Beyond AI Training: How Social and Experiential Learning Shape AI Adoption in the Workplace

Riya Sahni
riya.sahni@cs.columbia.edu
Columbia University
New York City, New York, U.S.A

Lydia B. Chilton
chilton@cs.columbia.edu
Columbia University
New York City, New York, U.S.A

Abstract

While organizations continue to invest in AI tools like M365 Copilot, little is known about how individual employees engage with these technologies once deployed. This study examines M365 Copilot adoption behaviors among a group of 10 experienced users across many industries. Findings reveal a strong preference for informal learning methods over structured training. Even though 9 out of 10 participants acknowledged that formal training for Copilot tools would be useful, 7 out of 10 stated that they ignored the Copilot onboarding videos provided to them, citing reasons such as time constraints, preference for self-guided learning, or reliance on external resources like ChatGPT. No participants used formal training as their primary learning method. Instead, experiential learning (trial and error, 8 participants) and social learning (peer discussions, 6 participants) emerged as dominant learning strategies. We discuss opportunities for promoting social learning of AI tools in the workplace.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**.

Keywords

AI productivity tools; social learning; discoverability; human cognition; technology adoption

1 Introduction

"In the future, your work is going to be defined not by the tasks you perform, but by the creativity you unleash", said Satya Nadella, painting a vision that encapsulates Microsoft's revolutionary approach with Microsoft 365 (M365) Copilot. Today, as AI continues to reshape the modern workplace, millions of paid users rely on M365 Copilot, positioning it as a pivotal tool in the AI race.

Prior research has extensively documented the role of organizational readiness and environmental pressures on AI integration in professional settings [1, 3]. However, research on the primary contributing factors of successful AI adoption at the individual level remains scarce [26]. Given AI's rapid integration into workplace environments and its evolving role in augmenting human decision-making, the lived experiences of employees as active participants in shaping AI use remain less understood. This study provides new insights into how social and experiential learning shape grassroots-level adoption of AI productivity tools, like M365 Copilot, in public U.S. organizations.

Our central research questions are:

- **What did participants use M365 Copilot for, and what are their general sentiments?**

- **What are the primary learning behaviors and patterns observable among these users?**

Biases in AI adoption emerge, as employees selectively trust AI recommendations based on perceived reliability and social influence [2, 6]. The strategic implementation of AI in organizations further depends on adaptive learning and knowledge diffusion rather than top-down training programs [5].

Our study, based on in-depth interviews with 10 seasoned professionals across diverse industries, reveals three key findings:

- **A Distinct Preference for Non-Traditional Learning:** Despite recognizing the value of formal training, participants overwhelmingly favor self-directed, experiential, and social learning methods.

- **Discoverability as a Reactive Process:** Users primarily discover Copilot's capabilities through ad hoc, informal channels—such as observing colleagues in action and spontaneous online searches—rather than through structured training.

- **Efficiency-Confidence Gap:** Although most participants report high perceived efficiency gains from using Copilot, a significant gap exists in their confidence in mastering its advanced features.

These findings align with broader trends in AI adoption, where employees engage in an iterative learning process, continuously refining their understanding through workplace interactions and social cues [3, 33]. On one hand, Copilot liberates users from routine busywork, allowing them to focus on more creative and strategic tasks; on the other, the constant evolution of its features and the reactive nature of learning in the workplace contribute to a pervasive uncertainty about fully harnessing its capabilities. As AI continues to permeate modern organizations, understanding how individuals learn, adapt, and shape these tools will be crucial for maximizing their potential impact.

2 Related Work

2.1 Informal Learning and Social Learning

Informal learning is learning that happens naturally, outside of formal training, and is spontaneous, unstructured, and often self-directed [16, 24]. It can consist of solitary learning experiences such as trial and error and social learning experiences like peer influence. Social learning specifically refers to learning that occurs through the interaction, observation, and imitation of others [7]. Social learning, especially through peer-to-peer exchanges, discussions with colleagues, and shared problem solving have been observed as integral parts of informal learning in professional contexts [16, 22].

Additional studies further emphasize the social nature of workplace learning. For example, Lave and Wenger describe learning as something that happens within communities of practice, while Boud, Keogh, and Walker point out that much professional development occurs through everyday interactions and reflective practice [10, 20]. More recently, Billett has shown that workplace participatory practices play a key role in ongoing professional growth [9]. Together, these studies suggest that informal learning at work is not just about individual effort, but also relies heavily on social interaction and collaboration. Building upon these insights, our study examines the roles of social and experiential learning—the two key components of informal learning—in the adoption of AI tools in industry.

2.2 AI Adoption At Work

Recent studies demonstrate that generative AI can substantially enhance workplace productivity. Research by Brynjolfsson et al. and various science productivity studies reveal that tools like Chat-GPT help professionals complete tasks faster and with improved quality [12]. In one experiment, 453 college-educated professionals using ChatGPT produced higher-quality work, and Dell et al. found that consultants with access to ChatGPT generated more innovative ideas [15]. Similar benefits are observed in technical domains, where studies by Cui and Peng report a 26% increase in weekly coding tasks [13, 28]. In a different creative context, ReelFramer demonstrates how a human-AI co-creative system can help journalists translate print articles into engaging, narrative-driven video reels, highlighting AI's potential to support content retargeting through multiple narrative framings [18, 21, 30, 34].

At the same time, several studies highlight important limitations of AI support. Research on creativity support tools by Tao et al. and investigations into tasks requiring core human skills by Long et al. indicate that while AI can assist with specific tasks, it does not fully replicate human judgment or creativity [23, 31]. Dell et al. further caution that for complex tasks, reliance on AI may lead to incorrect recommendations, while Kim documents challenges users face when integrating AI outputs into their work processes [15, 19]. Moreover, Randazzo suggests that AI tends to impact specific work processes rather than transforming entire industries [29].

