



Productivity gain, community strain: Stack overflow's community response to the AI initiatives

Title	Productivity gain, community strain: Stack overflow's community response to the AI initiatives
Author(s)	Scholtz, Dewan;Griva, Anastasia;Zamani, Efpraxia D.;Conboy, Kieran
Publication Date	2026-02-16
Publisher	Institute of Electrical and Electronics Engineers;Computer Society
Repository DOI	https://doi.org/10.1109/MS.2026.3665306

THEME ARTICLE: The Impact of AI on Productivity and Code

Productivity Gain, Community Strain: Stack Overflow's Community Response to the AI Initiatives

Dewan Scholtz, J.E. Cairnes School of Business and Economics; Lero Science Foundation Ireland Research Centre for Software; University of Galway, Galway, Ireland

Anastasia Griva J.E. Cairnes School of Business and Economics; Lero Science Foundation Ireland Research Centre for Software; University of Galway, Galway, Ireland

Efraxia D. Zamani, Durham University Business School, Durham University, Durham, UK

Kieran Conboy, J.E. Cairnes School of Business and Economics; Lero Science Foundation Ireland Research Centre for Software; University of Galway, Galway, Ireland

We examine how Stack Overflow's AI initiatives impact its open-source knowledge community. Motivated by visible community fracture, such as a moderator strike and content vandalism, we use netnography with trace data to follow contributor reactions to AI venture announcements, revealing challenges around attribution, compensation, autonomy, trust, and representation. We then provide a set of recommendations that community developers or moderators may implement to help address or mitigate these challenges.

INTRODUCTION

In the current landscape, AI tools like GitHub Copilot, Amazon Q, and OverflowAI are being promoted as breakthroughs in developer productivity, promising faster coding, smarter solutions, and streamlined workflows. However, this growing emphasis on productivity seems to forget about the open-source communities behind the scenes, raising the question of how AI tools can grow productivity without eroding the health of the communities that sustain and curate their training data.

Our research examines how this pursuit of greater developer productivity and AI advancement can overlook the open-source knowledge communities that contributed, curated, and sustained knowledge repositories, now used to train AI tools or models. Using Stack Overflow's development and deployment of OverflowAI as a case study, we analyse community responses to key initiatives and partnerships aimed at advancing AI tool productivity. These responses surface concerns around attribution, compensation, autonomy,

trust, and representation, potentially tied to community fracture, as indicated in recent protests, such as a moderator strike and the poisoning of data. Understanding these issues could prove crucial for identifying potential risks for community health and sustainability, thereby indirectly affecting the quality and reliability of the AI tools and models trained on their collective knowledge.

Stack Overflow's knowledge base is built through over 17 years of user-generated questions and answers, curated via community moderation. Users upvote quality posts, flag issues, and edit collaboratively, while elected moderators enforce rules, making this context ideal for investigating an established community behind AI training data for coding tools.

Related work

To contextualise current research, we organize the relevant literature into two domains: Productivity- and Community-focused.

1. Productivity-Focused Research

Prior research shows that AI tools can accelerate routine programming tasks, improve completion speed[1], [2], and support developers in overcoming basic coding hurdles[3]. These gains are accompanied by growing concerns about reduced code quality, subtle bugs, increased churn[4], [5], and uneven adoption in professional environments[6].

While these studies assess how AI tools improve short-term coding efficiency, their focus is not on the humans enabling these systems, specifically, the open-source knowledge communities whose contributions train and sustain them. Therefore, we now shift our focus from the question “does AI enhance productivity” to “who enables AI productivity”, referring to the contributors of its training data, and what the challenges that these communities now face, are.

2. Community-Focused Research

Studies focused on the community show that AI tools have not only changed where developers seek help but also how they participate in open-source knowledge platforms. For example, platform-level analyses[7], [8] reveal steep declines in Stack Overflow engagement after the release of ChatGPT, suggesting a shift away from community-based problem-solving. Others[9], [10] have identified disruptions in content quality, interaction patterns, and the flow of knowledge, especially as AI begins to replace lower-effort or repetitive contributions. Further studies[11] highlight divisions in developer attitudes, with mistrust in AI tools and dissatisfaction with how platforms implement them. Also, developers expect transparent and informed consent[12] while a significant portion of contributors’ work remains invisible and uncompensated[13].

