
IT’S ABOUT TIME: THE COPILOT USAGE REPORT 2025

THE TEMPORAL AND MODAL DYNAMICS OF COPILOT USAGE

A PREPRINT

Beatriz Costa-Gomes Sophia Chen Connie Hsueh Deborah Morgan
Philipp Schoenegger Yash Shah Sam Way Yuki Zhu
Timothé Adeline Michael Bhaskar Mustafa Suleyman Seth Spielman

Microsoft AI

December 2025

ABSTRACT

We analyze 37.5 million deidentified conversations with Microsoft’s Copilot between January and September 2025. Unlike prior analyses of AI usage, we focus not just on what people do with AI, but on how and when they do it. We find that how people use AI depends fundamentally on context and device type. On mobile, health is the dominant topic, which is consistent across every hour and every month we observed—with users seeking not just information but also advice. On desktop, the pattern is strikingly different: work and technology dominate during business hours, with “Work and Career” overtaking “Technology” as the top topic precisely between 8 a.m. and 5 p.m. These differences extend to temporal rhythms: programming queries spike on weekdays while gaming rises on weekends, philosophical questions climb during late-night hours, and relationship conversations surge on Valentine’s Day. These patterns suggest that users have rapidly integrated AI into the full texture of their lives, as a work aid at their desks and a companion on their phones.

1 Introduction

AI models are increasingly embedded in many aspects of modern work and life, with at least 10% of the adult population engaging with AI on a weekly basis (Chatterji et al. 2025). While much of the early work on AI usage emphasized the use of AI for productivity related tasks (Brynjolfsson, Li, and Raymond 2025), recent work has shown a rise in non-work related tasks. While we have an understanding of “what” people do with AI, we know less about when and how they do it. Analyzing how and when people use Copilot is critical for understanding the current economic and social impacts of AI, and what the future trajectory of AI adoption may look like.

Generative AI has diffused more rapidly than any other technology in history (Bick, Blandin, and Deming 2024; AI Economy Institute 2025). However, diffusion unfolds at multiple scales, spreading both through the economy and people’s daily lives (Rogers, Singhal, and Quinlan 2014; Wejnert 2002; Hagerstrand 1968). Only by understanding how people relate to these technologies can we truly build humanistic AI. In this large-scale analysis of Copilot usage we focus not just on what people do with AI, but on how and when they do it.

Recent work drawing on large-scale chat data has primarily focused on work use cases. For example, OpenAI (Chatterji et al. 2025) analyzed 1.1M sampled ChatGPT messages (May 2024–July 2025) with privacy-preserving, prompt-defined classifiers to label work status, intent, topic, and O*NET (a database of occupations) activity. They documented a rise in non-work messages (53% → 73% year over year), a dominance of practical guidance, seeking information, and writing (about 77% of conversations, with writing at about 40% of work use in June 2025), and an intent mix led by asking (51.6%) over doing (34.6%) and expressing (13.8%). Similarly focusing on work, Anthropic (Handa et al. 2025) drew on over 4M Claude.ai conversations (December 2024–January 2025) to look at AI usage at work. They found that usage of Claude is concentrated in IT/technology (~50%), followed by creative/cultural (~20%) and business/finance/customer service (~15%) use cases. In general, adoption is broad but shallow across occupations, and augmentative modes with necessary human guidance are more common than automation (57% vs. 43%). Automation is concentrated in content generation and coding “fix-and-retry”, while augmentation is more common in front-end development, professional communication, education, and validation.

Previous work from Microsoft analyzing 200,000 conversations from an earlier iteration of Copilot by Bing finds that the most common work activities that AI assists with are gathering information and writing (Tomlinson et al. 2025). Further work has also documented domain-specific AI use among students, programmers, and clinicians (Ammari et al. 2025; Klemmer et al. 2024; *Stack Overflow* 2025; Mandal et al. 2025), with specific use patterns emerging for different groups. Looking at online forum data, Zao-Sanders (2025) reported a shift toward emotional/support use cases, with technical troubleshooting declining.