While a number of studies have examined AI adoption at the organizational level—such as Neumann, Guirguis, and Steiner's comparative case study of Swiss public organizations that explored structural approaches to managing AI innovation—there remains a notable gap in understanding AI adoption from the individual perspective [26]. Recent calls in the literature urge research that considers AI adoption from the viewpoint of individual citizens and workers [8, 27]. In response, our research focuses on how employees in the United States adopt AI tools like M365 Copilot in their daily work by examining informal and social learning behaviors—such as trial and error, peer discussions, and self-directed exploration. This approach bridges the gap between top–down organizational strategies and grassroots-level experiences, providing a more complete understanding of AI adoption and offering practical guidance for designing user-centered training and support systems.

2.3 Microsoft 365 Copilot Tools

M365 Copilot is a large language model (LLM)-backed, chat-based AI productivity tool integrated with widely used applications such as Microsoft Word, PowerPoint, Power BI, and Teams. As one of the leading AI productivity tools in the industry, M365 Copilot has millions of paid users and is endorsed by numerous global enterprises for its significant contributions to efficiency, productivity, and cost savings [4, 17]. Its widespread adoption provides a unique opportunity to observe trends in AI tool adoption across a heterogeneous workforce. Consequently, M365 Copilot serves as an ideal proxy for understanding the broader learning and adoption behaviors associated with AI productivity tools across industries. For this reason, our study examines how individuals across various industries adopt M365 Copilot to reveal overall insights of grassroots-level workplace AI adoption.

3 Methodology

3.1 Study Procedure

3.1.1 Research Design and Approach. Between September and December 2024, we interviewed 10 industry workers in a cross-sectional, qualitative study. The purpose was to explore their experiences with learning and using M365 Copilot tools at work. The semi-structured format allowed the interviewer to flexibly adjust the conversation—skipping, reordering, or deepening questions—so as to elicit richer insights from each participant. The interview guideline, provided in Appendix A, focused on the following three primary themes:

- **Usage and Experience:** Participants' daily interactions with M365 Copilot tools, including their learning strategies, feature discoverability, and both positive and negative experiences.
- **Perceptions:** Participants' perceptions of M365 Copilot, its potential capabilities, and the extent to which these capabilities are realized in practice.
- **Learning Preferences:** Preferred learning styles and methods employed by participants to acquire knowledge about new Copilot features and functionalities.

Each interview lasted between 30 and 45 minutes and was conducted in English via Google Meet. All sessions were recorded and transcribed verbatim. Participants were not compensated monetarily for their time. The study received exemption status from the institutional review board, ensuring adherence to ethical research guidelines.

3.1.2 Sampling Strategy. Participants were recruited through professional networks. We then employed a combination of purposeful and snowball sampling methods to capture a diverse range of perspectives and behaviors. Purposeful sampling ensured that participants were seasoned working professionals and active M365 Copilot users, providing a solid foundation of relevant experiences. Snowball sampling further enriched the sample by incorporating voices from a variety of industries and organizational roles. Given the diversity of job functions and organizational structures, this approach allowed us to capture a wide array of M365 Copilot use cases and learning experiences. For example, the adoption and training

approaches of an executive manager supported by personal administrators may differ significantly from those of a senior manager working without administrative assistance. This sampling strategy, therefore, facilitated an examination of how M365 Copilot adoption behaviors and experiences may align or vary across different job functions, even within the same industry.

3.1.3 Participant Criteria and Selection. Participants were required to be active M365 Copilot users, engaging with the tool at least a few times per week for a minimum of two months. This criterion was established to ensure that the initial novelty effect had subsided and that participants' usage behaviors reflected established habits rather than temporary excitement. Individuals who exclusively used other Copilot tools—such as GitHub Copilot or Power BI Copilot—were excluded from the study. Focusing on M365 Copilot users was a deliberate choice, as this subscription group represents the largest population among the Copilot tools and, consequently, offers a particularly rich dataset for understanding adoption and learning behaviors.

3.1.4 Participant Demographics. Participants were sampled from a variety of industries, including Energy and Chemicals, Media and Entertainment, Transportation, and Consulting. The participants represent a broad spectrum of roles, from highly technical users to executive management, and exhibit varying levels of prior exposure to AI tools. Table 1 details key demographic and usage characteristics, including each participant's industry, role or job title, years of industry experience, duration of M365 Copilot usage, and the frequency of its use.

The sample comprises 4 women and 6 men. All participants reported using Teams Copilot at work, while 3 also utilized Bing Copilot, 6 employed Copilot for Outlook, 2 used Copilot for Word, and one participant also engaged with Power BI Copilot. Beyond these specifics, additional details regarding exposure to complementary AI tools reveal that, with the exception of two participants (P2 and P3) who have received formal training and possess extensive AI/ML experience through academic or previous professional roles, the remaining 8 participants primarily encountered AI through platforms such as Copilot, Claude, and ChatGPT. Notably, P6 leads a technical writing consulting firm that leverages various AI tools—including Quillbot and Bark—beyond the Microsoft Copilot Suite. Furthermore, P5 mentioned that she and some of her peers use Amazon Q in conjunction with Copilot at work.

3.2 Qualitative Coding & Analysis

The transcriptions were manually cleaned and reviewed to eliminate transcription errors and filler words while still preserving the participants' sentiments. Given that our objective was to examine how participant responses align with established learning frameworks rather than to uncover entirely new ones, we adopted a primarily deductive approach for the qualitative analysis [11].

