

**WestminsterResearch**

<http://www.westminster.ac.uk/westminsterresearch>

**AI in Software Programming: Understanding Emotional  
Responses to GitHub Copilot**

**Eshraghian, F., Hafezieh, N., Farivar, F. and De Cesare, S.**

This article is © Emerald Publishing Limited and permission has been granted for this version to appear here: <http://westminsterresearch.westminster.ac.uk/>

The final, published version in Information Technology & People, DOI: 10.1108/ITP-01-2023-0084, 2024 is available at:

<https://doi.org/10.1108/ITP-01-2023-0084>

This manuscript version is made available under the CC BY-NC 4.0 licence

<https://creativecommons.org/licenses/by-nc/4.0/>

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners.

# AI in Software Programming: Understanding Emotional Responses to GitHub Copilot

Farjam Eshraghian

Westminster Business School, University of Westminster, London, UK

Najmeh Hafezieh

School of Business and Management, Royal Holloway University of London, Egham, UK

Farveh Farivar

University of Tasmania, Hobart, Australia and  
Management and Marketing School, Curtin University, Perth, Australia

Sergio de Cesare

Westminster Business School, University of Westminster, London, UK

## Abstract

**Purpose** - The applications of Artificial Intelligence (AI) in various areas of professional and knowledge work are growing. Emotions play an important role in how users incorporate a technology into their work practices. The current study draws on work in the areas of AI-powered technologies adaptation, emotions, and the future of work, to investigate how knowledge workers feel about adopting AI in their work.

**Design/methodology/approach** - We gathered 107,111 tweets about the new AI programmer, GitHub Copilot, launched by GitHub and analysed the data in three stages. First, after cleaning and filtering the data, we applied the topic modelling method to analyse 16,130 tweets posted by 10,301 software programmers to identify the emotions they expressed. Then, we analysed the outcome topics qualitatively to understand the stimulus characteristics driving those emotions. Finally, we analysed a sample of tweets to explore how emotional responses changed over time.

**Findings** - We found six categories of emotions among software programmers: challenge, achievement, loss, deterrence, scepticism, and apathy. In addition, we found these emotions were driven by four stimulus characteristics: AI development, AI functionality, identity work, and AI engagement. We also examined the change in emotions over time. The results indicate that negative emotions changed to more positive emotions once software programmers

redirected their attention to the AI programmer's capabilities and functionalities, and related that to their identity work.

**Originality/value** - Our study makes a timely contribution to the discussions on AI and the future of work through the lens of emotions. In contrast to nascent discussions on the role of AI in high-skilled jobs that show knowledge workers' general ambivalence towards AI, we find knowledge workers show more positive emotions over time and as they engage more with AI. In addition, this study unveils the role of professional identity in leading to more positive emotions towards AI, as knowledge workers view such technology as a means of expanding their identity rather than as a threat to it.

**Practical implications** - Overall, as organisations start adopting AI-powered technologies in their software development practices, our research offers practical guidance to managers by identifying factors that can change negative emotions to positive emotions.

**Keywords** Artificial Intelligence, GitHub Copilot, Emotion, AI programmer, Twitter

**Paper type** Research paper

Figure 1. The research process (Source: Author's own work)

Figure 2. Coherence analysis (Source: Author's own work)

Figure 3. The connection between topic modelling (first stage) and qualitative analysis (second stage) (Source: Author's own work)

Figure 4. Screenshot of a Twitter bio from our sample (Source: Twitter.com)

Figure 5. Screenshot of a developer's contributions on GitHub (Source: GitHub.com)

Figure 6. Categories of emotional responses towards Copilot (Source: adapted from Beaudry and Pinsonneault (2010))

Figure 7. Stimulus characteristics and emotional change (Source: Author's own work)

Table 1. The outcome of topic modelling and labelling topics (Source: Author's own work)

Table 2. The coding process in the second stage of analysis (Source: Author's own work)

Table 3. Emotions associated with the topic modelling results (Source: Author's own work)

Information Technology & People

## 1. Introduction

As the economy is becoming increasingly knowledge-based, production and services depend on knowledge-intensive activities that rely on scientific and technological advancements. The nature of work in knowledge-based systems has changed immensely due to the adoption of Artificial Intelligence (AI) in information processing and data mining (Davenport and Kirby, 2016). In the past decade, many industries have increasingly adopted AI in their processes in, for example, data entry automation, recommendation search engines, disease diagnosis, financial wealth management, prediction of staff resignation, transaction fraud detection, and online customer support (Davenport, 2018). The revolutionary aspect of AI-powered technologies is that these technologies can perform tasks mainly without human intervention (Benbya *et al.*, 2020; Davenport and Ronanki, 2018). These capabilities have led experts to believe that AI disrupts employment stability and the workforce by replacing at least some human jobs (Jaiswal *et al.*, 2022; Larivière *et al.*, 2017).

Multinational corporations' heavy investment in AI-powered technologies intensifies the AI invasion perspective (Jaiswal *et al.*, 2022). Ramaswamy's (2017) study suggests that 47% of US employment could disappear by 2033. However, some jobs (e.g., accounting and legal) have already felt the impact of AI-powered technologies (Frey and Osborne, 2017). Reducing the need for human workers goes beyond blue-collar or physical jobs, as knowledge workers relying heavily on collecting and processing data (e.g., financial advisers, insurance advisers, engineers) can be superseded by AI-powered technologies (Frey and Osborne, 2017).

Knowledge workers' primary capital is knowledge, and their work mainly includes problem-solving that requires divergent (i.e., exploring the issue widely with many possibilities) and convergent (i.e., focused action to find a well-defined solution) thinking (Reinhardt *et al.*, 2011). However, AI-powered technologies have more potent abilities and capabilities in non-

linear and creative (convergent and divergent) reasoning than knowledge workers do. Thus, AI-powered technologies have threatened the roles of some knowledge workers (Cao *et al.*, 2021; Weiler *et al.*, 2019).

Although there has been an ongoing discussion on digital technologies and their role in replacing occupations, the primary focus was on blue-collar jobs (Frey and Osborne, 2017). With technology advancements and the rise of AI, this discussion has been extended to Industry 4.0 and this category of work (Schneider and Sting, 2020). With the gradual penetration of AI into more skilled roles, there have been concerns and emotional responses such as fear, excitement, or mixed emotions among workers (van Hoek *et al.*, 2019; Maier *et al.*, 2020). Such emotional responses can have negative and positive consequences for workers in terms of control over their roles, job performance, and turnover intentions (Yu *et al.*, 2022). However, most published studies have largely ignored knowledge workers' emotions and viewed them as passive recipients of technological advancement rather than as active players who can react to, contribute to, and shape the future of work in a meaningful way. Emotions fuel actions and decisions (Mayer and Salovey, 1997) and thus there have been recent calls for the need to study emotions and digitalisation of work and jobs (Panteli *et al.*, 2022).

In business studies, scholars have attempted to characterise the hybrid of humans and AI-powered technologies (Peeters *et al.*, 2021; Rai *et al.*, 2019). However, we argue that describing the relations between users and AI-powered technologies is impossible without investigating users' emotions. Although the role of emotions in IT use has been acknowledged for the past three decades, adopting an emotion-based approach to study the digitalisation of work is still emerging and not well developed (Gkinko and Elbanna, 2022; Panteli *et al.*, 2022). Emotions drive people to act upon them, thus we argue that to better understand how knowledge workers might interact with AI-powered technologies we need to

understand how they feel about them. Through an emotion-based lens and topic modelling, this study aims to address two main questions:

- 1) What emotions do knowledge workers express in relation to AI-powered technology?
- 2) How do knowledge workers' emotions guide their interaction with AI-powered technologies, and how do these emotions change over time?

## **2. Literature review**

### *2.1. AI-powered technologies*

AI can be characterised differently from other types of information technologies. The power of predictability in AI-powered technologies allows AI to appear akin to the realm of human cognition (Hu *et al.*, 2021). Machine Learning within the field of AI uses data to train, learn, and improve its reasoning and predictions using computational algorithms (Ghahramani, 2015). The ability to learn through datafication has contributed to the development of AI-powered technologies that are capable of automating parts of jobs, augmenting humans in carrying out their work, or taking over people's jobs. Datafication refers to a technological trend turning many aspects of life into data (Dourish and Gómez Cruz, 2018). Despite such potential, scholars do not have a consensus on how AI can impact the future of the labour market. The extent to which a job becomes automated (i.e., AI takes over some tasks of a job), AI pairs with humans to augment their capabilities, or AI replaces humans, depends mainly on the nature of the job (Brynjolfsson and Mitchell, 2017; Peeters *et al.*, 2021). There are also some concerns regarding AI-powered technologies. The first concern is the data quality used to train AI algorithms (Baumer *et al.*, 2017). Although datafication of human behaviour empowers AI algorithms, it can also translate human biases into AI-powered technologies that are supposed to perform cognitive tasks such as decision-making

(Johansen *et al.*, 2021). Moreover, access to accurate data for training algorithms has been a significant challenge (Johansen *et al.*, 2021). The second concern relates to the complexity of AI algorithms that perplexes humans. The self-learning aspect of some AI-powered technologies is a black box to most knowledge workers. Thus, knowledge workers are less likely to outsource cognitive parts of their work (i.e., decision making and reasoning) to AI (Asatiani *et al.*, 2020). Finally, the complexity of these algorithms may worry knowledge workers in terms of how the algorithms might perform in situations that do not precisely match the patterns used in the training data (Asatiani *et al.*, 2020).