While this body of work frames AI as the key driver of these shifts, our research focuses on open-source community fracture as a response not just to AI’s presence, but to how the platform treats its contributors, investigating how communities react to being co-opted for commercial AI use. There is also limited empirical research on community attitudes toward how their data is used and represented in AI systems, despite the potential for these attitudes to reveal important links between community health and the long-term quality of AI training data.

Through a netnographic analysis, backed with quantitative trace data of community responses, we investigate how the community is fractured through looking into developer dispositions, highlighting how issues of 1) attribution, 2) compensation, 3) autonomy, 4) trust, and 5) representation are playing out in real-time. We explore how disregarding contributor communities could risk undermining the quality of training data, as reduced engagement and acts of protest, like data poisoning or user strikes, can weaken the foundations of AI systems.

The paper is structured as follows. 1) *Community Fracture* lays out the motivation of the study by exploring

some indicators of decline in community health that co-occurred with Stack Overflow’s AI development and partnership initiatives. 2) “*Front-end*” of *OverflowAI* explores how Stack Overflow’s management publicly framed their AI tool, OverflowAI, and introduced its related initiatives. 3) “*Back-end*” of *OverflowAI* contrasts this public narrative with the internal community’s reactions to these announcements and initiatives. 4) *Suggestions and Discussion*: Identify key stakeholders and propose practical approaches to address the challenges of maintaining a healthy community while curating training data for AI systems.

METHOD

We applied netnography—“participant-observational research based in online fieldwork”[14]—to explore responses to Stack Overflow’s AI initiatives, combining active platform immersion from April 2023 to July 2025 with systematic documentation of community discourse. Ethnographic fieldwork was conducted within: Meta discussions and threads (focused on only the reactions towards four key posts, see section “*Back-End*” of *OverflowAI*), official communications, internal commentary, secondary surveys and reports, protest events and moderator actions, and corporate announcements. To triangulate contextual indicators and netnographic observations, we supplemented qualitative data, using Stack Exchange Data Explorer (SEDE, <https://data.stackexchange.com/>) queries to establish thresholds and validate reported patterns. On the query results, we ran a sentiment analysis and topic modelling in Python on comments, questions, and answers, relating to AI initiative posts using Hugging Face natural language processing models (<https://huggingface.co/models>).

The study focuses exclusively on content that is publicly accessible, without engaging directly with participants or accessing private or restricted materials.

Our data includes user comments, threaded discussions, and official communications from Stack Overflow. Based on the platform’s policy, only users with at least 50 reputation points; 37% active users can comment. Beyond this, for our netnographic analysis, we sampled comments from active users with more than 1000 reputation points (platform average points: 797; 9% active userbase), ensuring that our study draws from influential posts and contributors. Based on this threshold, a total of 81 posts (out of 328 in total), with >7900 net score (total upvotes minus downvotes) was thematically analysed, resulting in five inductively identified themes: attribution, compensation, autonomy, trust, and representation. Running a Hugging Face topic modelling model on all posts revealed similar themes.

Attitudes within these posts were categorised as positive (in agreement), negative (opposed), or neutral (indifferent or apathetic), first assigned manually, and then triangulated through a sentiment analysis on all the posts. We also coded

numeric polarity with votes combined with sentiment as a weighted understanding of approval or disapproval across topics, whereby the higher the net score, the greater the alignment with the announcement/post (e.g., post n has a net positive score showing community agreement, or post n with a net negative score showing disagreement). To upvote, users require 15 reputation points (55% active users), and to downvote, they require 125 reputation points (38% active users), showing further thresholds for user interactions. For more context, the top 100 posts (Apr 2023 – Jul 2025) have 303 net votes on average, further positioning the chosen posts as influential (seen below in the “Back-end” of OverflowAI section).