Table 2 represents a codebook that was developed based on predetermined learning framework categories, including:

- **Positive social learning:** instances where participants described collaborative learning experiences, knowledge sharing, and engaging in AI-related discussions with colleagues.

- **Negative social learning:** Instances characterized by resistance to knowledge sharing, limited access to peer discussions, or the dissemination of inaccurate information among colleagues.
- **Experiential learning:** Self-directed, hands-on experimentation with M365 Copilot to acquire new skills and insights.
- **Traditional learning:** Engagement with formal company-provided training resources such as internal documentation or onboarding videos.

Using the codebook as a guide, we systematically reviewed each transcription, sorting relevant participant quotes into the corresponding learning themes. To ensure coding accuracy and consistency, each transcript was reviewed twice. During this iterative review, ambiguous responses were either bolstered with additional context or flagged for further discussion. If sufficient context could not be found to confidently assign a quote to a specific category, that quote was excluded from the final analysis.

To enhance the credibility of our coding process, we maintained a detailed audit trail of coding decisions and conducted regular reflective sessions to address potential researcher bias. This systematic and reflective approach ensured that our analysis remained firmly grounded in the data and accurately reflected participants' experiences with M365 Copilot within the context of established learning frameworks.

4 Findings

4.1 General Use Cases, Sentiments, and Usability Barriers in M365 Copilot Adoption

After interviewing participants, we identified eight primary use cases for M365 Copilot and categorized each mention as positive or negative sentiment based on whether participants found Copilot helpful or frustrating. Figure 1 provides a breakdown of these mentions, illustrating both the most praised and most criticized functionalities.

4.1.1 What Participants Enjoyed About M365 Copilot. Among all 64 positive mentions, Figure 1 illustrates the most positively received functionalities were writing assistance (100% positive sentiment, 14 mentions), notetaking & summarization (70.5% positive, 31 positive out of 44 mentions), and information retrieval/explanations (56.3% positive, 9 out of 16 mentions). Other features, such as email/task prioritization (75.0% positive, 3 out of 4 mentions), meeting transcript verification (33.3% positive, 2 out of 6 mentions), scheduling/task generation (40.0% positive, 2 out of 5 mentions), multimodal capabilities (66.7% positive, 2 out of 3 mentions), and presentation assistance (16.7% positive, 1 out of 6 mentions), were noted but with lower positivity rates.

Writing assistance was the highest-rated use case, with zero negative mentions and 100% positive sentiment. Participants frequently used Copilot to refine emails and reports, adjust tone, and ensure professionalism. P3 noted, "I sometimes have the tendency to write more than I need to, and when I learned that it can do an effective summary for a wider audience, I basically started writing emails without paying a lot of attention to what I'm writing and then just dumping it into Copilot." Similarly, P4 emphasized the ease of adjusting tones, stating, "Now we can adjust the tones quite

Table 1: Participant Demographics (*P prefix in ID denotes "participant"*)

ID	Industry	Role	Experience (years)	Copilot Experience	Frequency of Usage
P1	Media & Entertainment	Senior Program Manager	27	3 months	Daily
P2	Energy & Chemicals	Digital AI Lead	12.5	3 months	Daily
P3	Energy & Chemicals	IT Digital Strategy Leader	27	1 year	Daily
P4	Energy & Chemicals	Well Engineer	20	8 months	Daily
P5	Transportation	Senior Manager Analytics	20	3.5 months	Weekly
P6	Consulting	Owner	22	4 months	Daily
P7	Energy & Chemicals	Program Manager	18	2 years	Daily
P8	Energy & Chemicals	Program Manager	38	1 year	Daily
P9	Consulting	Senior Manager	12	1.5 years	Weekly
P10	Energy & Chemicals	Country Head	26	2 years	Monthly

Figure 1: Distribution of Positive and Negative Sentiments in Copilot Use Cases

a bit just by selecting if it is a casual tone or more professional tone or if it needs to be concise."

Notetaking and summarization was the most frequently mentioned positive use case, with 31 positive mentions (70.5% of all mentions in this category). Participants valued how it allowed them to stay engaged in meetings without taking notes manually. P5 recounted, "I would say, the most memorable experience would be using Copilot to recap a meeting that I joined in 50 minutes late on a one-hour meeting." P8 highlighted Copilot's ability to track key discussions, saying, "I think if anything, the positive outcome has been that I'm able to track the volume of information in the meetings better because I can ask Copilot to go tell me what was said in this meeting…"

Information retrieval/explanations accounted for 9 positive mentions (56.3%), making it another widely used and appreciated feature. Participants cited its usefulness in surfacing relevant documents and definitions in real time. P8 described how it streamlined document searches: "I just started using it as a search feature and not just like searching the web, but searching SharePoint and searching [files]. I know I heard somebody say something in a meeting about some architecture. And so, I just put it in into Copilot and it goes

out and it says, 'Yeah, this was in this PowerPoint from this meeting on this date.'" P2 similarly noted, "As of right now, when I need to find information from internal [company] documents, I will use Copilot on a daily basis."

Overall, participants appreciated how M365 Copilot streamlined straightforward tasks like email writing and meeting summarization. These findings suggest that Copilot effectively meets both general and context-specific needs for routine tasks.

4.1.2 Issues That Participants Faced With M365 Copilot. While many participants found Copilot useful, Figure 1 also highlights 33 total negative mentions, indicating key areas of dissatisfaction. The use cases with the highest proportion of negative sentiment were presentation assistance (66.7% negative, 4 out of 6 mentions) and meeting transcript verification (66.7% negative, 4 out of 6 mentions). Notetaking and summarization (29.5% negative, 13 out of 44 mentions) and information retrieval/explanations (43.8% negative, 7 out of 16 mentions) had the highest absolute number of negative mentions, indicating that while these were the most used features, they were also among the most polarizing.