## 2.2 *AI-driven disruptions and labour market*

The learning nature of AI algorithms implies a sense of agency so that knowledge workers might perceive that AI has a choice, can perform independently, and replace some of their tasks (Ferràs-Hernández, 2018; Frey and Osborne, 2017). Thus, there have been ongoing discussions around AI and its roles in changing occupations (partially automating, augmenting, or replacing). While traditionally, it was argued that service jobs are less susceptible to automation, the advancement of AI allows knowledge-based services to be completed without human interference (Huang and Rust, 2018). AI-powered technologies can also impact specialised jobs involving pattern recognition (e.g., some legal and medical processes) (Frey & Osborne, 2017). For instance, Weiler *et al.* (2019) found that once the Self-Service Business Intelligence (SSBI) system was released to the market, data analysts reported concerns about their future of work since SSBI enables business users (e.g., sales, finance, HR departments) to explore data, forecast trends, and generate reports without a background in data mining and data analysis hence a potential threat to data analysts' jobs due to automation of data mining and reporting. Similarly, Brougham and Haar (2018) report



that even though introducing AI-powered technologies at work is not directly related to job insecurity, it might affect service workers' career planning and the future of work.

Regarding AI-driven decision-making disruptions, Cao *et al.* (2021) report that managers view AI-powered technologies as a threat that increases their vulnerability to making bad decisions and thus it could have negative impact on managers' attitudes in using AI. In addition, Maier *et al.* (2020) found that physicians were ambivalent about using AI-powered technologies for medical decision-making and resisted adopting them. On the contrary, van Hoek *et al.*'s (2019) study of radiologists, surgeons, and medical students showed that radiologists expected AI-powered technologies to increase their efficiency at work; thus, they tended to support the use of AI. In contrast, surgeons and medical students were sceptical about using AI-powered technologies. Moreover, in a recent study, Kim *et al.* (2021) explored social media conversations among the radiology community and found that there is both optimism (in terms of technological capabilities) and pessimism (more related to issues such as legal and ethical) in their discourse. Given the impact of AI in changing occupations and various emotions towards this technology, Gerli *et al.* (2022) noted that emotions are one of the psychological factors that can influence the skill development required for the adoption of smart technologies.

### 2.3 *Emotions and information technology use*

Affect-related concepts such as emotions related to technology or Information Technology (IT) use have a tradition in Information Systems (IS) discipline. The role of emotions has been recognised for the past three decades, initially through acceptance research. While the genesis of such research, investigating the link between emotions and IT use, is rooted in cognitive-based approaches, social cognitive models can hardly capture the full range of users' emotional reactions to IT (Beaudry and Pinsonneault, 2010; Stein *et al.*, 2015). A few studies reported on how emotions affect users' intentions, attitudes, and beliefs about

technology by focusing on some general emotions such as anger, fear, and anxiety (Venkatesh, 2000), as well as specific emotions such as enjoyment (Davis *et al.*, 1992), satisfaction (Bhattacharjee, 2001), and pleasure (Kim *et al.*, 2004). However, these studies were limited in scope as they only focused on how technology use triggered emotional responses. In addition, cognitive models emphasise sensemaking (Thompson, 2012), implying that perceptions shape the use of IT (Davis, 1989; Venkatesh *et al.*, 2003; Zhang, 2013). Although some authors have suggested that the emotional process is much like the perceptual process (Prinz, 2004; Tapollet, 2000), they are not similar (Deonna, 2006). A perception includes a sensemaking process resulting from sensory experience, so the perception of a new phenomenon (e.g., AI-powered technologies) results from the feature of the phenomenon (Niedenthal and Wood, 2019). In contrast, an emotion depends on unstable motivational sets of individuals who experience the phenomenon (Deonna, 2006). Thus, emotions drive people to act upon them, while perceptions do not necessarily exhort an action (Deonna, 2006). Furthermore, the perceptual process includes understanding and interpreting a phenomenon (in our case, AI-powered technologies), whereas emotions can be evoked by a phenomenon that is not consciously known (Niedenthal and Wood, 2019).

To shift from cognitive models and address the lack of an emotion-based approach in studying IT use, Beaudry and Pinsonneault (2010) developed a model suggesting four emotions toward using new technologies: challenge (e.g., excitement, hope), achievement (e.g., happiness, satisfaction), loss (e.g., anger, disappointment), and deterrence (e.g., anxiety, fear). They argued that people react to new technologies based on whether the new technology is a threat or an opportunity, and their perceived control over the outcomes created by the new technology. Beaudry and Pinsonneault (2010) also argued that challenge and achievement emotions (e.g., excitement and happiness) encourage employees to use the newly implemented technologies, while loss and deterrence emotions indirectly reduce IT

use. In relation to this, Kim and Kankanhalli (2009) also found that if users believe that the new system will be beneficial to them and that they will have the support of their organisation in learning and using it, they are less likely to have a negative emotion towards it to resist using it.

Based on Beaudry and Pinsonneault's (2010) model, Stein *et al.* (2015) suggested affective responses to technology use might be mixed rather than uniform, depending on how the technologies were introduced to employees and how much they participated in the change process (Stein *et al.*, 2015). These affective or emotional responses influence adaptation behaviours and strategies. When users show uniform emotions in response to the characteristics of technology, they use clear adaptation strategies. However, mixed emotions require combinations of different adaptation strategies (Stein *et al.*, 2015).

To summarise, introducing AI-powered technologies triggers different emotions (such as fear and threat) among workers in different industries and fields (Cao *et al.*, 2021; Kim *et al.*, 2021; Maier *et al.*, 2020). Yet, these nascent studies have examined neither the drivers or stimuli for such emotions nor how these emotions change over time. In addition to this gap, there have been wider calls for research on emotions and digitalisation of work as research in this area is still scarce, and "we need more research on the ways in which emotions are affected by and influence technology on an everyday basis" (Panteli *et al.*, 2022, p. 1685).

Therefore, in this study, we aim to explore knowledge workers' emotional responses towards a newly introduced AI-powered technology and understand how these emotions change through time.

### **3. Methods**

To explain the dynamics of emotions towards AI-powered technologies, we applied abductive reasoning to investigate computer programmers' feelings toward a new AI-powered technology named GitHub Copilot. Abductive reasoning involves developing

probable conclusions from certain known central premises and plausible minor premises (Sætre and Van de Ven, 2021). Abductive reasoning includes identifying unexplained observed experiences and generating and evaluating ‘hunches’ at individual and collective levels (Sætre & Van De Ven, 2021). Thus, abductive reasoning creates new ideas (Peirce, 1998) and helps researchers theorise the patterns extracted from data rather than merely generalise them (Vaast and Walsham, 2013). Therefore, this research strategy allows us to explain users’ emotions and recognise the general drivers triggering those emotions. Moreover, the version of the coding process in topic modelling that we used in the initial stage of data analysis is very close to the logic of abduction reasoning (Reichertz, 2007) as the coding process concentrates on searching for similarities and continuously creating new themes. Figure 1 portrays the process of our research.

-----  
Insert Figure 1 about here.  
-----

### *3.1 Research context*

To answer our research questions, GitHub Copilot was chosen as the empirical context. GitHub is the world’s most prominent host of source code. Launched in 2008, it allows software programmers to work together from anywhere to develop software programs. Two years after its launch, GitHub hosted over 1 million repositories of code written by software programmers worldwide. This number reached 200 million repositories including 39 million public repositories (GitHub.com, 2022) of software code in 2022. Since 2012, Microsoft has used publicly available source code on GitHub and finally acquired it in 2018. Microsoft introduced GitHub Copilot, its AI pair programmer, which writes software code using AI, in June 2021 for technical preview (Friedman, 2021). GitHub Copilot was officially launched as a subscription-based service for software programmers on 21 June 2022 (Dohmke, 2022).

