

The Fine Balance Between Helping With Your Job And Taking It: AI Code Assistants Come To The Fore

Cleidson R. B. de Souza¹, Gema Rodríguez-Pérez², Manaal Basha², Dongwook Yoon², Ivan Beschastnikh²

¹Universidade Federal do Pará, Belém, Brazil

²University of British Columbia, Canada

Abstract—AI code generation tools are reshaping the software engineering landscape. We report on the promises, use cases, and concerns about these tools through narratives that we have collected regarding two AI code generation tools, GitHub Copilot and Tabnine. For this, we study data we collected from Twitter over the last 18 months. Based on our analysis, we provide recommendations for software practitioners. We hope that our study inspires practitioners to try these tools but to also keep in mind some of the trade-offs in using these tools for professional software development.

I. INTRODUCTION

Programming is a non-trivial human activity. Until very recently, the idea of automating code generation was considered a pipe dream¹. This radically changed with the first generation of AI code assistants, all released within the past few years. These include Tabnine (released 2018), GitHub Copilot (released 2021), ChatGPT (2022), and Amazon CodeWhisperer (released 2022); all of these are described as “assistants” that are intended to help increase developer productivity.

These tools are trained on hundreds of millions of lines of source code. They work by translating a natural language prompt into a corresponding implementation. For example, prompting Copilot for “a function that computes the i^{th} element in the Fibonacci sequence” will result in a function (in the programming language of the project) that implements precisely this logic. These tools may seem like a novelty. But, large companies like Google have already reported astonishing results — a Google blog post reports that “3 percent of new code (measured in characters) is now generated by accepting ML completion suggestions” [2]. Similarly, CISCO reported 50% increase in productivity, albeit in a specific scenario [3].

Meanwhile, the fast-paced development of AI has generated some concern about the impact of this technology on humans and its implications for education [4], job losses [5], and other human aspects. This has fueled both public and scientific debate [6].

But, how are these tools perceived by developers? And, why do some people advocate for their use, while others argue for avoiding them? In this article, we aim to answer these questions by considering the public perception of these tools

¹Putting aside the advances in program synthesis, which is a topic at least 10 years old [1]

based on discussions on Twitter. We provide some quantitative results from our study and then detail four concrete narratives that we believe are representative. In short, we read many Twitter discussions so you do not have to. Let’s dig in!

II. DISTILLING NARRATIVES FROM TWITTER DISCOURSES

We have distilled 39k tweets about Copilot and Tabnine into a sample of 331 tweets. We focused on Copilot and Tabnine because they were the first two popular code-generation tools with the largest volume of related tweets. Our 331 tweets had a wide reach and were representative of the narratives generated by users of these tools. We analyzed this sample to understand the promises and concerns of these tools as reported by their users. Before detailing the methodology behind our study, it is important to mention why we used Twitter. Previous studies have shown that developers use Twitter to stay current with recent technological trends [7]. Specifically, Sharma and colleagues [8] report that the 3rd most popular category discussed by software developers on Twitter is exactly New Releases, i.e., “tweets announcing the release of a new software version, tool, etc”². Indeed, Copilot averaged 1,097 Tweets per month before its release on VS Code, after that, it averaged 1,175 Tweets per month which is a 7% increase in Tweets. Tabnine averaged 246 Tweets per month before its release on VS Code, after it averaged 329 Tweets per month which is a 34% increase in Tweets. Since we are interested in understanding software developers’ reactions to the first generation of AI code assistants, Twitter data proved valuable because it often included early adopters of these assistants as well as of other technologies³. Twitter is a starting point for this line of research and leaves room for future studies of other platforms, such as Stack Overflow and Reddit.

We used quantitative and qualitative approaches to analyze the public English discourse to build the narratives around AI code generation tools. We first retrieved the tweets of interest

²The first one is Article and Multimedia Sharing (Tweets sharing articles, blogs, tutorials, or videos related to software development), and the second is Technical Discussions (Tweets discussing some technical issues related to software development).

³https://blog.twitter.com/en_us/a/2016/new-research-8-ways-early-tech-adopters-use-twitter

using the Twitter API, querying it with keywords “*tabnine*” and “*github copilot*”⁴. We limited our search to retrieve tweets only published after the public releases of Tabnine (November 2018) and GitHub Copilot (June 2021). We obtained 39,361 tweets about GitHub Copilot (25,516) and Tabnine (13,845). We identified tweets in 81 different languages, including English (24,692), Spanish (18,99), German (479), etc. Out of the 39,361 tweets, 64.29% were posted by users who had associated location information in their profile. Figure 1 presents a map with the locations of these users. The map indicates that most tweets were from North America, Europe, and India. The countries with the most tweets were: Pakistan, USA, India, UK, and Canada (in decreasing order).




