interrupting-users-for-microtasks-with-reinforcement-learning>
- [23] Karen Holtzblatt and Hugh Beyer. 1997. *Contextual design: defining customer-centered systems*. Elsevier.
- [24] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 159–166.
- [25] ICC/ESOMAR. 2022. *ICC/ESOMAR International Code on Market, Opinion and Social Research and Data Analytics*. <https://esomar.org/codes-and-guidelines>

- Accessed on: June 09, 2025.
- [26] Andrei Cristian Ionescu, Sergey Titov, and Maliheh Izadi. 2025. A Multi-agent Onboarding Assistant based on Large Language Models, Retrieval Augmented Generation, and Chain-of-Thought. *arXiv preprint arXiv:2503.23421* (2025).
 - [27] Shamsi T Iqbal, Piotr D Adamczyk, Xianjun Sam Zheng, and Brian P Bailey. 2005. Towards an index of opportunity: understanding changes in mental workload during task execution. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 311–320.
 - [28] Shamsi T. Iqbal, Piotr D. Adamczyk, Xianjun Sam Zheng, and Brian P. Bailey. 2005. Towards an index of opportunity: understanding changes in mental workload during task execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA) (CHI '05). Association for Computing Machinery, New York, NY, USA, 311–320. doi:10.1145/1054972.1055016
 - [29] Maliheh Izadi, Jonathan Katzy, Tim Van Dam, Marc Otten, Razvan Mihai Popescu, and Arie Van Deursen. 2024. Language models for code completion: A practical evaluation. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.
 - [30] Imen Jaoua, Oussama Ben Sghaier, and Houari Sahaoui. 2025. Combining Large Language Models with Static Analyzers for Code Review Generation. arXiv:2502.06633 [cs.SE] <https://arxiv.org/abs/2502.06633>
 - [31] JetBrains. 2025. JetBrains AI Assistant. <https://www.jetbrains.com/junie/>.
 - [32] Ellen Jiang, Edwin Toh, Alejandra Molina, Kristen Olson, Claire Kayacik, Aaron Donsbach, Carrie J Cai, and Michael Terry. 2022. Discovering the Syntax and Strategies of Natural Language Programming with Generative Language Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 386, 19 pages. doi:10.1145/3491102.3501870
 - [33] Genovefa Kefalidou and Sarah Sharples. 2016. Encouraging serendipity in research: Designing technologies to support connection-making. *International Journal of Human-Computer Studies* 89 (2016), 1–23.
 - [34] Nadine Kuo, Agnia Sergeyyuk, Valerie Chen, and Maliheh Izadi. 2025. Supplementary Materials: Proactive AI Assistant Study. OSF. https://osf.io/t98h3/?view_only=f53356ba505e43f0bae1bc66b1fdcb01 Accessed: 2026-01-12.
 - [35] Thomas D LaToza and Brad A Myers. 2010. Developers ask reachability questions. In *Proceedings of the 32Nd ACM/IEEE International Conference on Software Engineering-Volume 1*. 185–194.
 - [36] James R Lewis. 2018. The system usability scale: past, present, and future. *International Journal of Human-Computer Interaction* 34, 7 (2018), 577–590.
 - [37] Q Vera Liao, Matthew Davis, Werner Geyer, Michael Muller, and N Sadat Shami. 2016. What can you do? Studying social-agent orientation and agent proactive interactions with an agent for employees. In *Proceedings of the 2016 acm conference on designing interactive systems*. 264–275.
 - [38] Tianjian Liu, Hongzheng Zhao, Yuheng Liu, Xingbo Wang, and Zhenhui Peng. 2024. Compeer: A generative conversational agent for proactive peer support. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–22.
 - [39] Christian Meurisch, Cristina A Mihale-Wilson, Adrian Hawlitschek, Florian Giger, Florian Müller, Oliver Hinz, and Max Mühlhäuser. 2020. Exploring user expectations of proactive AI systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020), 1–22.
 - [40] André N Meyer, Laura E Barton, Gail C Murphy, Thomas Zimmermann, and Thomas Fritz. 2017. The work life of developers: Activities, switches and perceived productivity. *IEEE Transactions on Software Engineering* 43, 12 (2017), 1178–1193.