GitHub claims this AI-powered technology can act as a software programmer to suggest code or develop an entire program (GitHub, no date). GitHub used public code on the GitHub platform to train Copilot's algorithm. Thus, any software programmer who published their code on the GitHub platform for free could participate in enhancing GitHub Copilot's capability in software coding. On 7 December 2022, GitHub finally announced that Copilot was available for businesses (Zhao, 2022).

The launch of GitHub Copilot triggered a wave of software programmers' reactions on social media, especially Twitter. As a micro-blogging platform, Twitter has been a powerful tool and platform for self-expression, sharing opinions and suggestions (Vaast *et al.*, 2017; Dindar and Dulkadir Yaman, 2018). Twitter stands apart from other platforms due to its emphasis on brevity and the immediate nature of tweets, which encourages users to openly share unfiltered opinions, feelings, and emotions. The character limit in each tweet, although recently removed, originally fostered this sense of immediacy and conciseness. In essence, the constrained character count compels users to distil the core of their messages and express unpolished emotions.

In addition, a "low level of reciprocal connections" (Shane-Simpson *et al.*, 2018, p. 277) on Twitter helps users share their feelings and promote their opinions without fear of judgment. Users are mainly followed by people whom they do not know personally. Such anonymity encourages them to share opinions and emotions more openly. Even these low-level reciprocal connections have raised Twitter to the top platform on which users share their work-related content and opinions (Van Zoonen, Verhoeven and Vliegthart, 2016).

In the context of software development and releases, Twitter has played a pivotal role, more than other platforms, for users to express their early opinions and initial reactions to newly released software. These early opinions and reactions have provided invaluable insights into the user experience for software developers (Guzman, Alkadhi and Seyff, 2017). Twitter's

affordances of facilitating raw and unfiltered expressions of emotions by users have significantly got the attention of developers in order to consider real-time feedback from this platform in refining and patching their software (Williams and Mahmoud, 2017). Despite Reddit being also used to share opinions about software, the platform hosts more in-depth discussions which do not necessarily express emotions and feelings about newly released tools and applications.

Thus, due to the above-mentioned capabilities of Twitter in facilitating unfiltered emotions, we selected Twitter as our data collection context and empirical setting. In addition, since GitHub Copilot was initially only available to GitHub users, who primarily consist of programmers, it was reasonable to infer that individuals commenting on this tool on social media, particularly Twitter, were likely programmers themselves.

Considering the context of Twitter, where each tweet is limited to 280 characters, it would be hard for the users of GitHub Copilot to express their feelings in detail in one tweet. To address the issue, firstly, we collected all the threads' tweets in which even one tweet included both keywords: 'GitHub' and 'Copilot'. Secondly, we collected data over a period of 11 months. This extended period of data collection helped us investigate if emotions evolved over time before the launch of Copilot's subscription-based version.

### *3.2 Data collection and analysis*

We used the Twitter academic Application Programming Interface (API:v2) to collect tweets related to GitHub Copilot. We collected 107,111 tweets that included both 'GitHub' and 'Copilot' keywords between 2 July 2021 and 1 June 2022. We intentionally ignored the tweets posted on the first three days of the release (29, 30 June and 1 July 2021) as the early tweets were about the news of the release or the excitement that descended after a few days. Our data analysis consists of three main stages.

### 3.2.1 First stage

Due to the size of the collected data, 107,111 tweets, it was impossible to use traditional approaches to code and analyse the textual data of the tweets (Baumer *et al.*, 2017; Hannigan *et al.*, 2019). Thus, using the Python programming language, we adopted the topic modelling method, a data-driven, computationally intensive analysis method for theory building (Berente *et al.*, 2019). The lexical framing in topic modelling approach is closely aligned with the grounded theory approach in explaining patterns in data by existing terminologies and concepts in the literature and generating theories (Berente *et al.*, 2019). Topic modelling is a statistical, exploratory text-mining technique that uses machine learning algorithms to discover topics and grounded conceptual relationships in textual data, even when researchers are oblivious of these existing topics in the text (Hannigan *et al.*, 2019). Topic modelling also allows researchers to reveal latent categories or dimensions of a phenomenon (Miranda *et al.*, 2022). In management studies, topic modelling is a common method utilised to organise, analyse, and make sense of online textual data that is considered extensive and unstructured (Hannigan *et al.*, 2019).

Following Hannigan *et al.* (2019), we used the Latent Dirichlet Allocation (LDA) technique as a topic modelling method. This generative probabilistic model classifies textual data and explains why some parts of the data are similar. The main advantage of LDA to other text mining approaches is that LDA does not need keywords to analyse textual data (Blei *et al.*, 2003). Thus, we uncovered themes (topics) that would have been tedious to obtain with a traditional qualitative thematic coding approach. The basic idea behind LDA is that a text is a random mixture of latent topics, where topics are characterised by specific words distributed across the text (Blei *et al.*, 2003).

Another popular technique for topic modelling is BERTopic (Grootendorst, 2022; Jeon, Yoon and Sohn, 2023). This technique, based on neural networks, uses pre-trained language models such as Bidirectional Encoder Representations from Transformers (BERT) (Jeon, Yoon and Sohn, 2023) to link different parts of a text to pre-learned topics. One of the advantages of this technique is its capability to manage noise (i.e., outlying and irrelevant text and documents) in the data (Sánchez-Franco and Rey-Moreno, 2022; Jeon, Yoon and Sohn, 2023).

Nonetheless, due to the nature of our data (i.e., limited character tweets with the keyword of “GitHub Copilot) and filters being applied by us to reduce noise (these filters are explained in the next few paragraphs), from this perspective, BERTopic did not seem to benefit our analysis more than LDA.

BERTopic and LDA categorise different pieces of text (i.e., documents, tweets, etc.) into different clusters (i.e., themes or topics), nonetheless, in BERTopic, bags of words do not directly represent generated topics, which consequently makes it challenging for researchers to interpret the topics (Grootendorst, 2022). On the other hand, the representation of topics using bags of words in LDA makes the topics more interpretable and explainable for researchers. But there is a trade-off – LDA requires significantly more preprocessing than BERTopic. This means researchers have to prepare the text (documents) before analysing it, a step known as 'rendering of corpora' (Hannigan *et al.*, 2019). During the preprocessing, the contextual knowledge of researchers can play a significant part such as removing unnecessary words and restricting words that appear in most documents in the dictionary based on which latent themes will be unravelled. This contrasts with BERTopic which relies on pre-trained language models and does not offer researchers the same level of control.

Therefore, as we aim to categorise software programmers' emotions about GitHub Copilot, the LDA model is appropriate for our research objective.



### *The LDA process and outcome*

Initially, we filtered our dataset based on two criteria: to include only English tweets and to remove the tweets that did not have any interactions (e.g., replies, likes, retweets) as they were more likely to have been shared by bots. In addition, we removed all the retweets of the same tweet, as the focus was on the content of the tweets. This led to a total of 16,130 relevant tweets written by 10,301 software programmers being included in our analyses. In order to conduct a topic modelling analysis on our data, we started by creating a word corpus. This involved standardising the forms of words used in the tweets. To achieve this standardisation, we applied lemmatisation and stemming techniques, transforming words into their root forms (Kobayashi *et al.*, 2018). Using the LDA algorithm, we then uncovered various topics present in our tweet collection. As mentioned above, the LDA algorithm unravels latent topics in the text (i.e., a set of documents) in the form of word vectors. Each vector lists the terms and their corresponding weights in representing a topic (Grimmer and Stewart, 2013).

The LDA algorithm also needs the number of topics required to be identified in the text as an input (Hannigan *et al.*, 2019). Although any number can be chosen, we used the coherence and perplexity measures to determine the optimal number of tweets in the collection of our tweets (Azzopardi, Girolami and Van Risjbergen, 2003). In order to sort, filter and prepare the tweets for the analysis, we used the Pandas library in Python. For the process of stemming and lemmatising the bag words in our tweets, the nltk library was employed. Finally, we conducted the main analysis in Python, using the Gensim library. Figure 2 illustrates that the optimal number of topics for our tweet collection was eight:

-----  
Insert Figure 2 about here.  
-----

We then interpreted the outcome of topic modelling by going through the tweets in each of the eight categories and comparing them with the keywords. This was an iterative process and was conducted by the first and second authors. This led to the emergence of the topics which are shown in Table 1. There were 645 tweets that the algorithm did not find relevant to any of the eight topics, accounting for 4% of the corpus.