Fig. 1. A map with locations of Twitter users for those users who had their location available in our dataset.

We first filtered out 14,669 retweets and non-English tweets. Given that English Tweets are the majority of the data and that English is regarded as the *lingua franca* of computing⁵, we assume the dataset sufficiently applies to the general public. Then, we identified the most significant tweets and qualitatively analyzed the narratives behind them. We consider a tweet to be significant if it belongs to the 1% of tweets that have the largest number of retweets, the largest number of replies, or the largest number of likes. In short, by significant, we mean tweets that have reached a large audience. This yielded 255 tweets about Copilot and 138 tweets about Tabnine. Table I overviews our final dataset.

The resulting codebook is a list of themes associated with each tweet. For instance, there are tweets describing performance issues with the AI-based code generation tools and tweets mentioning the low quality of the code suggested by the tools. Therefore, we organized these two codes (performance and quality) as examples of a larger category called *Concerns*. Similarly to these, we identified *Promises* as another category that encompasses the codes associated with the benefits of the tools, including usefulness, efficiency, integration with other software development tools, etc.

In addition, we identified the *Goal* of each tweet. In this case, tweets might be informative, provide recommendations,

⁴These were the tools publicly available when we collected the data in November 2022. Other tools were available either in beta or had recently been launched.

⁵https://en.wikipedia.org/wiki/English_in_computing

be funny, or be an advertisement about one of the tools. For instance, the following tweet recommends those writing PowerShell scripts to adopt GitHub Copilot: “*If you write #PowerShell scripts, you must check out #GitHub Copilot in #VSCode.*” Tweets marked as advertisements were those posted by the companies that build the AI tools and the people who work for these companies. Note that the categories are orthogonal, i.e., one tweet might be informative and, at the same time, describe a concern with a particular AI-based code generation tool. In fact, most tweets belonged to more than one category.

After creating our codebook, one of the authors coded the 393 tweets. We removed from our sample 57 tweets whose goal was to advertise the tools and 5 tweets that were not related to Tabnine or GitHub Copilot. In the end, we qualitatively analyzed 336 tweets⁶. Out of these 393 tweets, only 30 occurred after Twitter was acquired. This indicates that the change in Twitter’s ownership had at most a minor impact on the narratives we identified.

III. NARRATIVES

By creating the codebook and associating codes to Tweets, we were able to determine common narratives within the dataset while ensuring a level of consistency in our coding. Using the 31 unique codes from the codebook, the themes revealed a mixed sentiment towards AI code generation tools. Developers appreciated their productivity benefits and broad applicability, but also expressed concerns about code quality, security, and ethical implications. These tools were seen as potentially transformative for software development practices, with the fear of AI replacing human developers considered unwarranted. We then used our codebook and the narratives to assemble recommendations for developers in Section IV.

A. Promises: useful code suggestions

In general, several tweets mentioned the tools provide suggestions that are very accurate and useful (“*really good at predicting useful completions*”, “*crazy good*”, “*Pretty scary how accurate it is at predicting your code, especially after a while, it gets better at predicting your code.*”). These tweets often also reflected how *impressed* users are with the capabilities of the studied AI tools. For instance:

“*I was so surprised how good it [Tool] is*”

“*I’ve been trying [Tool] an auto-completion tool for VS Code that blew my socks off. Sometimes it seems like it’s just reading my mind*”

“*Been trying out the [Tool] with C# and @xunit in VS2022. It really is pretty mind blowing.*”

A related promise is associated with the broad applicability of AI-based code generation tools. For instance, we found tweets mentioning different programming languages (Java,

⁶Our final codebook with categories and codes, our final dataset, and the script we used to retrieve the data are available online at https://osf.io/xerz2/?view_only=ac60108ce22241b1bb9b51f40c781d86

TABLE I
DESCRIPTIVE STATISTICS OF OUR FINAL TWEETS DATASET.