 - [41] André N Meyer, Thomas Fritz, Gail C Murphy, and Thomas Zimmermann. 2014. Software developers' perceptions of productivity. In *Proceedings of the 22nd ACM SIGSOFT international symposium on foundations of software engineering*. 19–29.
 - [42] Microsoft. 2018. Visual Studio IntelliCode. <https://visualstudio.microsoft.com/services/intellicode/>.
 - [43] Microsoft. 2023. IntelliSense. <https://learn.microsoft.com/en-us/visualstudio/ide/using-intellisense?view=vs-2022>.
 - [44] Hussein Mozannar, Gagan Bansal, Adam Fourney, and Eric Horvitz. 2023. When to Show a Suggestion? Integrating Human Feedback in AI-Assisted Programming. In *AAAI Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:259108906>
 - [45] Hussein Mozannar, Gagan Bansal, Adam Fourney, and Eric Horvitz. 2024. Reading Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming. arXiv:2210.14306 [cs.SE] <https://arxiv.org/abs/2210.14306>
 - [46] 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 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.
 - [47] Hussein Mozannar, Valerie Chen, Mohammed Alsobay, Subhro Das, Sebastian Zhao, Dennis Wei, Manish Nagireddy, Prasanna Sattigeri, Ameet Talwalkar, and David Sonntag. 2024. The RealHumanEval: Evaluating Large Language Models' Abilities to Support Programmers. arXiv:2404.02806 [cs.SE] <https://arxiv.org/abs/2404.02806>
 - [48] Nachiappan Nagappan, Thomas Ball, and Andreas Zeller. 2006. Mining metrics to predict component failures. In *Proceedings of the 28th international conference on Software engineering*. 452–461.
 - [49] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an LLM to Help With Code Understanding. arXiv:2307.08177 [cs.SE] <https://arxiv.org/abs/2307.08177>
 - [50] Zhenhui Peng, Yunhwan Kwon, Jian Lu, Ziming Wu, and Xiaojuan Ma. 2019. Design and evaluation of service robot's proactivity in decision-making support process. In *proceedings of the 2019 CHI conference on human factors in computing systems*. 1–13.
 - [51] Dorin Pomian, Abhiram Bellur, Malinda Dilhara, Zarina Kurbatova, Egor Bogomolov, Andrey Sokolov, Timofey Bryksin, and Danny Dig. 2024. EM-Assist: Safe Automated ExtractMethod Refactoring with LLMs. In *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering (FSE '24)*. ACM, 582–586. doi:10.1145/3663529.3663803
 - [52] Kevin Pu, Daniel Lazaro, Ian Arawjo, Haijun Xia, Ziang Xiao, Tovi Grossman, and Yan Chen. 2025. Assistance or Disruption? Exploring and Evaluating the Design and Trade-offs of Proactive AI Programming Support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, 1–21. doi:10.1145/3706598.3713357
 - [53] 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.
 - [54] Advait Sarkar and Ian Drosos. 2025. Vibe coding: programming through conversation with artificial intelligence. *arXiv preprint arXiv:2506.23253* (2025).
 - [55] Jeff Sauro and James R Lewis. 2016. *Quantifying the user experience: Practical statistics for user research*. Morgan Kaufmann.
 - [56] SciPy Developers. 2023. *scipy.stats.wilcoxon - SciPy v1.11.3 Manual*. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.wilcoxon.html> Accessed: 2025-07-18.
 - [57] 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.
 - [58] Agnia Sergeyyuk, Ekaterina Koshchenko, Ilya Zakharov, Timofey Bryksin, and Maliheh Izadi. 2024. The Design Space of in-IDE Human-AI Experience. arXiv:2410.08676 [cs.SE] <https://arxiv.org/abs/2410.08676>
 - [59] Agnia Sergeyyuk, Ilya Zakharov, Ekaterina Koshchenko, and Maliheh Izadi. 2025. Human-AI Experience in Integrated Development Environments: A Systematic Literature Review. arXiv:2503.06195 [cs.SE] <https://arxiv.org/abs/2503.06195>
 - [60] Benjamin 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).