COMMUNITY FRACTURE

Two community protest actions suggest the existence of community fracture: the June 2023 moderator strike and data poisoning, coinciding with a sustained drop in user participation.

Moderator Strike

From June to August 2023, Stack Overflow moderators initiated a coordinated strike in protest of the platform’s decision to overturn its temporary ban on AI-generated content. The original ban had received the community’s support, registering +5249 net upvotes, the highest ever score recorded on Meta Stack Exchange, indicating that many community members stood against the inclusion of AI-generated responses. When the ban was reversed without

consultation, moderators issued a public statement and ceased all moderation activities (+1208 net votes). In a strike letter (<https://openletter.mousetail.nl/#signatures>), signed by 125 moderators, they stated, among others, that they will not close posts or handle flags. This coincided with a noticeable dip in questions closed (see figure 1, top left) and the backlog of tags, shared by a moderator as a screen capture during this period (see figure 1, bottom right), indicating that curation during this time was lacking.

Data Poisoning

Later contributors reported sabotaging their own content as a form of protesting the commodification of their knowledge. Reports indicated that users deliberately obfuscated, altered, or deleted content they had previously curated (e.g., in Figure 1, bottom: a user “deleted 1192 characters” from their post, and another changed their content to gibberish like the “hhhhh...”), stating it was done in response to the OpenAI and Google partnerships. Others began changing their usernames to messages like “X no longer feeds AI” to make their dissent publicly visible. A spike in moderator rollbacks, seen after the OpenAI partnership announcement, could indicate the impact of the data poisoning.

Mass Exit of Contributors

Earlier studies have suggested that there is a relationship between the decline (in terms of participants) and AI substitution[8], which could pose a threat to the “digital public

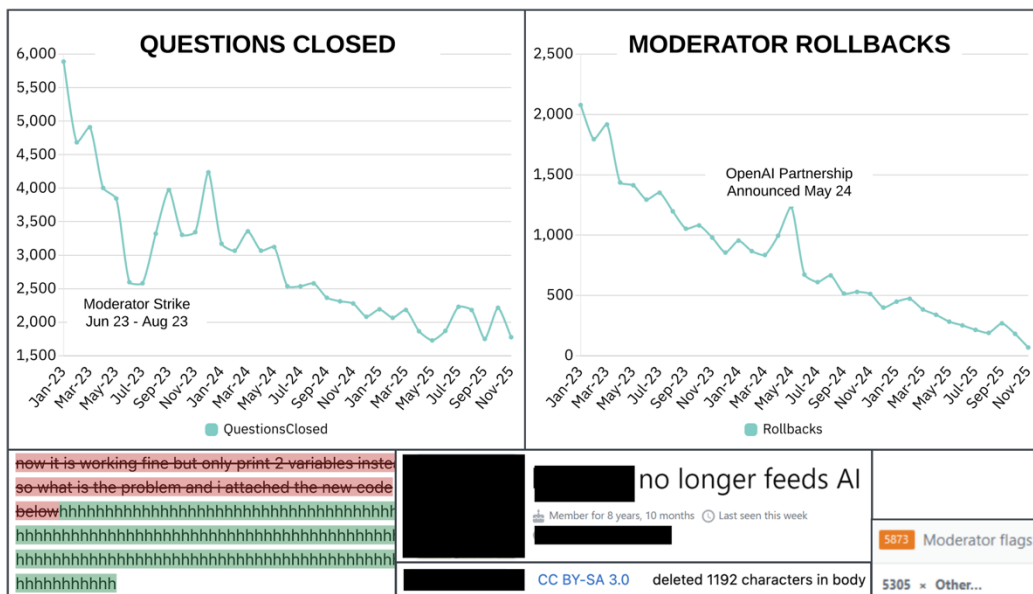


FIGURE 1. SEDE data showing a dip in questions closed during the moderator strike (top left), a spike in moderator rollbacks around the time of the OpenAI partnership (top right), and examples of data poisoning, and moderator flag backlogs (bottom) screen captured during these periods. Personal user information redacted for anonymity.

good”[7]. In this study, however, and while we cannot make any causal inferences, we do suggest that a mass exit of contributors might in parallel suggest the decline of the curating community within Stack Overflow, and which could then result over the longer term in further deterioration, and particularly in terms of high quality data for AI training due to absence of content moderation.