Despite this progress in understanding AI usage in the context of work tasks and productivity, two basic dimensions of use remain underexplored: how modality (mobile vs. desktop) and time dynamics shape usage. There is already a rich literature on how technology relates to and is integrated into daily life in general (Silverstone, Hirsch, and Morley 1992). Personal computing and mobile phones are notable instances of technology development and integration of the last few decades, and each iteration has impacted how and when users engage, access and shape information through their devices (Kamvar et al. 2009; Booth et al. 2023; Garg and Sengupta 2020). Prior work has not fully analyzed how AI use differs by device type or how patterns evolve temporally across the day or year.

We do not claim that the complex relationship between people and AI can be fully revealed by a classification of conversation logs. However, we do think that taking a look at how and when people use AI can provide insight into the social integration of the technology (Böhmer et al. 2011). We extend prior work by teasing out differences in AI usage by device type and time. We believe that this disambiguation enables a deeper understanding of people’s relationship with AI.

This paper is based on 37.5M deidentified conversations randomly sampled on a daily basis between January and September 2025 from Copilot. This allows us to analyze how AI usage has evolved over both the course of the year and within each day, across both mobile and desktop devices. These patterns illustrate that Copilot is utilized as a productivity tool when users are sitting at their work desktops, yet shifts into a conversational partner on questions of philosophy, health, and wellness when users are on their phones at night.

2 Classifying Copilot Usage

In line with industry best practice, Microsoft analyzes Copilot usage metrics in a privacy-preserving way, by running machine-based classifiers over deidentified conversation logs that are automatically scrubbed of any personally identifiable information (PII) (Tamkin et al. 2024; Chatterji et al. 2025). All data used in this paper was processed in accordance with Microsoft’s privacy policy (Microsoft 2025). Our machine-based classifiers are “eyes-off”, i.e., no human researcher sees the content of a conversation.

The data used in this analysis is from a random sample of user interactions with Copilot from January 7 to September 25, 2025. We exclude all enterprise traffic and users that have signed in on a commercial or educational enterprise account. The number of daily conversation samples was generally consistent throughout the collection period. We sample around 144,000 conversations per day from January 7 to September 25, 2025 (≈ 260 days). Conversations are sampled randomly and thus are representative of Copilot usage. Sampled conversations include coarse location (country), date, time, and the application on which the interaction occurred. Information about the application being used allows us to classify the usage as taking place either on a mobile device or a personal computer/desktop. Each sampled conversation is labeled by machine-based classifiers that assign the conversation a *topic* and an *intent* tag, or in the cases of ambiguous messages and low classifier certainty a no-topic/no-intent tag. These “No Topic”/“No Intent” conversations are mostly single or short conversations for which we cannot determine the user’s goal or subject of conversations (e.g. “Hi Copilot”) at sufficient certainty. These kinds of vague, short, exploratory conversations were common early in the year but diminished significantly over the study period.

Topics are the subjects of conversation. While there are many sets of possible topics, our list of topics adapts widely used Microsoft systems and third-party taxonomies for classifying web search activities to conversational AI (Private Advertising Technology Community Group 2024) (see Table 1).

Intents are the “jobs to be done”, or the tasks that the user would like to accomplish with Copilot. The intents for Copilot were developed via user research to capture tasks that are common on conversational AI systems. For our list of intents, see Table 2.

Each conversation is assigned a topic and intent independently. In total, this yields a system of ~ 300 topic-intent pairs. For example, “Seeking Advice” is an intent, while “Relationships” and “Work and Career” are topics. A user can be seeking advice about relationships or seeking advice about work and career, which would be captured by these respective topic/intent pairs. While far from exhaustive, the combination of topics and intents provides an overview of Copilot usage.