	COPILOT	TABNINE
USERS	Number of unique users	4634
	Top 5 countries	PAK, USA, IND, GBR, CAN
	Occupations of top 15 users	IT Services, Bloggers, Developers, Freelancers, Academics, Software Practitioners, Cofounders
		1890
		USA, IND, GBR, DEU, CAN
		IT Services, Bloggers, Developers, Developer Relations, Freelancers, Software Practitioners, Cofounders
TWEETS	Number of Tweets	255
	Min Number of Likes	1
	Mean Number of Likes	137
	Max Number of Likes	6247
	Min Number of Retweets	2
	Mean Number of Retweets	21
	Max Number of Retweets	1788
	Min Number of Comments	0
	Mean Number of Comments	6
	Max Number of Comments	186

Flutter, Dart, Python, and Rust) and frameworks (Drupal, React.js, etc.).

B. Beyond code suggestions: new work practices

Given the quality of the suggestions and their applicability, the AI tools *promise* to impact developers' work in different ways. First of all, they allow developers to save time during development⁷ (“*I feel like I've been able to write faster*”, “*it really saves me tons of time.*”, “*Saves me tons of typing.*”, “*the piece of software that most improved my productivity in the last months.*”, “*a noticeable reduction in workload.*”). AI tools also influence developers' work styles since some users reported one important side benefit of using these tools: better code documentation. This is illustrated with the tweets below:

“I've noticed that a nice secondary consequence of using [Tool] is that it's causing me to leave better comments because of the need to make sure the AI understands what I'm trying to do.”

*“By far the greatest benefit of using [tool] so far is I now don't have to be forced to document my code. I actively *want* to write great comments, because when I do, I get the dopamine hit of a good [Tool] suggestion.”*

Last but not least, we found tweets discussing how AI-based code generation tools can facilitate the *learning* of programming languages and frameworks. The following tweet illustrates this point:

“I'm using [Tool] like a private tutor to learn Haskell. Here it is telling me how to translate Python dicts to Haskell's maps. And of course it knows how to build a Markdown table. Pretty close to magic”

The user who posted the above tweet has a Ph.D. and has also previously worked for Google and Meta. But, we found

⁷Our methodology does not allow to make any claims about the impact of these tools on developers' productivity [9].

supportive tweets about AI-based code-generation tools from those with little programming experience, too. For example, we identified tweets with the hashtag #100daysofcode, which is used by someone who accepted the challenge to learn to code, improve one's code skills, by programming, and tweeting, during 100 days⁸. In this context, the following tweet is the beginning of a thread about how an AI tool could be used by beginner developers:

“Experienced devs like [Tool], but it is even more important for beginners!”

People learning to code often have ideas they want to implement, but spend a huge amount of time on simple things: reading data from files or visualizing it.

[Tool] can help”

C. Technical, and legal concerns

In parallel to the narratives about the promises of AI-based code generation tools, there was a narrative about the *concerns* that arise when using these tools. These are technical and legal ones.

Technical concerns are about the quality of the code suggestions by the AI tools because either they are incorrect (“*should never keep the output verbatim or trust it*”, “*this [Tool] completion ends where a \n should be escaping is hard*”, “*don't forget [Tool], an ML model that can't code*”) or insecure (“*I'd love it if you didn't suggest vulnerable code snippets like the [example]*”, “*Careful! [Tool] doesn't always generate the most secure code.*”). Whenever posting about the incorrect suggested code snippets, users often posted screenshots or short videos capturing their use of the tool.

There is one major legal concern discussed in the tweets. As explained, AI-based code generation tools are trained with code from millions of open-source projects. However, we

⁸<https://www.100daysofcode.com>

found tweets mentioning the possible unethical aspects of these tools, for instance, because *“People’s source code ... is feed into these machines then turned into a paid service without these original creators even agreeing to this let alone getting compensated.”*

The following tweet makes a parallel between computer-generated images by systems like DALL-E and AI code-generation tools:

“This most likely violates licenses of artworks similar to how code was violated by” [Tool.] [Twitter-handle] is doing a class action lawsuit against them for violating licenses, do with that information what you want.”

The above tweet also mentions a lawsuit against a major technology company, which is a popular concern identified in the tweets. Some developers even “tested” AI-based tools to find out whether their code was inappropriately suggested by these tools. For instance, the following tweet even included screenshots of such a test:

[Tool] with “public code” blocked, emits large chunks of my copyrighted code, with no attribution, no LGPL license. (...) My code on left, [Tool] on right. Not OK.”

Both technical and legal concerns led many Twitter users to either abandon or not adopt AI-based tools. For instance, a developer reported the following about the technical aspects of the tool he tried:

“So I finally got around to trying [Tool.] My overall experience: it can be remarkable at times, but for the most part, it slowed me down and made my work more stressful. I felt a significant sense of relief when I uninstalled it.”