“FRONT-END” OF OVERFLOWAI

Since 2023, Stack Overflow has publicly framed OverflowAI as a developer-first initiative designed to responsibly integrate AI with community knowledge. Through blog posts, feature launches, and partnerships, the platform emphasized trust, attribution, and productivity, positioning AI as a complement to human expertise. We present a chronological overview of these ‘front-end’ initiatives, which we later contrast with the ‘back-end’ community responses.

April 2023: In a blog post, “Community is the future of AI” (<https://stackoverflow.blog/2023/04/17/community-is-the-future-of-ai/>), CEO Prashanth Chandrasekar framed Stack Overflow’s AI strategy around its community: “Grounding our responses in the knowledge... Stack Overflow helps users to understand the provenance of the code they hope to use”. Chandrasekar stressed that Stack Overflow would reward contributors and maintain human curation: “AI models are powerful because they were trained on... Stack Overflow Q&A. We want to craft models that reward the users who contribute and keep the knowledge base... open and growing.”

July 2023: At the 2023 WeAreDevelopers conference (stackoverflow.blog/announcing-overflowai/), Stack Overflow formally unveiled OverflowAI, its umbrella for all AI-powered features, including:

- ➔ Integrated Developer Environment (IDE) integrations, applications that help programmers develop software code efficiently: trained on all surfaced “answers from 58+ million posts”, freely contributed and curated by over 25 million experienced Stack Overflow users, directly within workflows.
- ➔ Auto-ingest internal knowledge into Stack Overflow, improving current features and increasing developer productivity on the platform itself.

The company repeatedly stressed that AI would augment, not replace, the community: “if you can’t find what you need within our extensive data, the Stack Overflow community is readily available to fill in the gaps beyond AI’s capabilities”. OverflowAI was said to be designed to boost developer productivity by surfacing relevant, cited answers faster.

May 2024: Stack Overflow and OpenAI announced a strategic collaboration (stackoverflow.co/openai-partnership/). OpenAI will gain access to Stack Overflow’s content via OverflowAPI, a subscription-based Application

Programming Interface (API) service that provides continuous access to Stack Overflow’s public dataset to train and fine-tune a large language model, and Stack Overflow will use OpenAI’s models to refine its own products. The release said OpenAI would “surface validated technical knowledge from Stack Overflow directly into ChatGPT,” ensuring answers in ChatGPT are “trusted, attributed, [and] accurate,” thanks to the millions of community contributions.

Around this time, in a blog post, CPO Ryan Polk announced that OverflowAI features were available (stackoverflow.blog/introducing-overflowai-a-new-era-of-community-driven-ai-at-stack-overflow/), also stating that Stack Overflow’s recent developer survey showed most developers want AI but many distrust it, and Polk promised continuous improvement based on community feedback, with the vision of making OverflowAI “an indispensable tool for developers”.

November 2024: Stack Overflow publicized that OverflowAPI won the 2024 API World Award for “Best AI API” (stackoverflow.co/stack-overflow-overflow-api-award/). A press release framed OverflowAPI as “empowering developers through innovative, socially responsible AI” by granting AI builders access to Stack Overflow’s public dataset. It also highlighted Stack Overflow’s AI ethics stance: partners using OverflowAPI must commit to attribution to contributors and use high-quality data in alignment with the company’s “socially responsible AI” framework.

July 2025: At the 2025 WeAreDevelopers conference, Stack Overflow unveiled a new vision and architecture (stackoverflow.blog/a-new-era-of-stack-overflow/). Its data-licensing business was rebranded as Knowledge Solutions, with partners like OpenAI and Google Cloud using its content to power AI. Their renewed mission is “to be the world’s most vital source for technologists,” leveraging community expertise alongside AI.