Table 1: Conversation topics and descriptions

Topic	Description (what the conversations are about)
Art and Design	Art, design, and creative projects.
Autos and Vehicles	Cars, motorcycles, and other vehicles.
Beauty and Fashion	Beauty and fashion.
Books and Literature	Books, literature, and reading.
Business and Industry	Business or industry such as strategy, planning, and entrepreneurship.
Education	Education, learning, and school. Excludes conversations about specific school subjects that have their respective topics.
Entertainment	Movies, music, TV shows, and other forms of entertainment. Excludes conversations about video games, and other forms of games which should be in Games.
Food and Drink	Food, cooking, and beverages.
Games	Video games, role-playing games, and other forms of gaming.
Health and Fitness	Health, fitness, and medical concerns.
Hobbies and Leisure	Hobbies, crafts, and leisure activities. Excludes conversations about Games, Books and Literature, Entertainment, Travel, Sports and other hobbies/leisure activities that have an assigned topic.
Home and Garden	Home improvement, gardening, and DIY projects.
Internet and Social Media	Internet, social media. Excludes conversations covered in other topics, like News and Current Events.
Language Learning and Translation	Learning a natural language or simple translation requests. Excludes conversations about programming languages, which should be in Programming.
Local	Nearby events, restaurants, institutions, attractions.
Math and Logic	Mathematics and logic.
Money	Money, personal finance, and investments.
News and Current Events	News, current events of regional, national, or global interest, and politics.
Personal Growth and Wellness	Mental health, personal development, self-improvement, and self-help.
Pets and Animals	Pets, animals, and wildlife.
Programming	Computer programming.
Relationships	Relationships, family, dating, and love.
Religion and Philosophy	Religion, spirituality, and philosophy.
Science	Science and scientific research.
Society, Culture, and History	Society, culture, and history. Excludes conversations about news and current events.
Sports	Professional sports.
Technology	Technology, gadgets, and software.
Travel	Travel, vacations, and destinations.
Work and Career	Work, career, and professional development.
No Topic	A "No Topic" conversation has a vague or unclear subject.

Table 2: Conversation intents and descriptions. Bold corresponds to the name used for the intent.

Intent	Description
Creating Text or Images	Creating text, imagery, code, or any content from scratch.
Editing Text or Images	Editing text, imagery, code, or any content from an original source.
Entertainment	Using the AI to pass-time, play games, have undirected chit-chat types of conversations. In Entertainment conversations, often the conversation itself is the goal.
Getting Feedback or Advice	Getting feedback, advice, emotional support, or brainstorming on anything NOT related to Shopping and Product Research.
Information Summary	Asking AI to summarize content, extract key points. If the summarization task is to help a user learn a new skill or related to schoolwork, it should be in the Learning category; if it is about summarizing products, it should be in the Shopping and Product Research category.
Learning	Helping the user with schoolwork or learning a new skill.
Planning	Help with planning tasks, travel, or events; organizing one’s time for increased efficiency.
Searching for Information	Searching for general information about a topic, business, person, and/or place. Seeking information related to products or services that one might purchase, including product comparisons and reviews, should be in the Shopping and Product Research category. Questions related to schoolwork or learning a new skill should be in the Learning category.
Shopping and Product Research	Searching for information about a product or service. Assisting with shopping decisions, especially learning about products, categories of products, or comparing products.
Technical Support	Help with computers, electronics, and other devices. Providing technical support for software, hardware, or other technical issues.
No Intent	A “No Intent” conversation is a short exchange between the user and the AI for which we cannot determine the user’s goal(s).

3 How and When People Use Copilot

3.1 The General Overview

We look at Copilot usage by ranking topics and intents: the most common topic of conversation for that time is marked as rank 1, and the least common as rank 30. Intents are ranked 1–11. For topic-intent pairs, we consider the rank of the conversation for each pair, resulting in ranks ranging from 1 to 330.

In Table 3, we show the top 5 topics and top 5 intents across the entire dataset. We find that the most common topics of conversation on Copilot are “Technology” as well as “Work and Career”, suggesting strong productivity-aimed uses of Copilot. However, conversations about “Health and Fitness” also rank highly, indicating that people also turn to Copilot for more personal questions. “Searching” is the most common intent across conversations, followed by “Advice”.