-----  
Insert Table 1 about here.  
-----

### 3.2.2 *Second stage*

Following our abductive strategy, we compared the emerging topics with Beaudry and Pinsonneault's (2010) framework of four emotions and found that six topics were related to four emotions within their framework. In this process of interpreting the topics through comparison with the literature and Beaudry and Pinsonneault's (2010) framework, we went back and forth between tweets and the topics to elicit emotions related the framework. As a result, we found that four topics showed two emotions revealing the same emotions, which led to six emotions related to eight topics. These were loss, deterrence, challenge, achievement (according to the framework) and two new emotion categories: scepticism and apathy. This iterative process of comparing the topics with the literature also led us to discover the factors that were driving those emotions. Therefore, we continued this interpretation and further qualitative analysis of tweets in each of the six emotions. This formed the second stage of our analysis to identify which stimulus cues or factors led to those emotions. Table 2 depicts the coding process of this stage.

-----  
Insert Table 2 about here.  
-----

Utilising in vivo terms derived from our dataset, we formulated first-order codes such as “intellectual property”, “collaborating in code writing”, “improvement in productivity”, “creativity unleash”, and “skill development”. In the initial step of this analysis, we cross-referenced our emerging codes and categories with insights from the existing literature. Moreover, although we did not directly incorporate literature themes and concepts into our analysis, in accordance with the recommendations of Gioia et al. (2012), we referenced the literature (e.g., Stein *et al.*, 2015) to contextualize our findings, discerning their relevance, and identifying potential new themes in the process. The second step of analysis concentrated on the comparison and categorisation of the initial codes, leading to the development of second-order codes. Some instances of these codes include: “AI training”, “human AI pair”, “programming proficiency”, and “learning and education”.

During the third step, we extended our comparison of emerging codes to explore what stimulated software programmers’ emotions by developing connections between the second-order codes. Throughout this process, four principal themes surfaced as key elements driving software programmers’ emotions. This was through further grouping of second-order codes into the more abstract themes of “AI development”, “AI functionality”, “Identity work”, “AI engagement” as the four drivers of emotions. To elucidate the connection between drivers and the emotions detected during the first stage of our analysis, we developed Figure 3. This visual representation illustrates the flow from our initial qualitative interpretation of tweet analysis (based on topic modelling) within each emotion category to the subsequent identification of drivers fuelling those emotions.

-----  
Insert Figure 3 about here.  
-----

### 3.2.3 *Third stage*

In the third stage of our analysis, we conducted a more in-depth analysis of the temporal changes of emotions concerning the use of the AI programmer. To achieve this objective, we identified all tweets posted by each tweet author. As mentioned above, our dataset contained 16,310 tweets posted by 10,301 distinct authors. 2,210 of them had more than one tweet about GitHub Copilot in our dataset and among them, 1,473 authors had at least two tweets which were one month or more apart. For our temporal analysis, we focused on these 1,473 authors. These authors tweeting about GitHub Copilot over time (i.e., having at least two tweets apart for one month or more) can indicate that they engaged with the tool more often than others and such engagements might have impacted their feelings and emotions about this AI-enabled programming pair. As a result, it provided us with the opportunity to investigate changes of emotions over time.

For this temporal analysis, we had 1,473 who tweeted 6,282 times about GitHub Copilot over the timeframe we mentioned above. Nonetheless, instead of analysing all these 6,282 tweets and comparing the tweets of each author from the subset of 1,473 authors or any random sampling, we decided to adopt a purposive sampling approach. This approach empowered us to take into account the unique context (Miles, Huberman and Saldana, 2014) of GitHub Copilot which was particularly designed for pair programmers. In other words, as this AI-powered tool was designed for coders and developers, as one of our sampling criteria for the purpose of the temporal analysis, the authors of the tweets needed to be experienced programmers. As discussed above, the other sampling criterion was to have tweeted over time about the tool. The reason for considering experienced developers in this stage is due to their (more likely) continuous interaction with software development projects over time, therefore, it would be more likely that their use of GitHub Copilot would be more sustained and intensive. In this respect, research has shown that experienced software programmers

have the capability to predict software change more accurately (Lindvall and Sandahl, 1998; Tóth et al., 2010).

The process started by going down the list of the 1,473 authors who had tweeted over time about GitHub Copilot. In order to consider the tweets of an author in the temporal analysis, the bio information of the author on their Twitter account was checked. For example, in the following screenshot of an author's Twitter bio (Figure 4), the person referred to working for a well-known technology company previously and his involvement in the development of two famous pieces of software. In addition, in the Twitter bio, there was also a link to their website on which more details of their experience in programming were provided.

-----  
Insert Figure 4 about here.  
-----

On the website, we could also find a link to their GitHub profile. Having visited the GitHub profile, it was possible to check how long they have been coding and contributing to open-source projects, what these projects were and what the level of their activities was to make these contributions. For instance, the following screenshot from the above author's GitHub profile (Figure 5) indicates the number of contributions they had on each day to different open-source projects in 2021 and different projects they contributed to:

-----  
Insert Figure 5 about here.  
-----

The above example shows the robust approach we adopted in order to check whether an author in our dataset was an experienced developer who commented on GitHub Copilot or not. If an author met both of the sampling criteria for the temporal analysis, the person's tweets over time were compared to recognise any change in expressed emotions about this AI-powered tool. We conducted our analysis until we reached theoretical saturation (Miles,

Huberman and Saldana, 2014). Having compared 311 tweets from 97 distinct authors, we were convinced that no more insights would be gained by further analysis, nonetheless we continued our analysis until we reached 106 authors who shared 336 tweets.

#### 4 Findings

In addition to Beaudry and Pinsonneault's (2010) four emotions, our analysis revealed another emotion that we named 'scepticism' (Table 3, Figure 6). There were also a group of software programmers who only shared news about the launch of GitHub Copilot and did not show any specific emotions. Thus, we grouped them under the 'apathy' theme.

-----  
Insert Table 3 about here.  
-----

Furthermore, our analysis revealed four factors that triggered software programmers' emotional reactions to the AI-powered technology. These factors were related to: 1) how GitHub Copilot was developed (AI development); 2) the features and capabilities of GitHub Copilot (AI functionality); 3) how the technology is related to their position and performance (e.g., identity work, Stein *et al.* (2015)), and; 4) how much computer programmers engaged with GitHub Copilot (AI engagement).

-----  
Insert Figure 6 about here.  
-----

##### 4.1 AI development

AI development relates to how the GitHub Copilot algorithm was developed and the data that was used to train the AI according to GitHub (no date), they have utilised the repository of

codes on the platform to train Copilot. Since the launch of the GitHub platform in 2008, individual software programmers developed and publicly shared these codes. Thus, some felt using others' code for training the Copilot algorithm was problematic and a violation of copyright. We identified that this characteristic resulted in two types of emotions: loss and deterrence. 'Loss' resulted from anger, frustration, and a lack of control that a group of software programmers showed on Twitter. This is evident in the following examples:

We can all agree on GitHub Copilot is not a dumb copy / paste machine with an 'AI-powered' sticker on it. These technologies are inevitable and will likely pressure the entire world to bend copyright law until it is considered acceptable. As of right now, I just don't think it is! (T14077)

Who called it GitHub 'Copilot' and not GitHub 'Copolyot'? (T15333)

Guess there's more reason to remove my repos off GitHub. Abuses like these are terrible. Even though I license my code under BSD, the code is still required to retain my copyright notice, and Copilot violates that license by failing to include it. (T14215)

'Deterrence' originated from the fear and shock some programmers felt as open-source resources were used to build the GitHub Copilot algorithm. This group believed that it was unethical for Microsoft uses open-sourced code to develop its AI programmer tool and then sell it to individuals and businesses.

Last Copilot hot take. MS/GitHub [Microsoft/GitHub] have used their dominant market position as the world's largest source of public code to build a closed proprietary service, without their customers' explicit consent. (T60442)

How long until GitHub a) changes the license of Copilot to not allow public sharing, or b) just cancels that whole stupid experiment? Also, can I be compensated for my contributions to their AI model? I mean, they are using all our hard work for profit. (T16024)

#### 4.2 *AI functionality*

Our analysis showed that the functions of Copilot and its capabilities also drove emotional responses. Software programmers were interested to see how this new AI-powered

technology might assist in the process of coding. The analyses revealed that software programmers' emotional responses to AI functionality were linked to 'challenge' and 'achievement' emotions. Both groups saw the AI-powered technology as an opportunity and were excited about its launch. However, one group perceived a lack of control over its functionality (i.e., achievement). In contrast, the other group who tested it demonstrated control over its unexpected consequences (i.e., challenge), as shown in the following tweets:

Seriously... Seriously, GitHub Copilot is the most mind-blowing tool I've ever tried. It seriously helps me out a ton, and it's amazing how non-generic the suggestions are. It's so good. (T9260)

Okay, I know Github Copilot is in technical preview and the developer is in control, etc. But good GOD it is unnerving seeing it accurately guesses what you're going to do next. Five out of six times, and on the sixth, it was something I hadn't thought of but was needed. (T309)