 - [61] Tabnine. 2020. Tabnine. <https://docs.tabnine.com/main>.
 - [62] Xin Tan, Xiao Long, Xianjun Ni, Yinghao Zhu, Jing Jiang, and Li Zhang. 2024. How far are AI-powered programming assistants from meeting developers' needs? arXiv:2404.12000 [cs.SE] <https://arxiv.org/abs/2404.12000>
 - [63] Adam Tornhill and Markus Borg. 2022. Code red: the business impact of code quality—a quantitative study of 39 proprietary production codebases. In *Proceedings of the International Conference on Technical Debt*. 11–20.
 - [64] Priyan Vaithilingam, Elena L Glassman, Peter Groenewegen, Sumit Gulwani, Austin Z Henley, Rohan Malpani, David Pugh, Arjun Radhakrishna, Gustavo Soares, Joey Wang, et al. 2023. Towards more effective ai-assisted programming: A systematic design exploration to improve visual studio intellicode's user experience. In *2023 IEEE/ACM 45th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 185–195.
 - [65] Nalin Wadhwa, Jui Pradhan, Atharv Sonwane, Surya Sahu, Nagarajan Natarajan, Aditya Kanade, Suresh Parthasarathy, and Sriram Rajamani. 2024. CORE: Resolving Code Quality Issues using LLMs. *Proceedings of the ACM on Software Engineering* 1 (07 2024), 789–811. doi:10.1145/3643762
 - [66] Justin D Weisz, Michael Muller, Jessica He, and Stephanie Houde. 2023. Toward General Design Principles for Generative AI Applications. In *Joint Workshops on Human-AI Co-Creation with Generative Models and User-Aware Conversational Agents*.
 - [67] Justin D. Weisz, Michael Muller, Stephanie Houde, John Richards, Steven I. Ross, Fernando Martinez, Mayank Agarwal, and Kartik Talamadupula. 2021. Perfection Not Required? Human-AI Partnerships in Code Translation. In *26th International Conference on Intelligent User Interfaces (IUI '21)*. ACM, 402–412. doi:10.1145/3397481.3450656
 - [68] Windsurf. 2023. Windsurf AI Code Editor. <https://windsurf.ai>.
 - [69] Di Wu, Fangwen Mu, Lin Shi, Zhaoqiang Guo, Kui Liu, Weiguang Zhuang, Yuqi Zhong, and Li Zhang. 2024. iSMELL: Assembling LLMs with Expert Toolsets for Code Smell Detection and Refactoring. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering (Sacramento, CA, USA) (ASE '24)*. Association for Computing Machinery, New York, NY, USA, 1345–1357. doi:10.1145/3691620.3695508
 - [70] Yisen Xu, Feng Lin, Jinqiu Yang, Tse-Hsun, Chen, and Nikolaos Tsantalis. 2025. MANTRA: Enhancing Automated Method-Level Refactoring with Contextual RAG and Multi-Agent LLM Collaboration. arXiv:2503.14340 [cs.SE] <https://arxiv.org/abs/2503.14340>

- org/abs/2503.14340
- [71] Aiko Yamashita and Leon Moonen. 2013. Do developers care about code smells? An exploratory survey. In *2013 20th working conference on reverse engineering (WCRE)*. IEEE, 242–251.
- [72] 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.
- [73] Zhongzhi Yu, Mingjie Liu, Michael Zimmer, Yingyan Celine Lin, Yong Liu, and Haoxing Ren. 2025. Spec2RTL-Agent: Automated Hardware Code Generation from Complex Specifications Using LLM Agent Systems. arXiv:2506.13905 [cs.AR] <https://arxiv.org/abs/2506.13905>
- [74] Sebastian Zhao, Alan Zhu, Hussein Mozannar, David Sontag, Ameet Talwalkar, and Valerie Chen. 2025. CodingGenie: A Proactive LLM-Powered Programming Assistant. arXiv:2503.14724 [cs.HC] <https://arxiv.org/abs/2503.14724>