We find substantial differences between how AI is used on different types of devices, and there has been more volatility in usage patterns on desktops and laptops than on mobile devices, as seen in Figures 1 and 2. The most frequent topic-intent pairs appear more stable on mobile; “Health and Fitness” / “Searching” is unchanged as the most frequent pair,

Table 3: General top 5 ranks for Topic and Intent.

Rank	Topic	Rank	Intent
1	Technology	1	Searching
2	Work and Career	2	Advice
3	Health and Fitness	3	Creating
4	Language learning and translation	4	Learning
5	Society, Culture, and History	5	Technical Support

and four pairs remain within the top four for the majority of the year. While the stability in usage is an interesting observation, we think that dominance of “Health and Fitness” on mobile, regardless of time of day or date suggests that users engage with Copilot as a confidant for personal topics and as a companion for personal improvement. This rise of advice-seeking conversations further highlights a growing user trust in Copilot, as individuals increasingly view it not only as a source of information but as a reliable source of advice. On personal computers, Copilot usage has been more volatile than on phones; 20 distinct topic-intent pairs entered the top 10 rank over the course of the year, showing that relative use of Copilot is changing month-to-month at a high rate. This is compared to 11 distinct topic-intent pairs for Copilot mobile use over the year.

On desktop, we can see how several intents are paired with the “Work and Career” topic, which does not feature at all on mobile. Given that our dataset explicitly excludes enterprise-authenticated traffic, this finding highlights the permeability of the boundary between professional and personal computing. Interestingly, “Art and Design / Creating” holds ranks 3 and 4 on desktop for two months, and then disappears from the top 10 while it consistently remains in the top 10 on mobile, with a drop only in September. This pattern might be related to popular public interest around image generation at this time in 2025. In the summer months, we observe certain subjects related to school work declining in prevalence, such as “Language learning and Translation” and “Science”, with “Entertainment” increasing in rank.

Overall, the data suggests that, as might be expected, Copilot users reach for their phones for the topics that are more personal or pique their interest beyond work. The consistency is nonetheless marked. However, to provide a complete picture, we must also examine the temporal dynamics of usage.

3.2 Temporal Dynamics of Copilot Usage

In order to further visualize the differences between January and September, we calculate the delta (difference in rank between two months) for the top 10 topics for every hour, averaged across every day (see Figure 3).

The biggest difference is how “Programming” is much more common in January than in September, with the largest variance towards the latter being “Society, Culture, and History”. This illustrates how, throughout the year, the center of gravity of Copilot usage has shifted from purely productivity-focused conversations towards more social topics. This likely reflects a dual dynamic: the broadening of habits among existing users, and the democratization of the user base as mainstream adopters—who may have less technical priorities than the developer-heavy cohort of early January—joined the platform. This includes, for example, having conversations that range from understanding pop culture to historical events to international politics.

It’s not just about changes across the year, but also from month to month. For example, two consecutive months that showed a large variation in ranks between them in the previous analysis (in Figures 1 and 2) were June and July 2025. For a more in-depth breakdown between these two months, see Figures 4 and 5 for a detailed hourly average of June, and Figures 6 and 7 for how that shifted in comparison to July, on desktop and mobile respectively.

By analyzing hourly patterns across devices, we observe three distinct modes of interaction: the workday, the constant personal companion, and the introspective night.

The Workday (8 a.m. to 5 p.m.): On desktop, the data mirrors the standard economic workday. Between 8 a.m. and 5 p.m., “Work and Career” overtakes “Technology” for the top rank, while education and science-focused conversations rise significantly compared to nighttime hours. This professional focus is further highlighted by the inverse pattern on mobile: “Entertainment” shows a consistent U-curve, remaining high in the evenings and late nights but dropping substantially during business hours. Interestingly, we also observe a rise in “Travel” during these active hours (see Figure 8), suggesting that users may be utilizing their work tools to plan trips or prepare for commutes.

The Constant Companion: While the workday dictates the rhythm for most topics, one subject defies temporal patterns. On mobile, “Health and Fitness” remains the most frequent topic across every hour of the day, invariant to




Figure 1: Monthly top 10 ranks for topic-intent pairings on desktop. The lines correspond to when the pairing is present in consecutive months.