Table 2: Qualitative Analysis Codebook with Example Quotes

Code Name	Definition	Inclusion Criteria	Example Quotes
Positive Social Learning	Instances where participants reported gaining valuable knowledge about M365 Copilot tool(s) by observing, mimicking, or speaking with colleagues.	• Mentions learning about M365 Copilot through colleagues • Expresses appreciation for shared workplace knowledge about M365 Copilot	"One guy accidentally typed his Copilot prompt in the chat, and I saw it. I pinged him: 'Did that give you what you wanted?' He said yes and showed me, and I'm like, 'Son of a gun—alright…'"- P8
Negative Social Learning	Instances where participants found social learning ineffective.	• Complaints about limited knowledge-sharing resources • Resistance to knowledge sharing • Colleagues providing inaccurate info	"Copilot says 'you have one of 30 responses.' I knew that meant conversation length, but some people thought they only got 10 total tries—so they barely used it, afraid of running out." - P2
Experiential Learning	When participants participated in self-led or hands-on experimentation to learn about M365 Copilot.	• Self-learning through trial & error • Testing features without formal guidance	"I've kind of changed my thought patterns to say, I wonder can Copilot do that instead of going and trying to figure it out on my own. And so I figure out a way to ask Copilot and see if Copilot can do it." - P4
Traditional Learning	When participants use formal training resources like internal documentation or onboarding videos.	• Watching company-published trainings/tutorials • Reading official documentation	"There was another team with Enterprise AI who trained some of us as early adopters to use [Copilot] with Teams and they had a slack channel…"- P5

Presentation assistance and meeting transcript verification, each receiving 4 negative mentions (66.7%), were particularly criticized for inaccuracy. P7 expressed disappointment in Copilot's Power-Point assistance, stating, "The PowerPoint outputs felt like a generic Google answer, not something tailored to our work." Similarly, P10 mentioned issues with meeting transcript summarization, noting, "The Copilot transcript summary didn't quite capture the essence of what people were saying."

Notetaking and summarization had the highest absolute number of negative mentions (13 negative mentions, 29.5% of all mentions in this category). Several participants reported concerns about inaccurate or incomplete meeting summaries. P1 recalled, "There was one meeting. And I'm glad I looked at the notes relatively early because there was a whole section of a discussion that was really important. And it wasn't included in the Copilot notes at all." P5

stated that her team stopped using Copilot for technical discussions, saying, "We have stopped using [Teams Copilot] to capture anything for code reviews, or anything at all for technical now based on that one experience because we don't want to get it wrong."

Information retrieval/explanations, with 7 negative mentions (43.8%), was another polarizing feature, with users reporting irrelevant search results and difficulty retrieving past decisions. P2 recounted, "I asked Copilot to retrieve past meeting decisions, but it said it had no recollection. I manually found the transcript with the needed information."

Other use cases had lower but still notable levels of negative sentiment: scheduling/task generation (60.0% negative, 3 out of 5 mentions), email/task prioritization (25.0% negative, 1 out of 4 mentions), and multimodal capabilities (33.3% negative, 1 out of 3 mentions). P6 reported difficulties with automation workflows,

saying, "I had a workflow set up to automate table sharing, but it kept glitching." P5 noted an issue with Copilot's data processing, stating, "If the table name was called A_B, it incorrectly merged it into ABCD."

Overall, participants' primary issues with M365 Copilot centered on inaccuracies and inconsistencies in its more advanced features, such as presentation assistance and meeting transcript verification. These challenges suggest that while Copilot excels at streamlining routine tasks, its performance on complex or technical functions remains unreliable, possibly dissuading users from leveraging these use cases.

4.2 High Perceived Efficiency using M365 Copilot at Work but Low Confidence in its Mastery

Perceived Efficiency vs. Confidence of Mastery Ratings for M365 Copilot
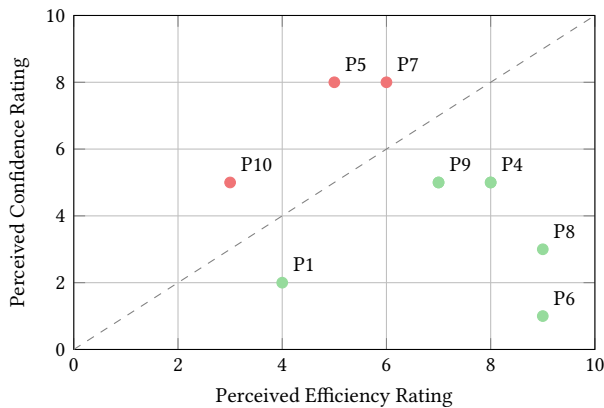
Figure 2: Scatter plot of Efficiency versus Confidence ratings for each participant. Points are labeled with the participant ID. Points above the reference line (where Confidence exceeds Efficiency) are colored red, indicating lower confidence relative to efficiency.

The data reveal a discrepancy between how much participants perceive M365 Copilot to benefit their efficiency at work and their confidence in fully utilizing its capabilities. 7 out of 10 participants rated Copilot's impact on their efficiency as 6 or above on a 10-point scale, indicating a strong perceived benefit. In contrast, only 20% reported high confidence (a rating of 7 or above) in their ability to leverage all of Copilot's features. This gap is illustrated by the scatter plot in Figure 2, where many points lie significantly below the $y = x$ dashed line—demonstrating that while users feel the tool boosts their efficiency, they remain uncertain about fully mastering it.