As shown above, OverflowAI’s public narrative centres on ethical, community-driven AI that enhances developer workflows while preserving trust and transparency. While Stack Exchange Inc. presented OverflowAI through a series of carefully worded announcements that emphasized productivity, innovation, and community alignment, the response from the Stack Overflow contributor base has painted a vastly different picture.

“BACK-END” OF OVERFLOWAI

To understand the contrast between the public narrative and the community’s response, we examined user reactions to initiatives that signaled the company’s strategic AI direction. These initiatives were communicated to the internal community through four key posts:

- ➔ **2025 Developer Survey: 13306 views, -131 net score**

- **ChatGPT API Integration: 23820 views, -216 net score**
- **Google AI Partnership: 25520 views, -459 net score**
- **OpenAI Partnership: 56117 views, -1065 net score**

These announcements/posts were not received as collaborative milestones but as top-down decisions and were met with disapproval from their respondents. By manually identifying themes based on our netnographic observations and triangulation through topic modelling, we show below how each announcement triggered hundreds of comments and related questions (328 across all posts), from long-time contributors, moderators, and developers, many of whom expressed concern about the health of their community, which we explore below through five themes.

Community Attribution

Attribution emerges as one of the most criticized aspects of Stack Overflow's OverflowAI initiative. Some contributors' comments reflect a sense of exploitation and disenfranchisement, who feel their voluntary labour has been monetized without acknowledgment or benefit.

Out of the posts where this theme appeared, the reactions were 98% negative, and only 2% positive, indicating that the respondents find current attribution mechanisms insufficient for their efforts.

Many describe a lack of meaningful attribution, where their content is harvested to train AI models or fuel commercial APIs without direct credit, visibility, or compensation. Some respondents are especially frustrated that attribution is relegated to vague legal footnotes or buried in API documentation, rather than made visible at the point of use (ChatGPT API Integration, seen in Figure 2, right), for example:

"We are literally curating Google's data set for free", net score: +107.

"So what exactly do we get for curating a multi-billion dollar corporation's data set for them, other than they will make us something cool that we can pay them for?", net score: +314.

Such comments show how some users perceive an imbalance where labour is extracted without reciprocal benefit or recognition.

This attitude clashes with the company's public-facing commitments to attribution. In the public announcement of the OpenAI partnership, Stack Overflow emphasized that content shared with OpenAI would be "trusted, attributed, [and] accurate." Similarly, in the public Google partnership announcement, the company positioned itself as a proponent of "socially responsible AI," with a stated commitment to fair data use and attribution. The sentiment towards each of these announcements (Figure 2, right) reveals a contrary narrative.

Additionally, all Stack Overflow content is licensed under CC BY-SA 4.0, which requires attribution and share-alike licensing. Respondents note that these protections do not translate effectively when content is used to train AI models. Designed for human-readable reuse, the license's attribution and reciprocity requirements are often invisible or impractical in AI contexts, leaving some contributors feeling that the spirit of the policy is bypassed.

Community Compensation

The theme of community compensation reveals frustration among respondents, whereby there is an attitude that Stack Overflow is commercializing user-generated content without fair compensation or transparent benefit-sharing mechanisms. From all posts coded under this theme, 99% of reactions were negative. (see Figure 2, left). This shows that some contributors feel extracted from, not included in, the monetization of the community's collective knowledge.

Based on these interactions, contributors who have voluntarily curated, edited, and answered technical questions for years now see that same content being sold or licensed to




FIGURE 2. Percentage sentiment per theme (left) and average sentiment score towards each post, based on related questions, comments and answers (right). Data gathered from SEDE, cleaned and coded in Python, themes and sentiment triangulated through topic modelling and sentiment analysis using Hugging Face, neutral sentiments excluded.

corporations like OpenAI and Google, generating revenue with no tangible return for those who produced it.