Figure 2: Monthly top 10 ranks for topic-intent pairings on mobile. The lines correspond to when the pairing is present in consecutive months.




Figure 3: Average hourly rank difference for the top 10 ranks in January vs September, across all devices.




Figure 4: Top 10 average rank of conversation topic per hour, for the month of June, on desktop.




Figure 5: Top 10 average rank of conversation topic per hour, for the month of June, on mobile.




Figure 6: Average hourly rank difference for the Top 10 ranks on desktop, in June vs July.




Figure 7: Average hourly rank difference for the Top 10 ranks on mobile, in June vs July.

the time. This suggests a device-specific usage pattern where the phone serves as a constant confidant for physical well-being, regardless of the user’s schedule.

The Introspective Night: Copilot usage shifts dramatically when the workday ends. As the night deepens, users turn to AI to ask the big questions of life that are not always easy to answer—or perhaps to ask them at a time when others are not awake. “Religion and Philosophy” shifts up in rank during the depths of the night through dawn. This late-night introspection is visible across devices, though with monthly nuances; in June, for instance, distinct introspective topics pierced the top 10 that did not appear in July: “Religion and Philosophy” on mobile and “Personal Growth and Wellness” on desktop.




Figure 8: Average hourly rank for Religion and Philosophy and Travel, for all devices, across all months.

The idea of looking into daily patterns is not just about the hours but also about the cadence of days. When observing the daily average per month, we found two particularly interesting patterns.

First, on a direct productivity versus leisure comparison, we found that we can pinpoint which days were weekdays by looking at the comparison between “Programming” and “Games”, as in Figure 9. Further solidifying the usage pattern

effects across the day in line with traditional work schedules, “Programming” conversations are more common than “Games” during the weekdays, and vice versa during the weekend, mirroring each other.




Figure 9: Average daily rank of Programming vs Games conversations, in August. The boxes correspond to the weekends.

Second, we see that our users trust Copilot with questions close to their hearts: in the month of February (see Figure 10), our users had a rise in conversations about “Personal Growth and Wellness” in the run up to Valentine’s Day, with “Relationships” spiking on the day itself.




Figure 10: Average daily rank of Relationships and Personal Growth and Wellness, in February.

4 Discussion

Our analysis of 37.5 million conversations reveals that Copilot is no longer just a productivity tool; its use is shaped by the temporal and physical context of the user. By disentangling seasonality, daily rhythms, and device-level differences, we move beyond the monolithic view of “AI usage” to reveal a technology that has integrated into the full texture of human life. The contrast between the desktop’s professional utility and the mobile device’s intimate consul-

tation suggests that users are engaging with a single system in two ways: a colleague at their desk and a confidant in their pocket.

This bifurcation has significant implications for the design of generative AI. The industry has largely treated the “chatbot” as a uniform experience across endpoints. However, our finding that mobile users prioritize health and fitness—regardless of the hour—indicates that the mobile form factor signals a shift toward personal conversations and self-improvement. This suggests the need for interfaces that are context-aware, differentiating not just in UI, but in personality and capability. A desktop agent should optimize for information density and workflow execution, while a mobile agent might prioritize empathy, brevity, and personal guidance.

Furthermore, the temporal rhythms we observed underscore the generality of Large Language Models, their ability to sync with human circadian rhythms in a way no previous technology has. The rise in philosophical and existential queries during the late night suggests that as the sun sets, the user’s need shifts from external productivity (‘doing’) to internal reflection (‘being’).

Ultimately, these patterns paint a picture of rapid and deep social integration. Users have tacitly agreed to weave AI into the fabric of their daily existence, turning to it for code reviews at 10 a.m. and existential clarity at 2 a.m. This is not merely a story of adoption, but of adaptation. The data suggests that we are not just using AI to do our work faster; we are using it to navigate the complexities of being human, one prompt at a time.