For example, P6 rated efficiency at 9 but expressed a confidence level of only 1, while P8 gave an efficiency rating of 9 and a confidence rating of 3. These cases underscore a recurring sentiment among many participants: despite acknowledging the practical benefits of M365 Copilot, there is a pervasive uncertainty about unlocking its full potential. Some participants, like P5, employed a

variety of learning approaches—including formal workshops, reading documentation, and engaging in peer collaboration—to familiarize themselves with Teams Copilot. Yet even she admitted, "I don't know what the new features are, but I think confidently, we have tried at least what was made aware to us."

Conversely, P8, who rated his efficiency highly but his confidence very low, remarked, "Using it in Teams just makes me, I don't want to say smarter, but it makes me more engaged and helps me understand more about what's going on." However, after attending one of Satya Nadella's keynotes, he noted: "There is so much more that can be done that I don't even understand yet." Further emphasizing the challenge, P7 commented,

> "Even the name changing of the software and the lineage of the applications is really difficult for people who aren't deeply engaged. It almost becomes an obsession to stay connected with what's happening."

Collectively, these insights suggest that while M365 Copilot is widely regarded as a powerful tool for enhancing work efficiency, its constantly evolving features and frequent updates present significant barriers to mastery.

4.3 Why Users Avoid Training and Feel They Underutilize M365 Copilot

Many users expressed that they sensed they were not making full use of its capabilities. However, this sentiment was often also paired with a sense of lack of motivation to formally explore additional features. Many users acknowledged that M365 Copilot had more to offer but were unsure whether they truly needed to expand their knowledge. As P9 put it,

> "If I need to do something and I can't find an answer, I try to look up videos and do it. But if I don't need anything, I don't know what else it can do. You don't know what you don't know."

Some participants recognized their own reluctance to explore Copilot further but cited time constraints and competing priorities as the primary reason for avoiding formal training. P1 admitted, "There was some [formal training] made available. I haven't used them.". P6 echoed his sentiment, saying "I haven't really had the time to go through it personally, but I'm hoping during Christmas holidays I might be able to play around more with it."

Another factor contributing to the underutilization was a rigid mental model of AI that framed Copilot as maladaptive tool. P7 framed this concept well, explaining,

> "if you're not going in the direction that feels like it's going to be successful,... you abandon that use case and you move on... you just say 'I'm not going to waste any more time with this... I'd rather do it manually'."

Participants noted that it was easy to get frustrated with Copilot quickly, with a couple of participants explicitly defining this as a 15-minute period of patience. Some users even mentioned abandoning features after one unsuccessful attempt instead of iterating or seeking guidance.

Based on these observations, it seems that underutilization is not necessarily due to a lack of Copilot's capabilities, but rather a combination of competing priorities, limited patience for experimentation, and a lack of perceived urgency to learn more. In many cases, users had access to training but did not prioritize it, reinforcing the idea that merely offering training may not be enough.

4.4 Social and Experiential Learning in AI Adoption
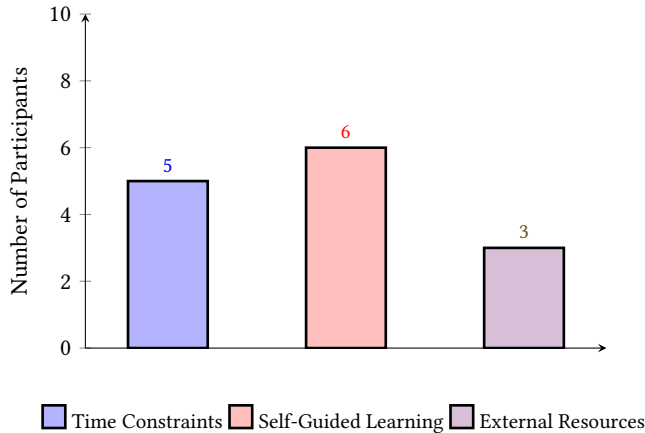
Figure 3: Bar chart displaying the percentage of participants citing time constraints (50%), self-guided learning (60%), or reliance on external resources (30%) as reasons for ignoring onboarding videos.

4.4.1 Participants Ignored Formal Trainings. Although 9 out of 10 participants acknowledged that formal training for Copilot tools could be beneficial, 7 explicitly reported ignoring official onboarding videos. Figure 3 shows that 6 participants mentioned ignoring formal training (e.g., watching videos and reading documentation) simply because they preferred self-guided, hands-on learning instead (e.g., trial and error, ad hoc exploration, and social learning). Time constraints were another key factor, with 5 participants citing lack of time as a reason for avoiding training. P10 stated, "There were onboarding classes and even some PowerPoint presentations on how Copilot can be used, but I simply did not have the time to review any of it."

Self-guided learning in the form of social learning also played a significant role in participants ignoring formal training. P4 mentioned teaching peers informally, saying, "... and before lunch was even over, I sent the job description out, and they're like, 'How did you do that?' I said 'it's Copilot.'...and then I showed them. They're like 'wow'...". Similarly, P5 found Slack discussions with other teams valuable:

> ""I realized we could share screenshots back and forth and say, 'that's how you do that', that's 'how you do a ticket integration...' we learned from each other by sharing our screen during some of the office hours..." - P5"

However, sometimes this self-guided learning was also motivated because of time constraints. P8 recounts:

> "I didn't spend a whole lot of time [trying to get Copilot to generate a PowerPoint presentation]... it was probably 15 minutes or so. It wasn't enough time to go and do [an internal] search to see if there was [formal training] out there to tell me how to do it."

Essentially, participants often intentionally bypassed formal training for Copilot because of time constraints and a strong preference for self-guided learning, which was also sometimes motivated by time constraints.