I think this year I've sent enough GitHub Copilot work into our production! I use it every day; sometimes it annoyed me, but most of the time it saves time! Thanks @GitHubCopilot. (T5247)

Besides the emotion of challenge, another group of software programmers expressed their happiness and satisfaction, representing the emotion of achievement. Such emotions were associated with the use of Copilot within their practice and their positive experience of how they could save time by using Copilot, as the following tweets show:

I just played a couple of hours with GitHub Copilot, and I really like it. It can fail like 50% times in its guesses, but what's good is that the other 50% times, it guesses well, and allows me to write code much faster than before, so it's definitely worth it. (T478)

After giving the function name, I just keep pressing 'Tab', thinking GitHub Copilot will write something and it actually does exactly what I want. (T954)

### 4.3 Identity work

Our analysis revealed some software programmers reflected on their knowledge and expertise to express their opinions about Copilot and whether this AI-powered programmer is



sufficiently capable to automate the whole coding process and potentially replace human programmers. Once individuals' use of technology is related to their power, status, and performance, their emotions are triggered by 'identity work' (Stein *et al.*, 2015). Topic modelling showed that identity work triggered two different emotions. First, by relating it to their core knowledge, expertise, and status, some computer programmers demonstrated enjoyment and satisfaction as they believed the AI programmer assisted them in carrying out routine tasks and enabled them to spend more time on innovative tasks. This behaviour comes from the achievement emotion, as shown in the following tweets:

...the best thing about GitHub Copilot, for me, is that it reduces the amount of 'reinventing the wheel' I need to do in a project. (T5145)

GitHub Copilot has been amazing. I built this crazy VR [Virtual Reality] thing in 7 days with it, and it really makes coding a joy. (T10932)

The second group, who compared their performance and status with this technology, questioned the AI programmer's abilities and capabilities. They showed worry and surprise emotions as they were sceptical about whether this could ultimately automate software programming process. We believe this emotion does not fit into Beaudry and Pinsonneault's (2010) four categories of emotions. Therefore, we added a new category to describe these emotions as 'scepticism'. Instances of this category can be seen in the following tweets:

....You don't have to know music theory to play an instrument. You just have to mimick and feel. I highly doubt that this work with software (or web) dev. (T10173)

What is the minimum amount of code in bytes or lines that can be considered licensable and defensible? I doubt a 10 line function would be most of suggestions are less than that. I only see GitHub Copilot as being an overhyped autocompleter. (T14098)

#### 4.4 AI engagement

In this study, AI engagement refers to how software programmers engaged with Copilot and worked with it to estimate its capabilities and values in coding practice. One group we identified felt excited and hopeful (i.e., challenge emotion). The more they used the AI programmer, the more they identified that it could have more extensive use beyond acting as a pair for software programmers. One important use they mentioned was that Copilot could be used to educate software programming students. They suggested that Copilot can help novice programmers or students understand the logic behind the generated codes. For example, the tweets below show that some software programmers believed the AI programmer could be a game changer as it detected their errors, helped them learn more, and assisted them in teaching programming.

GitHub's Copilot is gonna be a game changer for learning how to code. Instead of spending time lost in syntax or clerical errors, you can focus on understanding the code logic. Can't wait till it goes public so all my students can use it. (T5415)

New feature with GitHub CoPilot, where it explains existing code. One of the things I spend a bit of time doing while teaching programming is helping the student build a strong mental model of the code, i.e., understand what each line of code does. This could be helpful! (T5954)

GitHub's Copilot has gotten considerably better over the past few months. I'm not sure why, but I have a hunch the input it takes is now broader than the just the current file content (which I thought was the case when I first started using it), either way works really well! (T999)

The second group were the programmers who used Copilot and watched its development process closely. However, we found no particular emotions expressed by this group of software programmers (i.e., apathy). They tweeted updates regarding the AI programmer. For example, one shared the release of the Copilot extension for Microsoft Visual Studio:

GitHub Copilot technical preview for Visual Studio is available. #github #visualstudio #GitHubCopilot ... (T3163)

GitHub Updates Copilot, 'AI Pair Programmer' and Codespaces (Online VS Code). Editor support now includes Neovim and JetBrains IDEs, especially focused on the latest versions of IntelliJ IDEA and PyCharm. #githubcopilot #githubcodespaces #githubuniverse". (T8590)

#### 4.5 *Stimulus characteristics, emotions, and their dynamics over time*

Figure 7 shows the association between the stimulus features and emotional responses explained in the previous section. As depicted in this figure, the initial loss and deterrence emotions that originated from AI development transformed into challenge and achievement emotions as programmers shifted their attention to the technology's functionality.

-----  
Insert Figure 7 about here.  
-----

Furthermore, as they engaged more with this new AI programmer and related this with their expertise and experience in the practice of programming, they expressed emotions of challenge, achievement and scepticism. Therefore, the two main stimulus features that guided emotional change were AI engagement and identity work. This is evident from the tweets below shared by one programmer over seven months:

**Tweet 1:** "We've started getting nonsensical PRs to GitHub that are too bad but oddly somewhat valid to have been written by a human. They seem like they're Copilot-authored. Somethines the author opens & then closes them. One had mystery commit messing with GH Actions yml. Shady. Banned." (T9900) **Date: 02/09/2021**

**Tweet 2:** "I'm absolutely loving GitHub Copilot lately. I started off thinking it was pretty gimmicky, but it's surprisingly good most of the time. And even when it's wrong, it's funny. Cool comment suggestion here:" (T1700) **Date: 10/04/2022**

**Tweet 3:** "I continue to be blown away by how damn good GitHub Copilot is. I know I've said this all a few times, and it sometimes gets things hilariously wrong, but it's really, really impressive in general. " (T454) **Date: 23/05/2022**

Based on the three tweets, the author's view on GitHub Copilot has gone through a significant shift, from initial deterrence to growing appreciation and even admiration. In the first tweet

(September 2, 2021), Copilot-generated code is seen as "nonsensical" and "shady" with concern about automation-induced authorship mistakes and potential security risks. The immediate response the author suggested was to ban Copilot-generated contributions. However, tweet 2 (April 10, 2022) shows a positive shift as Copilot is now described as "absolutely loving" and "surprisingly good". Humor is found even in Copilot's occasional errors, suggesting a more accepting attitude. This specific example highlights Copilot's helpfulness in suggesting comments. Tweet 3 (May 23, 2022) continued positive reinforcement and the author praised Copilot as "damn good" and "really, really impressive". The author's journey with Copilot showcases a change from initial negative emotion to growing acceptance and even positive endorsement.

Another group of software programmers associated their initial loss with their experience and knowledge (i.e., identity work) and expressed scepticism. They viewed the AI programmer as a new tool that learns and improves itself over time, however, its performance significantly depends on training data and resources. Therefore, they expected expert maintainers to play an important role by filtering and designating resources that such tools would learn from. Such a view limits the expectations of knowledge workers from AI technologies and algorithms to a certain level due to algorithms' performance depending on the quality of training resources:

I'd rather like to see GitHub Copilot learning from designated resources. Imagine that it looks for testable examples or designated examples directories. Then, it may actually motivate maintainers to produce high-quality examples. (T14100)

However, a combination of identity work and AI engagement as stimulus characteristics led to an emotional shift from loss to achievement. An example is the series of tweets below:

**Tweet 1:** "I could be very wrong but I was thinking the issue was that copilot was sharing other people's GitHub code with the user, occasionally verbatim, regardless of how it was licensed" (T13864) **Date: 08/07/2021**

**Tweet 2:** ... tools can reduce cognitive load and increase productivity. GitHub Copilot is a great example. When it can figure out what I'm doing, it provides a pretty good starting point for a given method or class." (T5087) **Date: 03/01/2022**

This example illustrates the overall shift in author's emotion. The author's initial tweet centered on potential licensing and code ownership issues, reflecting an ethically-minded approach and loss emotion. In the later tweet (after about six months), the focus shifts to the practical benefits of Copilot in terms of productivity and cognitive load reduction. This suggests a growing recognition of Copilot's potential value as a tool (achievement emotion). The following set of tweets similarly show a similar transition:

**Tweet 1:** "Using GitHub copilot is like playing a video game with cheat codes." (T9512) **Date: 17/09/2021**

**Tweet 2:** "GitHub Copilot is getting pretty good at providing test recommendations" (T1077) **Date: 29/04/2022**

The first tweet indicates an initial ambivalence. The "cheat codes" analogy suggests initial ethical concerns about the use of Copilot and its impact on coding practices and skills. However, the second tweet shows a growing appreciation, which indicates a more positive emotion (achievement), focusing on a specific area where Copilot excels (test recommendations). This suggests a growing recognition of its potential benefits.