“I’m all for us getting financially compensated since we’re the ones who authored (and curated) the content being sold.”, net score: +104.

A few comments argue that this isn’t just about money, it’s about the principle of equitable value recognition. The platform is viewed as having shifted its allegiance from its contributor base to enterprise clients, particularly through the OverflowAPI licensing model.

Community Autonomy

Respondents critique the company for not obtaining explicit consent or even informing contributors about the nature of data usage in AI partnerships (Figure 2, right). Here, 84% of votes expressed negative attitudes (Figure 2, left).

Some question who owns the data they contribute and whether they ever truly agreed to its current uses. Others raise concerns about privacy, platform overreach, and the absence of opt-in/opt-out mechanisms for licensing or training AI, for example:

“I wish I had the option to cast a ‘veto vote’ regarding the partnership with Google”, Net score: +4.

“I feel violated, cheated upon, betrayed, and exploited... Humans are meant to exploit machines, not the other way round...”, net score: +460.

What some of these comments suggest is a feeling of powerlessness: contributors are subject to platform decisions over how their content, often years' worth of accumulated technical labour, is being sold, repurposed, and embedded in external AI systems without their explicit agreement. As mentioned, Stack Overflow content is licensed under CC BY-SA 4.0, and some respondents feel AI companies ignore or cannot enforce these terms, disregarding user autonomy by bypassing license obligations. These respondents feel they have lost control over their work, fuelling frustration despite lacking legal ownership. For many, it’s about respect and being acknowledged for building the platform’s knowledge.

Community Trust

Attitudes around community trust on Stack Overflow reflect the lack of respondents’ confidence that the platform aligns its actions with its stated values and involves the community in meaningful decision-making (Figure 2, right).

Across the responses, attitudes on all the posts were negative (Figure 2, left). This indicates that, among the users who responded, there is a lack of confidence in the platform’s positioning and narrative of being a responsible AI actor.

The respondents rejected Stack Overflow’s public-facing language of “socially responsible AI”, “trusted attribution”,

and “commitment to community values”, seeing these as hollow or performative, considering decisions made without community consultation or endorsement.

Some contributors criticized Stack Overflow’s partnerships and AI ambitions as profit-motivated and opaque, cloaked in vague ethical language, for example:

“This assumes that Google is ‘Socially Responsible’ simply because they say so. I see evidence that is contrary to this claim”, net score: +40.

Community Representation

Community representation on Stack Overflow refers to the extent to which users feel their voices, concerns, and values are acknowledged and reflected in platform decisions.

For representation, the attitude was negative (Figure 2, left), with many respondents reporting a sense of exclusion, marginalization, and disenfranchisement.

The frustration reportedly stems from a breakdown in the platform’s traditional community contract, where users once played a central role in shaping moderation, policy, and direction. Some comments express that the platform’s decisions are imposed, and claims of community-centeredness ring hollow, for example:

“It’s been pretty evident for a while now that the commitment is indeed to AI, not so much the community/ies.”, net score: +56.

Notably, the responses don’t just reflect the absence of representation, but also some claiming active misrepresentation, particularly in the framing of Stack Overflow’s 2025 Developer Survey (Figure 2, right). Respondents have accused the company of push-polling, structuring survey questions to elicit favourable responses toward AI integration, despite community resistance. Though these survey results were cited in blog posts (<https://survey.stackoverflow.co/2025/>) to justify AI expansions like OverflowAI, some users rejected the claim that these reflected the true developer stance, for example:

“A lottttt of AI questions on this survey. Not surprising but they seem to be worded very favorably towards a rather optimistic view of LLM coding capabilities”, net score: +94.