4.4.2 Discoverability and AI Adoption. Figure 4 illustrates the distribution of participants' primary learning methods: 60% relied on social/peer learning and 30% identified trial-based learning as their main approach. No participants preferred formal training, and one participant (P10) had minimal engagement (so minimal learning time) with the tool, since they delegated most Copilot tasks to two executive assistants.

Social learning played a significant role in Copilot adoption. P8 described learning a useful prompt by observing a colleague's mistake in a Teams chat:

> "One of the guys that was in the meeting thought he was trying to use Copilot but instead he wrote out his prompt in the chat and I saw it and I'm like okay, and then he realized what he'd done and then he moved it over and. So I pinged him on the side, and I'm like, 'Did that thing [the prompt] did it give you what you wanted?' and he's like, 'Yeah,' and he showed me some things and them I'm like, 'Son of a gun, all right...'"

Similarly, P1 became interested in Copilot after seeing colleagues using it: "I had peers, co-workers who were part of the pilot program. They would be on conference calls, running Copilot, and I was like, 'I want that.'" P5 highlighted how peer-driven experimentation encouraged feature adoption:

> "Our internal group learned from each other... It was peer pressure. Others were testing [Teams Copilot], so we were too. We were challenged to try new features every day. When we were on our own, like with Power BI Copilot, we weren't as motivated."

Without a peer network, P5 noted that learning Power BI Copilot felt isolating, underscoring the role of social learning in fostering engagement with AI tools.

5 Discussion

5.1 Preference towards Social Learning over Formal Training for M365 Copilot

Our study shows that professionals prefer social learning over formal training when adopting M365 Copilot. Informal learning theories and social learning theory explain that employees learn by trial and error, peer exchanges, and collaborative problem solving rather than through structured modules [7, 16, 24]. Studies by Livingstone, Lave and Wenger, Boud, Keogh, and Walker, and Billett also show that everyday interactions and reflective practices are central to
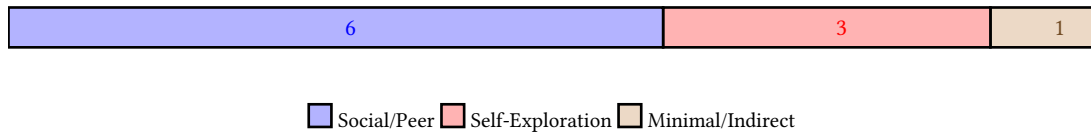
Figure 4: Stacked bar chart representing the distribution of participants by primary learning source: Social/Peer Learning (6), Self-Exploration (3), and Minimal/Indirect Exposure (1). Each bar segment is labeled with the number of participants.

workplace learning. These findings confirm that work learning relies on social interaction, not just indivduals [9, 10, 20, 22].

In AI adoption, the preference for social learning has important implications. Formal training may not fully address the challenges of rapidly evolving tools like M365 Copilot. Our findings suggest that users learn more effectively through direct interaction and observation. This grassroots approach helps bridge the gap between initial exposure and full mastery, which is an issue that formal training alone may not solve.

These results offer practical steps for organizations. Companies might create communities of practice or discussion forums where employees share tips and solve problems together. Informal workshops, peer-led sessions, or mentorship programs can also encourage workers to learn about M365 Copilot capabilities and use cases collaboratively.

5.2 Perceived Underuse of M365 Copilots for More Complex Tasks

Many participants expressed low confidence in their proficiency with M365 Copilot despite acknowledging its high benefits, and many believed they were severely underutilizing its capabilities. This gap between immediate perceived benefits and deep mastery has deep groundings in technology complexity theory. According to the Technology Acceptance Model [14], perceived ease of use is a key factor in technology adoption—if a tool is seen as too complex, users are less likely to adopt it fully. Similarly, the Unified Theory of Acceptance and Use of Technology [32] identifies "effort expectancy" as a critical determinant, where higher perceived complexity reduces the likelihood that users will invest the effort required to master the technology. Our findings, where users benefit from Copilot yet struggle to achieve deep mastery, align with these models by illustrating that the inherent complexity of the tool can hinder full utilization, despite its clear benefits.

Looking forward, these findings have important implications for the future of AI adoption at work. To bridge the gap between high perceived benefits and low mastery, AI tools like M365 Copilot need to be more transparent about their capabilities. Enhancements such as guided walkthroughs, real-time feedback, and contextual prompts can help users discover and utilize underused features.

In addition, our interviews revealed that some users are more inclined to experiment and "tinker" with Copilot than others. This aligns with Moore's (1991) Crossing the Chasm, which argues that a small group of early adopters play a critical role in learning and championing new technologies [25]. Organizations could identify these early adopters—our internal tinkerers—and empower them as champions or mentors. Their success stories and practical tips can then help encourage and guide their peers in mastering the tool.

6 Future Work

This section outlines research directions to improve AI tool adoption. We propose studies on interventions to boost user confidence, reasons users forgo training, and gaps between perceived and actual Copilot usage.

6.1 Potential Design Interventions to Improve AI Confidence

Many users lacked confidence in using M365 Copilot, not due to missing features but unclear expectations, misinformation, and lack of feedback. This uncertainty may lead to disengagement.

Future research could test design interventions, such as:

- **Real-time feedback and validation**
- **Gamification and confidence metrics**: integrating some sort of self-assessment features that allow users to track their progress in mastering the system without having to engage in formal training. Maybe a confidence dashboard or achievement-based learning approach might encourage users to engage more deeply with Copilot.

By designing and testing these interventions, future work could help determine which UI/UX strategies build user confidence to improve AI adoption.