## 5 Discussion

This research reports how software programmers as a group of knowledge workers felt about an AI-powered technology called GitHub Copilot and how their emotions changed over time. In the era of work digitalisation, the adoption and application of AI have been increasingly widespread across occupations and professions. Thus, it is necessary to understand how professionals and knowledge workers respond to it. We explored how software programmers voluntarily engaged with the AI programmer (i.e., GitHub Copilot) without organisational pressure. Thus, we expand the existing literature by focusing on emotional reactions to AI-powered technologies at the individual level rather than the organisational level. Consistent

with previous studies (e.g., Cao *et al.*, 2021; Weiler *et al.*, 2019), our research shows how knowledge workers express mixed emotions about AI adoption, either positive and negative or positive and neutral.

### 5.1 *Emotional responses and AI use in knowledge workers' practice*

Based on Beaudry and Pinsonneault's (2010) framework of emotion categories, we classified the range of emotions knowledge workers expressed and categorised them into five groups (not considering the apathetic group). We extended on this framework by introducing a new emotion category called scepticism. Scepticism is neither positive nor negative, and it is neutral regarding the perception of control over AI consequences. The analysis revealed that scepticism was reported by software programmers who relied on their identity work (i.e., experience and knowledge) but did not use or test the AI programmer thoroughly. Based on their expertise and how GitHub used open access code to train the AI programmer, they argued that the AI programmer would not be creative and innovative like a human programmer. Moreover, our results suggest such affective or emotional responses were based on the four stimulus characteristics of AI development, AI functionality, level of engagement and identity work.

Those knowledge workers who engaged more with the AI technology and were experienced (more potent in professional identity) were positive towards this new AI tool, showing challenge or achievement emotions. This is contrary to recent research on the use of AI technologies by other professional communities, such as medical professionals. Two studies of physicians and radiologists revealed that implementing new AI-powered technologies at work resulted in both optimism and pessimism among them (Kim *et al.*, 2021; Maier *et al.*, 2020). Moreover, unlike the existing body of research on the role of automation through technologies such as AI on jobs and employment, which can be perceived as a threat and

cause negative emotions (e.g., anxiety) (Goethals and Ziegelmayer, 2022; Strich *et al.*, 2021), our research identified a more positive outlook to automation of knowledge workers' practices.

An important difference between the extant literature and our study is the context of our research and the specific group of knowledge workers we examined. While most prior studies focused on non-IT knowledge workers, we investigated IT professionals who can directly relate to AI-powered technologies and are perhaps emotionally attached to them. Therefore, they have a better sense of how the technology works and what role it may play in their practice, and thus, their professional identity plays a role in how they feel about AI.

## 5.2 *The interplay between AI automating and augmenting and knowledge workers' identity*

Since we studied knowledge workers' emotions toward an AI technology which is related to their professional identity, we have found that through time and with more engagement with AI technology, knowledge workers viewed AI less as an identity threat but rather as identity expansion (Selenko *et al.*, 2022). A group of software programmers in this study indicated that the AI programmer could act as a learning tool and potentially help them and their trainees and mentees upskill. In addition, since more routine parts of programming practice would be automated, that would enable the augmentation of experienced and qualified knowledge workers. This supports the recent discussion by Raisch and Krakowski (2020) who suggested that automation and augmentation are not separate, but interdependent, as AI technologies can conduct a whole task or complement knowledge workers' competencies in fulfilling the task (Rai *et al.*, 2019).

Similarly, our study indicated that a group of knowledge workers did not view AI as replacing human actors but as complementing their skills and expertise. Depending on their level of experience and authority in their field, the dynamic interaction of automating and

augmenting creates an AI-knowledge worker hybrid leading to their identity expansion or enhancement (Selenko *et al.*, 2022).

Our findings also contrast with Strich *et al.*'s (2021) study that suggests the adoption of AI-powered technologies can deskill qualified employees, lowering the required skills needed for their job. We argue that our positive findings could originate from our level of analysis (individual level) and the point that the software programmers in our study were not forced by their employers to use the AI programmer. Therefore, we argue that the voluntary adoption of AI-powered technologies might decrease users' negative feelings despite some general concerns over AI threats, and might enable users to get closer to their ideal work selves.

## **6 Conclusions**

### *6.1 Theoretical contributions*

The contribution of our study is threefold. First, our findings contribute to the literature on emotion and IT use (Beaudry and Pinsonneault, 2010; Stein *et al.*, 2015). We not only confirmed Beaudry and Pinsonneault's (2010) framework of emotions toward technology use, but we also add a new emotion (i.e., scepticism) to their four categories (i.e., achievement, challenge, loss and deterrence). This new emotion category is neutral regarding viewing AI technology as an opportunity or threat, or perception of control over AI's consequences.

In addition, we found four stimulus characteristics (i.e., AI development, AI functionality, AI engagement, identity work) that triggered the five emotions and drove their change through time (see Figure 7). The first shift (from loss and deterrence to achievement and challenge) occurred when software programmers interested in the AI development process were influenced by AI functionality stimulus. The second emotional shift was from scepticism to



achievement, as the software programmers who were initially sceptical changed their emotions after using the AI programmer and perceived it as an opportunity.

Second, we contribute to the growing body of literature on the role of AI technology in the future of work. As applications of AI are expanding from the automation of routine tasks to ones requiring more analytical and intuitive intelligence (Huang and Rust, 2018), there has been more concern regarding the impact of AI on higher-skilled jobs (Weiler *et al.*, 2019). Recent studies examining the role of AI technologies in managerial decision-making (Cao *et al.*, 2021), surgery operations (van Hoek *et al.*, 2019), and radiology (Kim *et al.*, 2021) showed that there is a combination of positive and negative attitudes and feelings towards AI tools. Our findings build on this discussion by showing that despite both optimism and pessimism towards AI, experienced professionals show more positive emotions over time as they perceive AI as enhancing their work and skills.

Third, this study contributes to the role of technology in professional identity. Unlike some studies on the impact of AI on employment (Goethals and Ziegelmayer, 2022; Strich *et al.*, 2021) that suggest negative perceptions of AI resulting in negative emotions (e.g., anxiety, stress), we found negative perceptions could change to positive emotions once knowledge workers understood AI functionality and engaged with it in their practice. In contrast to the previous studies of AI in work practices (Strich *et al.*, 2021), we studied the emotional response of a particular group of knowledge workers who understand technology, are more emotionally connected to it, and can directly relate to it. Therefore, despite the initial sense of identity threat, the interdependence between automating and augmenting could lead to 'identity expansion' (Selenko *et al.*, 2022) as threats became less evident over time.

## 6.2 *Practical implications*

This study has several practical contributions. First, our findings show how the relationship between adopting AI-powered technologies and knowledge workers is complex. Specifically, the way in which specialised users emotionally engage with AI systems is subject to change in time. Understanding the evolving nature of knowledge workers' emotions towards AI is crucial for organisational leaders. Managers can utilise these insights to proactively plan the introduction of new AI systems into the workplace. By anticipating and addressing initial negative emotional attitudes, organisations can implement more effective change management strategies.

Second, our study suggests that AI adoption should be designed with an awareness of the emotional evolution among knowledge workers. Our study shows that experimenting with new AI tools is important in incorporating the technology into users' day-to-day practices. It allows knowledge workers to see how the AI tool works and understand its benefits and drawbacks. Thus, organisations can use an experimental approach to tailor their AI adoption process that minimises potential resistance and fosters a positive reception. This may involve phased introductions, extensive training programs, or targeted communication strategies to highlight the benefits of AI tools over time.

Third, our research shows the positive evolution in emotions towards the AI coding tool over time. Thus, organisations can leverage this insight to promote continuous learning among knowledge workers. Encouraging employees to explore and experiment with AI tools not only enhances their skill sets but also contributes to a culture of innovation within the organisation. Finally, some tweets suggested the significant contribution of AI, such as GitHub Copilot, to the education and training of software programmers. Thus, we argue that organisations and educational institutions can integrate AI technologies into training programs, allowing novice programmers to benefit from AI-generated coders and insights.

This dual-use of AI for both productivity and educational purposes aligns with the evolving landscape of tech-driven workplaces.

### *6.3 Research limitations and future research*

This study has a few limitations which could guide future studies. First, we acknowledge that the difference between the extant literature and our study is the context of our research and the specific group of knowledge workers we examined. We focused on software programmers, while the existing studies focus mainly on non-IT knowledge workers (e.g., managers and radiologists). Since software programmers' job is directly related to technology, it's tempting to assume software programmers' tech-savviness directly translates to heightened emotional attachment to AI, but their perspective demands closer examination. While professional identity and knowledge undoubtedly influence their feelings, specific AI characteristics add a nuanced layer to this context.