SUGGESTIONS & DISCUSSION

The following recommendations are directed at stakeholders most affected by or responsible for AI integration in software development:

- ➔ AI companies and platform operators (e.g., Stack Overflow, GitHub, OpenAI)
- ➔ Software development teams and CTOs
- ➔ Open-source contributors and community moderators
- ➔ Policymakers and digital rights advocates
- ➔ AI ethics researchers and academic institutions

Recommendations

Suggestion	What It Entails	What Stack Overflow Shows Us	What Other Open-Source Examples Show Us
Proactive Community Engagement	Include contributors in decision-making; transparent governance.	Frustrations due to exclusion from AI decisions and AI training. Users rejected the initiative because they were excluded during the partnership discussions and roll-outs.	Wikipedia's Request for Comment (RFC) process involves the community in policy decisions, building trust (meta.wikimedia.org). This could be adapted to include them in discussions surrounding new initiatives.
Align Community Health with Data Quality	Protect contributor engagement to maintain fresh, accurate data.	Community fractures, vandalism, and moderator strikes degrade data quality. There's a connection between the health of the community and the quality of its data.	GitHub's community health tools (discussions, code reviews, contributor agreements) encourage respectful, constructive engagement (github.com). This foundation also allowed the community to be more included in their AI initiatives.
Consent and Opt-In/Out Mechanisms	Let users refuse AI training on their content.	Users explicitly want control over AI use of their contributions. They express their frustration with a lack of an opt-in/out mechanism.	Mozilla Common Voice platform allows contributors to donate voice data with clear consent and opt-out options (commonvoice.mozilla.org). Adopting such a system, before allowing third-party AI training, will alleviate any further community fracture.
Attribution and Recognition Mechanisms	Credit contributors for AI-trained content.	Frustration over the lack of contributor credit and uncertainty about whether their attribution policies will be enforced externally.	OpenStreetMap's attribution guidelines require crediting contributors for data use (wiki.openstreetmap.org). Whether or not this can be enforced on AI tools is yet to be explored.
Partnership and Initiative Transparency	Publicly disclose AI partnerships and data use terms.	Distrust due to opaque corporate AI collaborations and non-transparency surrounding what these ventures entail.	Wikimedia's Transparency Report details partnership terms and data requests, fostering openness (wikimediafoundation.org). Including the community in all aspects of partnerships and initiatives during and before decisions.

TABLE 1. Suggestions for ensuring community health and alleviating fractures when training AI on open-source contributions.

This study offers a focused analysis of four influential posts, and its insights remain bounded by the specific interactions and respondents examined. The approach cannot support platform-wide causal inference, yet it surfaces dynamics that merit broader scrutiny. The findings suggest that attempts to secure productivity gains through AI training can obscure the underlying dependencies on community-generated datasets. Additionally, the observed and previously reported decline in community engagement raises concerns about the long-term stability of such models. Future research should extend these observations using longitudinal data, examine cross-platform patterns of community deterioration, and assess how shifts in contributor participation influence model reliability. Understanding these linkages could be essential for designing AI systems that remain credible as their social foundations evolve.

REFERENCES

- [1] A. Moradi Dakhel, V. Majdinasab, A. Nikanjam, F. Khomh, M. C. Desmarais, and Z. M. (Jack) Jiang, "GitHub Copilot AI pair programmer: Asset or Liability?," *J. Syst. Softw.*, vol. 203, p. 111734, Sept. 2023, doi: 10.1016/j.jss.2023.111734.
- [2] Q. Gallea, "From Mundane to Meaningful: AI's Influence on Work Dynamics -- evidence from ChatGPT and Stack Overflow," 2023, arXiv. doi: 10.48550/ARXIV.2308.11302.
- [3] M. Kazemitabaar, X. Hou, A. Henley, B. J. Ericson, D. Weintrop, and T. Grossman, "How Novices Use LLM-based Code Generators to Solve CS1 Coding Tasks in a Self-Paced Learning Environment," in *Proceedings of the 23rd Koli Calling International Conference on Computing Education Research*, Koli Finland: ACM, Nov. 2023, pp. 1–12. doi: 10.1145/3631802.3631806.