6.2 Comparing Self-Reported Underutilization with Objective Usage Data

Users often felt they underutilized Copilot, though actual usage may differ. Some may use it effectively without realizing, while others overestimate engagement. Understanding this gap can inform training strategies.

Future research should quantify and compare self-reported vs. actual usage through:

- **Usage data analytics**: detailed personal tracking key metrics including frequency of Copilot interactions, diversity of features used, and session duration to determine a score that demonstrates how effectively users actually engage with the tool
- **Task-based evaluations**: designing controlled experiments where users are given specific tasks to complete using Copilot. By measuring their efficiency, confidence levels, and actual usage behaviors, researchers can identify gaps between perception and actual proficiency.
- **Longitudinal tracking**: studying how users' confidence and usage patterns evolve over time, particularly after targeted training interventions or UX updates.

7 Limitations

7.1 Industry Representation and Generalizability

Our sample included participants from a range of industries, including Media and Entertainment, Energy and Chemicals, Transportation, and Consulting. These industries were selected due to their varying levels of AI adoption and integration, allowing us to explore how different workplace enviornments influence Copilot usage and adoption. While this diversity strengthens our findings, it also presents a limitation in terms of generalizability. Some industries, particularly those with strict compliance regulations or less exposure to AI-driven tools, may exhibit different adoption behaviors that were not captured in this study. Future research should explore Copilot adoption across a broader range of industries, partiularly in sectors where AI implementation is more regulated.

7.2 Sampling Methodology and Bias

This study employed a combination of purposeful sampling and snowball sampling, which allowed us to gather insights from professionals actively using Copilot. However, these methods may introduce bias by overrepresenting users who are already engaged with AI tools and underrepresenting those who may have disengaged early. Additionally, with a sample size of only 10 participants, individual perspectives may not fully reflect broader organizational or industry-wide trends. To mitigate these concerns, we included participants from diverse backgrounds and roles, but future studies should incorporate larger-scale sureys or longitudinal studies to validate and expand upon these findings.

8 Conclusion

Our work identifies three key factors shaping Copilot adoption: a strong preference for experiential and social learning, the role of informal discovery in learning new features, and a persistent gap between perceived efficiency gains and confidence in mastery. These insights imply that organizations should look beyond traditional training methods. Embedding real-time feedback, contextual tips, or interactive modules directly within the tool may help bridge the gap between immediate productivity and long-term proficiency.

References

[1] Mohamed Gamal Aboelmaged. 2014. Predicting E-readiness at firm-level: an analysis of technological, organizational and environmental (TOE) effects on e-maintenance readiness in manufacturing firms. *International Journal of Information Management* 34, 5 (2014), 639–651. https://doi.org/10.1016/j.ijinfomgt.2014.05.002

[2] Saar Alon-Barkat and Madalina Busuioc. 2022. Human-AI interactions in public sector decision-making: "Automation Bias" and "Selective Adherence" to algorithmic advice. *Journal of Public Administration Research and Theory* (2022).

[3] Sulaiman Alsheibani, Yen Cheung, and Chris Messom. 2020. Winning AI Strategy: Six-Steps to Create Value from Artificial Intelligence. In *AMCIS 11*. https://aisel.aisnet.org/amcis2020/adv_info_systems_research/adv_info_systems_research/1

[4] Judson Althoff. 2025. *The value of AI: How Microsoft's customers and partners are creating differentiated AI solutions to reinvent how they do business today*. https://blogs.microsoft.com/blog/2025/01/28/the-value-of-ai-how-microsofts-customers-and-partners-are-creating-differentiated-ai-solutions-to-reinvent-how-they-do-business-today/ Microsoft Blog.

[5] H. Igor Ansoff et al. 2019. *Implanting Strategic Management* (3rd ed.). Springer International Publishing. https://doi.org/10.1007/978-3-319-99599-1

[6] David Antons and Frank T. Piller. 2015. Opening the black box of 'not invented here': attitudes, decision biases, and behavioral consequences. *Academy of Management Perspectives* 29, 2 (2015), 193–217. https://doi.org/10.5465/amp.2013.0091

[7] Albert Bandura. 1977. *Social Learning Theory*. Prentice Hall, Englewood Cliffs, NJ.

[8] Paul F. Langer Bernd W. Wirtz and Carolina Fenner. 2021. Artificial Intelligence in the Public Sector - a Research Agenda. *International Journal of Public Administration* 44, 13 (2021), 1103–1128. https://doi.org/10.1080/01900692.2021.1947319

[9] Stephen Billett. 2001. Learning in the Workplace: Strategies for Effective Practice. (01 2001).

[10] David Boud, Rosemary Keogh, and David Walker. 1985. *Reflection: Turning Experience into Learning* (1st ed.). Routledge. https://doi.org/10.4324/9781315059051

[11] Bonnie S Brennen. 2021. *Qualitative research methods for media studies*. routledge.

[12] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. 2024. Generative AI at Work. arXiv:2304.11771 [econ.GN] https://arxiv.org/abs/2304.11771

[13] Zheyuan Cui, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2024. *The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers*. Working Paper. SSRN. https://doi.org/10.2139/ssrn.4945566 Available at SSRN: https://ssrn.com/abstract=4945566.

[14] Fred D. Davis. 1989. Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly* 13, 3 (1989), 319–340. https://doi.org/10.2307/249008

[15] Fabrizio Dell'Acqua, Edward McFowland, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R Lakhani. 2023. *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality*. Technical Report 24-013. Harvard Business School Technology & Operations Management Unit.