Legal concerns, however, not only led users to abandon the tools, but also led them to discourage *others* to adopt these tools.

“Any employer who allows their employees to use [Tool] should be aware they might be sued for copyright intrusion - and anybody putting code on [Tool] should be aware [Tool] will ignore your license and try to steal it to make \$\$\$. Shameful.”

According to a Twitter user, this legal concern was exacerbated by a US-centric view of tool developers, i.e., we observed complaints about US companies and their narrow view of privacy, contrasting it with the European Union (EU) and its privacy law (GDPR):

“The internet isn’t just US, why does everyone forget the EU exists. gdpr exists for a very good reason”

D. AI impact on software development

Another narrative we identified in the tweets was about the future of software development work given the impact and evolution of AI in software development tools. This narrative spans three aspects.

First, many tweets mentioned the integration of voice interaction with IDEs so that a developer can speak her intentions

and the IDE, with the AI-based code tool, would generate the requested code. Additional tweets noticed the opportunity to increase accessibility and inclusion in software development, as exemplified in the following tweet:

“This [voice-based interaction with code generation] will be an amazing #Accessibility tool for my limited mobility friends AND any of you with carpal tunnel, broken wrists, etc.”

We note that most of the tweets about voice-based interaction were inspired by the beta launch of a feature in this context.

The second aspect concerns the future of AI-based tools. In this case, we identified tweets that speculated about potential new development tools to be built inspired by AI. For instance, supporting pull requests and code reviews, command-line interface, scripts, etc. As an experienced developer tweeted:

“We can’t be far off AI being able to automate tedious refactorings of giant codebases. It’s a shame we mostly went through the Python 2-3 migration a few years before it would be feasible to have some slightly more advanced version of [Tool] do the whole thing for us”

Finally, the last aspect includes AI anxiety [6]: concerns about being replaced by automated AI tools. The following tweet expresses this idea:

It’s really scary how good that shit is getting.” “Man ... seeing how [Tool] automates coding it really IS a question if coders are even going to be needed at all in a few years :P

It’s really scary how good that shit is getting.”

IV. RECOMMENDATIONS FOR PRACTITIONERS

Software engineering is a continuously changing field. Rather than perceiving AI code assistants as a threat or a chore (another new technology to install/learn/use), we believe practitioners should see these tools as an opportunity. While we found a narrative criticizing the quality of AI suggestions, there is a contrasting narrative that captures how users were also impressed by the suggestions. Previous research [1], [10] on AI code assistants indicates that even bad suggestions can be beneficial as they provide a “starting point” for developers to complete their tasks.

At this point, it is unclear why there is a difference between the good and bad suggestions, i.e., whether that is related to the professionals using the tools, the quality of the prompts, the programming language used, a combination of these elements or something else. We recommend that users experiment with these tools to see for themselves if these AI tools are useful *to them*, given their industry, individual use case, and satisfaction with the generated suggestions. In short, *Give your AI coding assistant a try.*

It is also important to manage expectations instead of blindly assuming that, because the tool worked well for

someone else, it will also work for me. In other words, *Decide it for yourself*. Finally, one should not discount the fast-paced development of AI technologies, which suggests that tools rapidly improve and might prove useful after just a couple more months, so *Be patient with your AI coding assistant*.

Some of the narratives we identified suggest that AI-based tools will make software development more inclusive and accessible, especially through voice-based interfaces. New AI-based tools are already being created, extending the tools we consider in this paper, as well as tools that address other aspects of the software development landscape, see for instance [11], [12]. AI assistants also have the potential to influence the learning process of software developers who constantly need to learn new technologies. Arguably, in simple projects, these tools might allow a single software engineer to develop the entire project without the need to hire a specialized engineer with expertise in, for instance, a particular type of database. In this case, our recommendation is *To be on the lookout for new AI coding assistants*.

An issue, with existing AI assistants, is that there is no way to distinguish generated code from human-written code [13]. Yet, retaining traceability is important: which assistant was used to generate the code, which prompt was used, was the generated code inspected and/or modified by a human, etc? This information, if currently maintained, is tracked manually. But, without this information, the provenance of the code-base quickly deteriorates. We recommend that practitioners annotate their code with comments containing traceability information, i.e., *Document your AI coding sessions*.

Besides the tools coming out of large companies, there are already open-source alternatives like GPT-Code-Clippy⁹. We expect more such alternatives in the future. These open-source tools are also likely to be easier to review and use in commercial contexts and are less likely to face legal concerns.