Unlike other technological advancements, AI can potentially upend core processes within entire professions (regardless of type of knowledge work), not just automate individual tasks (von Krogh, 2018). In addition, due to AI systems' inherent learning and adaptability, they are evolving and their behaviour can dynamically shift across situations and tasks.

Furthermore, AI's hidden and often inscrutable algorithms (Dourish, 2016) pose a challenge even for any knowledge worker (including even IT workers or tech-savvy programmers).

This lack of transparency creates an opaque assistance system, potentially fostering similar emotional responses in both programmers and non-IT knowledge workers. While programmers may grasp how AI was trained, the "why" behind its outputs remains unknown, similar to their non-IT counterparts. Therefore, we suggest future research is needed to explore how software programmers' professional identity and knowledge play a part in their emotions and also to focus on comparative studies between different types of knowledge

workers and their emotions in relation to how AI systems are impacting their practices. This could shed light on why some of our findings were inconsistent with existing reports.

Second, although our study revealed that the continuous use of AI technologies might elicit more positive emotions, one area of future research is to explore this phenomenon longitudinally and see how the incorporation of AI may change over time as the functionality and capacity of AI-powered technologies such as Copilot's learning algorithm will develop over time. In addition, algorithms behind AI technologies are dynamic but opaque; thus, these technologies might have unintended or unknown long-term consequences (Faraj *et al.*, 2018). Third, further research is required regarding ethical concerns about creating a proprietary AI programmer based on open-source code and marketing the final product.

Fourth, future research might also explore other professions' affective responses towards Github Copilot and how they use and incorporate AI tools in their practice. This would be an interesting and important line of enquiry as research has shown that users' responses might differ when they are not closely related to AI technology. In addition, what we studied was voluntary use of AI as the technology was only available in preview version (for free) and the specific group of knowledge workers we studied were experimenting the new Github copilot tool. Future studies can compare the voluntary and non-voluntary use of AI and their impact on users' perceptions and emotions, especially within the context of organisations and teams. This is important as these AI tools are being developed and adopted rapidly and organisations might enforce their use to enhance productivity.

Finally, as recently suggested by Panteli *et al.* (2022), more research is required to focus on gender differences in emotions and AI, as male and female knowledge workers might express different emotions towards such tools which in turn might influence their interactions and collaborations in software development projects.

## References

- Asatiani, A. *et al.* (2020) 'Knowledge workers' reactions to a planned introduction of robotic process automation—Empirical evidence from an accounting firm', in Hirschheim, R., Heinzl, A., and Dibbern, J. (eds) *Information Systems Outsourcing: The Era of Digital Transformation*. Springer, pp. 413–452.
- Azzopardi, L., Girolami, M. and Van Risjbergen, K. (2003) 'Investigating the relationship between language model perplexity and IR precision-recall measures', in *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pp. 369–370.
- Baumer, E. *et al.* (2017) 'Comparing grounded theory and topic modeling: Extreme divergence or unlikely convergence?', *Journal of the Association for Information Science and Technology*, 68(6), pp. 1397–410.
- Beaudry, A. and Pinsonneault, A. (2010) 'The Other Side of Acceptance: Studying the Direct and Indirect Effects of Emotions on Information Technology Use', *MIS Quarterly*, 34(4), pp. 689–710.
- Benbya, H., Davenport, T. H. and Pachidi, S. (2020) 'Artificial intelligence in organizations: current state and future opportunities', *MIS Quarterly Executive*, 19(4).
- Berente, N., Seidel, S. and Safadi, H. (2019) 'Data-Driven Computationally Intensive Theory Development Data-Driven Computationally Intensive Theory Development', *Information Systems Research*, 30(1), pp. 50–64.
- Bhattacharjee, A. (2001) 'Understanding information systems continuance: An expectation-confirmation model', *MIS quarterly*. JSTOR, 25(3), pp. 351–370.
- Blei, D. M., Ng, A. Y. and Jordan, M. I. (2003) 'Latent dirichlet allocation', *Journal of machine Learning research*, 3(Jan), pp. 993–1022.
- Brynjolfsson, E. and McAfee, A. (2011) *Race against the machine: How the digital*

*revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy.* Lexington, MA: Digital Frontier Press.

Brynjolfsson, E. and Mitchell, T. (2017) 'What can machine learning do? Workforce implications', *Science*, 358(6370), pp. 1530–1534.

Cao, G. *et al.* (2021) 'Understanding managers' attitudes and behavioral intentions towards using artificial intelligence for organizational decision-making', *Technovation*. Elsevier Ltd, 106, p. 102312.

Davenport, T. H. (2018) *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.

Davenport, T. H. and Kirby, J. (2016) 'Just how smart are smart machines?', *MIT Sloan Management Review*, 57(3).

Davenport, T. H. and Ronanki, R. (2018) 'Artificial intelligence for the real world', *Harvard business review*, 96(1), pp. 108–116.

Davis, F. D. (1989) 'Perceived usefulness, perceived ease of use, and user acceptance of information technology', *MIS quarterly*, 13(3), pp. 319–340.

Davis, F. D., Bagozzi, R. P. and Warshaw, P. R. (1992) 'Extrinsic and intrinsic motivation to use computers in the workplace', *Journal of applied social psychology*. Wiley Online Library, 22(14), pp. 1111–1132.

Deonna, J. A. (2006) 'Emotion, perception and perspective', *dialectica*. Wiley Online Library, 60(1), pp. 29–46.

Dindar, M. and Dulkadir Yaman, N. (2018) '#IUseTwitterBecause: content analytic study of a trending topic in Twitter', *Information Technology and People*, 31(1), pp. 256–277. doi: 10.1108/ITP-02-2017-0029.

Dourish, P. (2016) 'Algorithms and their others: Algorithmic culture in context', *Big Data & Society*, 3(2), p. 205395171666512.

- Dourish, P. and Gómez Cruz, E. (2018) 'Datafication and data fiction: Narrating data and narrating with data', *Big Data and Society*, 5(2), pp. 1–10.
- Faraj, S., Pachidi, S. and Sayegh, K. (2018) 'Working and organizing in the age of the learning algorithm', *Information and Organization*. Elsevier, 28(1), pp. 62–70.
- Ferràs-Hernández, X. (2018) 'The future of management in a world of electronic brains', *Journal of Management Inquiry*. SAGE Publications Sage CA: Los Angeles, CA, 27(2), pp. 260–263.
- Frey, C. B. and Osborne, M. A. (2017) 'The future of employment: How susceptible are jobs to computerisation?', *Technological Forecasting and Social Change*. Elsevier B.V., 114, pp. 254–280. doi: 10.1016/j.techfore.2016.08.019.
- Gerli, P. *et al.* (2022) 'The hidden power of emotions: How psychological factors influence skill development in smart technology adoption', *Technological Forecasting and Social Change*. Elsevier Inc., 180(April), p. 121721. Available at: <https://doi.org/10.1016/j.techfore.2022.121721>.
- Ghahramani, Z. (2015) 'Probabilistic machine learning and artificial intelligence', *Nature*. Nature Publishing Group, 521(7553), pp. 452–459.
- GitHub.com (2022) *GitHub number of public repositories*.
- Gkinko, L. and Elbanna, A. (2022) 'Hope, tolerance and empathy: employees' emotions when using an AI-enabled chatbot in a digitalised workplace', *Information Technology and People*.
- Goethals, F. and Ziegelmayer, J. L. (2022) 'Anxiety buffers and the threat of extreme automation: a terror management theory perspective', *Information Technology and People*, 35(1), pp. 96–118. doi: 10.1108/ITP-06-2019-0304.
- Grimmer, J. and Stewart, B. M. (2013) 'Text as data: The promise and pitfalls of automatic content analysis methods for political texts', *Political analysis*. Cambridge University Press,

21(3), pp. 267–297.

Grootendorst, M. (2022) ‘BERTopic: Neural topic modeling with a class-based TF-IDF procedure’. Available at: <https://arxiv.org/abs/2203.05794v1>.

Guzman, E., Alkadhi, R. and Seyff, N. (2017) ‘An exploratory study of Twitter messages about software applications’, *Requirements Engineering*. Springer London, 22(3), pp. 387–412.

Hannigan, T. R. *et al.* (2019) ‘Topic Modeling in Management Research’, *Academy of Management Annals*, 13(2), pp. 586–632.

van Hoek, J. *et al.* (2019) ‘A survey on the future of radiology among radiologists, medical students and surgeons: Students and surgeons tend to be more skeptical about artificial intelligence and radiologists may fear that other disciplines take over’, *European Journal of Radiology*. Elsevier, 121(November), p. 108742.