[16] Michael Eraut. 2004. Informal learning in the workplace. *Studies in Continuing Education* 26, 2 (2004), 247–273. https://doi.org/10.1080/158037042000225245

[17] J. P. Gownder, James McQuivey, Rowan Curran, Camille Floyd, and Demi Starks. [n. d.]. *Build Your Business Case for Microsoft 365 Copilot*. https://www.forrester.com/report/build-your-business-case-for-microsoft-365-copilot/RES180016

[18] Anna Kantosalo and Hannu Toivonen. 2016. Modes for creative human-computer collaboration: Alternating and task-divided co-creativity. In *Proceedings of the seventh international conference on computational creativity*. 77–84.

[19] Jini Kim and Hajun Kim. 2024. Unlocking Creator-AI Synergy: Challenges, Requirements, and Design Opportunities in AI-Powered Short-Form Video Production. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.

[20] Jean Lave and Etienne Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.

[21] Vivian Liu, Han Qiao, and Lydia Chilton. 2022. Opal: Multimodal Image Generation for News Illustration. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 73, 17 pages. https://doi.org/10.1145/3526113.3545621

[22] D.W. Livingstone. 2001. Adults' Informal Learning: Definitions, Findings, Gaps and Future Research. (01 2001).

[23] Tao Long, Katy Ilonka Gero, and Lydia B Chilton. 2024. Not Just Novelty: A Longitudinal Study on Utility and Customization of an AI Workflow. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 782–803. https://doi.org/10.1145/3643834.3661587

[24] Victoria J. Marsick and Karen E. Watkins. 1990. *Informal and Incidental Learning in the Workplace*. Routledge, London.

[25] Geoffrey A. Moore. 1991. *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*. Harper Business, New York, NY.

[26] Oliver Neumann, Katharina Guirguis, and Reto Steiner. 2024. Exploring artificial intelligence adoption in public organizations: a comparative case study. *Public Management Review* 26, 1 (January 2024), 114–141. https://doi.org/10.1080/14719037.2022.204

[27] Irina Pencheva, Marc Esteve, and Slava Jankin Mikhaylov. 2020. Big Data and AI – A transformational shift for government: So, what next for research? *Public Policy and Administration* 35, 1 (2020), 24–44. https://doi.org/10.1177/0952076718780537 arXiv:https://doi.org/10.1177/0952076718780537

[28] Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590* (2023).

[29] Steven Randazzo, Hila Lifshitz-Assaf, Katherine Kellogg, Fabrizio Dell'Acqua, Ethan R Mollick, and Karim R Lakhani. 2024. *Cyborgs, Centaurs and Self Automators: Human-Genai Fused, Directed and Abdicated Knowledge Co-Creation Processes and Their Implications for Skilling*. Technical Report. SSRN.

[30] Johannes CP Schmid and Johannes CP Schmid. 2021. Narrative Framing: Storytelling, Structures, and Perspectives. *Frames and Framing in Documentary Comics* (2021), 201–274.

[31] Yanda Tao, Changseung Yoo, and Animesh Animesh. 2023. AI Plus Other Technologies? The Impact of ChatGPT and Creativity Support Systems on Individual

Creativity. In *Proceedings of the International Conference on Information Systems*.

[32] Viswanath Venkatesh. 2008. Technology Acceptance Model 3 and a Research Agenda on Interventions. *Decision Sciences* 39, 2 (May 2008), 273–315. https://doi.org/10.1111/j.1540-5915.2008.00192.x

[33] V. Venkatesh, F. D. Davis, and G. B. Davis. 2003. User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly* 27, 3 (2003), 425. https://doi.org/10.2307/30036540

[34] Sitong Wang, Samia Menon, Tao Long, Keren Henderson, Dingzeyu Li, Kevin Crowston, Mark Hansen, Jeffrey V Nickerson, and Lydia B Chilton. 2024. ReelFramer: Human-AI co-creation for news-to-video translation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–20.

A Methods

A.1 Interview Guidelines

Phase 1: Learnability & Onboarding

(1) Did you have specific onboarding or training to learn how to use Copilot?

(2) Through which mediums have you mostly learned how to use Copilot? (e.g., trainings in your workplace, workshops, online videos, social media, your own explorations, trial and error, etc.)

(2)1. Was it through any of these mediums?

(3) Have you had experience using other AI tools either at work or outside of your workplace (e.g., Chat-GPT, Gemini, Claude, etc.)?

(3)1. If yes, how long have you been using these tools? What have you used them for?

(3)2. How have your experiences with these other AI tools helped you understand how Copilot works (e.g., what it does, how it works, its limitations)?

(4) On a scale of 1 to 10, with 1 being not at all and 10 being completely, how confident do you feel that you truly maximize the capabilities and features offered by Copilot when you work?

(5) Do you and your colleagues at work ever share new information about how to use Copilot (e.g., best practices, new features, discovered limitations)?

(6) What has motivated you to incorporate Copilot into your work practice? (e.g., genuine curiosity, leadership expectations, efficiency improvements)

Phase 2: Last Memorable Bad Time Using Copilot

(1) Can you recall a time when you were excited to use Copilot at work to help you accomplish something, but it fell short of your expectations?

(2) In general, what would you say are the biggest limitations of Copilot that you have personally experienced so far?

Phase 3: Last Memorable Good Time Using Copilot

(1) Can you recall a time when you were excited to use Copilot at work and it met or exceeded your expectations?

(2) In general, what would you say are the greatest strengths of Copilot that you have personally experienced so far?

Phase 4: General Use of Copilot

(1) What are your most common use-cases for Copilot at work?

(2) On a scale of 1 to 10 (1 = not at all, 10 = extremely well), how well does Copilot help you accomplish your tasks on average?

(3) On a scale of 1 to 10 (1 = worse off, 5 = neutral, 10 = extremely beneficial), how has Copilot affected your efficiency at work?