- [4] N. Perry, M. Srivastava, D. Kumar, and D. Boneh, “Do Users Write More Insecure Code with AI Assistants?,” in Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, Copenhagen Denmark: ACM, Nov. 2023, pp. 2785–2799. doi: 10.1145/3576915.3623157.
- [5] D. Nam, A. Macvean, V. Hellendoorn, B. Vasilescu, and B. Myers, “Using an LLM to Help With Code Understanding,” in Proceedings of the IEEE/ACM 46th International Conference on Software Engineering, Lisbon Portugal: ACM, Apr. 2024, pp. 1–13. doi: 10.1145/3597503.3639187.
- [6] S. Pan et al., “‘I Don’t Use AI for Everything’: Exploring Utility, Attitude, and Responsibility of AI-empowered Tools in Software Development,” 2024, arXiv. doi: 10.48550/ARXIV.2409.13343.
- [7] M. del Rio-Chanona, N. Laurentsyevea, and J. Wachs, “Are Large Language Models a Threat to Digital Public Goods? Evidence from Activity on Stack Overflow,” 2023, doi: 10.48550/ARXIV.2307.07367.
- [8] G. Burtch, D. Lee, and Z. Chen, “The consequences of generative AI for online knowledge communities,” *Sci. Rep.*, vol. 14, no. 1, p. 10413, May 2024, doi: 10.1038/s41598-024-61221-0.
- [9] L. Da Silva, J. Samhi, and F. Khomh, “LLMs and Stack Overflow discussions: Reliability, impact, and challenges,” *J. Syst. Softw.*, vol. 230, p. 112541, Dec. 2025, doi: 10.1016/j.jss.2025.112541.
- [10] L. Zhong and Z. Wang, “Can LLM Replace Stack Overflow? A Study on Robustness and Reliability of Large Language Model Code Generation,” *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 19, pp. 21841–21849, Mar. 2024, doi: 10.1609/aaai.v38i19.30185.
- [11] J. R. Cheng, R. Wang, T. Zimmermann, and D. Ford, “‘It would work for me too’: How Online Communities Shape Software Developers’ Trust in AI-Powered Code Generation Tools,” *ACM Trans. Interact. Intell. Syst.*, vol. 14, no. 2, pp. 1–39, June 2024, doi: 10.1145/3651990.
- [12] G. Pistilli and B. Trevelin, “Can AI be Consentful?,” 2025, arXiv. doi: 10.48550/ARXIV.2507.01051.
- [13] J. Meluso, A. Casari, K. McLaughlin, and M. Z. Trujillo, “Invisible Labor in Open Source Software Ecosystems,” Jan. 10, 2025, arXiv: arXiv:2401.06889. doi: 10.48550/arXiv.2401.06889.
- [14] R. V. Kozinets, “Netnography,” in *The International Encyclopedia of Digital Communication and Society*, John Wiley & Sons, Ltd, 2015, pp. 1–8. doi: 10.1002/9781118767771.wbiedcs067.

DEWAN SCHOLTZ is a Ph.D. researcher in Business Information Systems at the University of Galway, Ireland, financially supported by the Hardiman Scholarship. His research focuses applies a temporal lens, researching how time and control manifests in open knowledge communities, particularly focused on the impact of Generative AI. Contact him at d.scholtz1@universityofgalway.ie.

ANASTASIA GRIVA is an Associate Professor at the University of Galway, Ireland. Her research focuses on Business Analytics, Responsible AI, and Data-Driven Innovation. Griva received her Ph.D. from Athens University of Economics and Business, Greece. Contact her at anastasia.griva@universityofgalway.ie.

KIERAN CONBOY is a Professor at the University of Galway, Ireland. His research interests include contemporary technology management and design, including concepts such as temporality, flow, open innovation, and agility. Kieran is a member of the Royal Irish Academy, previously worked with Accenture and leads numerous industry projects in this area. Contact him at kieran.conboy@universityofgalway.ie.

EFPRAXIA D. ZAMANI is a Professor of Information Systems at Durham University Business School, Durham, UK. Her research focuses on the intersection of societal and organizational implications of Information Systems and emerging technologies. She has received her PhD from Athens University of Economics and Business, Greece. Contact her at Efpraxia.zamani@durham.ac.uk.