References

- AI Economy Institute (Nov. 2025). *AI Diffusion Report: Where AI is Most Used, Developed, and Built*. Tech. rep. Accessed: 2025-11-17. Microsoft Research. URL: <https://www.microsoft.com/en-us/research/group/aiei/ai-diffusion/>.
- Ammari, Tawfiq et al. (2025). “How Students (Really) Use ChatGPT: Uncovering Experiences Among Undergraduate Students”. In: *arXiv preprint arXiv:2505.24126*.
- Bick, Alexander, Adam Blandin, and David J Deming (2024). *The rapid adoption of generative AI*. Tech. rep. National Bureau of Economic Research.
- Böhmer, Matthias et al. (2011). “Falling Asleep with Angry Birds, Facebook and Kindle: A Large Scale Study on Mobile Application Usage”. In: *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*. ACM, pp. 47–56.
- Booth, Frederick et al. (2023). “A mental health and well-being chatbot: user event log analysis”. In: *JMIR mHealth and uHealth* 11, e43052.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025). “Generative AI at work”. In: *The Quarterly Journal of Economics* 140.2, pp. 889–942.
- Chatterji, Aaron et al. (2025). *How people use chatgpt*. Tech. rep. National Bureau of Economic Research.
- Garg, Radhika and Subhasree Sengupta (2020). “He is just like me: a study of the long-term use of smart speakers by parents and children”. In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4.1, pp. 1–24.
- Hagerstrand, Torsten (1968). *Innovation diffusion as a spatial process*.
- Handa, Kunal et al. (2025). “Which economic tasks are performed with ai? evidence from millions of claude conversations”. In: *arXiv preprint arXiv:2503.04761*.
- Kamvar, Maryam et al. (2009). “Computers and iPhones and mobile phones, oh my! A logs-based comparison of search users on different devices”. In: *Proceedings of the 18th international conference on World wide web*, pp. 801–810.
- Klemmer, Jan H et al. (2024). “Using ai assistants in software development: A qualitative study on security practices and concerns”. In: *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 2726–2740.
- Mandal, Soumik et al. (2025). “Utilization of Generative AI-drafted Responses for Managing Patient-Provider Communication”. In: *npj Digital Medicine* 8.1, p. 591.
- Microsoft (2025). *Microsoft Privacy Statement*. <https://www.microsoft.com/en-gb/privacy/privacystatement>. Accessed: November 2025.
- Private Advertising Technology Community Group (2024). *The Topics API*. <https://github.com/patcg-individual-drafts/topics>. Individual draft proposal. Accessed: 2025-10-24. URL: <https://github.com/patcg-individual-drafts/topics>.
- Rogers, Everett M, Arvind Singhal, and Margaret M Quinlan (2014). “Diffusion of innovations”. In: *An integrated approach to communication theory and research*. Routledge, pp. 432–448.

- Silverstone, Roger, Eric Hirsch, and David Morley (1992). "Information and communication technologies and the moral economy of the household". In: *Consuming Technologies: Media and Information in Domestic Spaces*. Ed. by Roger Silverstone and Eric Hirsch. London: Routledge, pp. 15–31.
- Stack Overflow (2025). *2025 Stack Overflow Developer Survey*. Stack Overflow. URL: <https://survey.stackoverflow.co/2025/ai/> (visited on 10/14/2025).
- Tamkin, Alex et al. (2024). *Clio: Privacy-Preserving Insights into Real-World AI Use*. arXiv: 2412.13678 [cs.CY]. URL: <https://arxiv.org/abs/2412.13678>.
- Tomlinson, Kiran et al. (2025). "Working with AI: measuring the applicability of generative AI to occupations". In: *arXiv preprint arXiv:2507.07935*.
- Wejnert, Barbara (2002). "Integrating models of diffusion of innovations: A conceptual framework". In: *Annual review of sociology* 28.1, pp. 297–326.
- Zao-Sanders, Marc (Apr. 9, 2025). *How People Are Really Using Gen AI in 2025*. Harvard Business Review. URL: <https://hbr.org/2025/04/how-people-are-really-using-gen-ai-in-2025> (visited on 10/14/2025).