Finally, it is all too easy to entertain the notion of AI assistants that come to replace human engineers. AI assistants are indeed highly sophisticated tools, and they stand to become increasingly important in our profession. However, engineering robust software requires complex human-in-the-loop tasks. So, we consider the replacement fear to be unwarranted. So, put your fears aside, *Make your AI coding assistant your friend* and see what this amazing technology has to offer!

V. FINAL REMARKS

AI code generation tools have great potential to impact the software development process. We extracted tweets about two of these tools, GitHub Copilot and Tabnine, to identify the different narratives being discussed about these tools. We extracted 39k tweets made from June 2021 to December 2022, filtered them, and analyzed the 331 tweets that have reached the largest audience, measured by the number of likes, retweets, and comments. Our analysis yielded different narratives describing the promises of these tools, the technical and legal concerns around their usage, and the potential impact

⁹<https://github.com/CodedotAI/gpt-code-clippy>

of AI on the software development process. A limitation of this study is our sample size of the tweets, which can affect the generalizability and comprehensiveness of our recommendations. We conclude our paper with recommendations for practitioners about how to go about adopting these tools so as to benefit their careers.

We are currently analyzing other sources of information (StackOverflow, Reddit, etc) to find out whether the narratives around AI-based code assistants on these platforms are similar to the ones we identified on Twitter.

ACKNOWLEDGMENTS

The authors would like to thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for the Discovery Grant RGPIN-2014-04870 and funding reference number GR023171. The first author would like to thank the National Council for Scientific and Technological Development (CNPq) for grant numbers 420406/2023-9 and 442779/2023-2.

REFERENCES

- [1] P. Vaithilingam, T. Zhang, and E. L. Glassman, "Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models," in *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: <https://doi.org/10.1145/3491101.3519665>
- [2] "ML-enhanced code completion improves developer productivity," <https://ai.googleblog.com/2022/07/ml-enhanced-code-completion-improves.html>, accessed: 2022-12-12.
- [3] "Learn how tabnine enables "pair programming" with ai," <https://blogs.cisco.com/developer/tabninepairprogramming01>, accessed: 2022-12-12.
- [4] Y. N. Harari, "Reboot for the AI revolution," *Nature*, vol. 550, no. 7676, pp. 324–327, Oct. 2017. [Online]. Available: <https://doi.org/10.1038/550324a>
- [5] J. Manyika, S. Lund, M. Chui, J. Bughin, J. Woetzel, P. Batra, R. Ko, and S. Sanghvi, "Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages," McKinsey Global Institute, Tech. Rep., 2017, accessed: 2022-12-13.
- [6] J. Li and J.-S. Huang, "Dimensions of artificial intelligence anxiety based on the integrated fear acquisition theory," *Technology in Society*, vol. 63, p. 101410, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0160791X20300476>
- [7] L. Singer, F. Figueira Filho, and M.-A. Storey, "Software engineering at the speed of light: How developers stay current using twitter," in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: Association for Computing Machinery, 2014, p. 211–221. [Online]. Available: <https://doi.org/10.1145/2568225.2568305>
- [8] A. Sharma, Y. Tian, and D. Lo, "What's hot in software engineering twitter space?" in *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2015, pp. 541–545.
- [9] N. Forsgren, M.-A. Storey, C. Maddila, T. Zimmermann, B. Houck, and J. Butler, "The space of developer productivity: There's more to it than you think," *Queue*, vol. 19, no. 1, p. 20–48, mar 2021. [Online]. Available: <https://doi.org/10.1145/3454122.3454124>
- [10] J. D. Weisz, M. Muller, S. Houde, J. Richards, S. I. Ross, F. Martinez, M. Agarwal, and K. Talamadupula, "Perfection not required? human-ai partnerships in code translation," in *26th International Conference on Intelligent User Interfaces*, ser. IUI '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 402–412. [Online]. Available: <https://doi.org/10.1145/3397481.3450656>
- [11] A. A. Sawant and P. Devanbu, "Naturally!: How breakthroughs in natural language processing can dramatically help developers," *IEEE Software*, vol. 38, no. 5, pp. 118–123, 2021.
- [12] M. Pradel and S. Chandra, "Neural software analysis," *Commun. ACM*, vol. 65, no. 1, p. 86–96, dec 2021. [Online]. Available: <https://doi.org/10.1145/3460348>

- [13] "Apparently i'm a robot." <https://mobile.twitter.com/JanelleCShane/status/1601729685385535488>, accessed: 2022-12-13.