Hu, Q. *et al.* (2021) ‘Can AI artifacts influence human cognition? The effects of artificial autonomy in intelligent personal assistants’, *International Journal of Information Management*. Elsevier Ltd, 56(March 2020), p. 102250.

Huang, M. H. and Rust, R. T. (2018) ‘Artificial Intelligence in Service’, *Journal of Service Research*, 21(2), pp. 155–172.

Jaiswal, A., Arun, C. J. and Varma, A. (2022) ‘Rebooting employees: upskilling for artificial intelligence in multinational corporations’, *International Journal of Human Resource Management*. Routledge, 33(6), pp. 1179–1208.

Jeon, E., Yoon, N. and Sohn, S. Y. (2023) ‘Exploring new digital therapeutics technologies for psychiatric disorders using BERTopic and PatentSBERTa’, *Technological Forecasting and Social Change*, 186, p. 122130.

Johansen, J., Pedersen, T. and Johansen, C. (2021) ‘Studying human-to-computer bias transference’, *AI & SOCIETY*. Springer, pp. 1–25.



Kim, B. *et al.* (2021) 'How does the radiology community discuss the benefits and limitations of artificial intelligence for their work? A systematic discourse analysis', *European Journal of Radiology*. Elsevier B.V., 136, p. 109566. doi: 10.1016/j.ejrad.2021.109566.

Kim, H.-W. *et al.* (2004) 'Understanding the balanced effects of belief and feeling on information systems continuance', in Agarwal, R., Kirsch, L., and DeGross, J. I. (eds) *Proceedings of the 25th International Conference on Information Systems*. Washington, DC, pp. 297–310.

Kim, H.-W. and Kankanhalli, A. (2009) 'Investigating User Resistance to Information Systems Implementation: A Status Quo Bias Perspective', *MIS Quarterly: Management Information Systems*, 33(3), pp. 567–582.

Kobayashi, V. B. *et al.* (2018) 'Text classification for organizational researchers: A tutorial', *Organizational research methods*. SAGE Publications Sage CA: Los Angeles, CA, 21(3), pp. 766–799.

von Krogh, G. (2018) 'Artificial Intelligence in Organizations: New Opportunities for Phenomenon-Based Theorizing', *Academy of Management Discoveries*, 4(4), pp. 404–409. doi: 10.5465/amd.2018.0084.

Larivière, B. *et al.* (2017) "'Service Encounter 2.0": An investigation into the roles of technology, employees and customers', *Journal of business research*. Elsevier, 79, pp. 238–246.

Lindvall, M. and Sandahl, K. (1998) 'How well do experienced software developers predict software change?', *Journal of Systems and Software*. Elsevier, 43(1), pp. 19–27.

Maier, S. B., Jussupow, E. and Heinzl, A. (2020) 'Good, bad, or both? Measurement of physician's ambivalent attitudes towards AI', *27th European Conference on Information Systems*.

Miles, M., Huberman, A. M. and Saldana, J. (2014) 'Qualitative data analysis: A methods

sourcebook'. Thousand Oaks, California: Sage Publications, Inc.

Niedenthal, P. M. and Wood, A. (2019) 'Does emotion influence visual perception? Depends on how you look at it', *Cognition and Emotion*, 33(1), pp. 77–84.

Panteli, N., Giæver, F. and Engesmo, J. (2022) 'Guest editorial: Emotions in the digitalised workplace', *Information Technology and People*, 35(6), pp. 1677–1692.

Peeters, M. M. M. *et al.* (2021) 'Hybrid collective intelligence in a human–AI society', *AI and Society*. Springer London, 36(1), pp. 217–238.

Rai, A., Constantinides, P. and Sarker, S. (2019) 'Editor's Comments- Next-Generation Digital Platforms: Toward Human–AI Hybrids', *Mis Quarterly*, 43(1), pp. iii–ix.

Raisch, S. and Krakowski, S. (2020) 'Artificial Intelligence and Management: The Automation-Augmentation Paradox', *Academy of Management Review*, pp. 1–48. doi: 10.5465/2018.0072.

Ramaswamy, S. (2017) 'How companies are already using AI', *Harvard Business Review*, 14(April), p. 2017.

Reichertz, J. (2007) 'Abduction: The logic of discovery of grounded theory', in Charmaz, K. and Bryant, A. (eds) *The SAGE handbook of grounded theory*. SAGE Publications Ltd, pp. 214–228.

Reinhardt, W. *et al.* (2011) 'Knowledge worker roles and actions—results of two empirical studies', *Knowledge and process management*, 18(3), pp. 150–174.

Sætre, A. S. and Van de Ven, A. (2021) 'Generating theory by abduction', *Academy of Management Review*. Academy of Management Briarcliff Manor, NY, 46(4), pp. 684–701.

Sánchez-Franco, M. J. and Rey-Moreno, M. (2022) 'Do travelers' reviews depend on the destination? An analysis in coastal and urban peer-to-peer lodgings', *Psychology and Marketing*. John Wiley and Sons Inc, 39(2), pp. 441–459. doi: 10.1002/mar.21608.

Schneider, P. and Sting, F. J. (2020) 'Employees' Perspectives on Digitalization-Induced

Change: Exploring Frames of Industry 4.0', *Academy of Management Discoveries*, 6(3), pp. 406–435.

Selenko, E. *et al.* (2022) 'Artificial Intelligence and the Future of Work: A Functional-Identity Perspective', *Current Directions in Psychological Science*, 31(3), pp. 272–279.

Shane-Simpson, C. *et al.* (2018) 'Why do college students prefer Facebook, Twitter, or Instagram? Site affordances, tensions between privacy and self-expression, and implications for social capital', *Computers in Human Behavior*. Elsevier Ltd, 86, pp. 276–288.

Stein, M.-K. *et al.* (2015) 'Coping with information technology: Mixed emotions, vacillation, and nonconforming use patterns', *MIS Quarterly*, 39(2), pp. 367–392.

Strich, F., Mayer, A. S. and Fiedler, M. (2021) 'What Do I Do in a World of Artificial Intelligence? Investigating the Impact of Substitutive Decision-Making AI Systems on Employees' Professional Role Identity', *Journal of the Association for Information Systems*, 22(2), pp. 304–324.

Thompson, M. (2012) 'People, practice, and technology: Restoring Giddens' broader philosophy to the study of information systems', *Information and Organization*. Elsevier, 22(3), pp. 188–207.

Tóth, G. *et al.* (2010) 'Comparison of different impact analysis methods and programmer's opinion: an empirical study', in *Proceedings of the 8th International Conference on the Principles and Practice of Programming in Java*, pp. 109–118.

Vaast, E. *et al.* (2017) 'Social media affordances for connective action: An examination of microblogging use during the Gulf of Mexico oil spill', *MIS Quarterly: Management Information Systems*, 41(4), pp. 1179–1206. doi: 10.25300/misq/2017/41.4.08.

Vaast, E. and Walsham, G. (2013) 'Grounded theorizing for electronically mediated social contexts', *European Journal of Information Systems*. Nature Publishing Group, 22(1), pp. 9–25.

Venkatesh, V. (2000) 'Determinants of Perceived Ease of Use: Integrating Control, Intrinsic Motivation, and Emotion into the Technology Acceptance Model', *Information Systems Research*, 11(4), pp. 342–365.

Venkatesh, V. *et al.* (2003) 'User acceptance of information technology: Toward a unified view', *MIS quarterly*. JSTOR, 27(3), pp. 425–478.

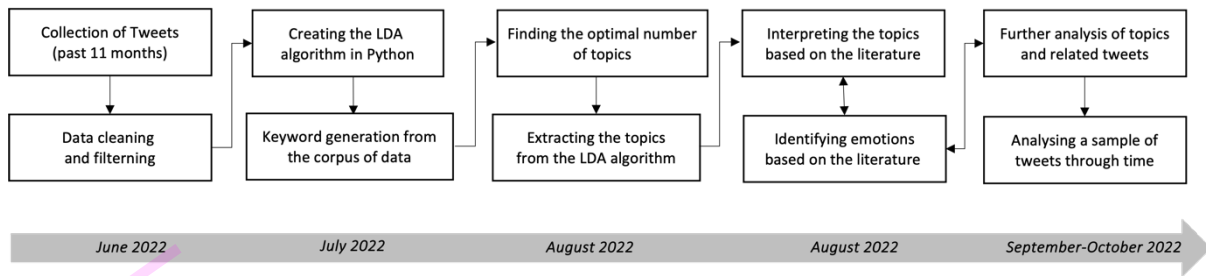
Weiler, S., Matt, C. and Hess, T. (2019) 'Understanding user uncertainty during the implementation of self-service business intelligence: A thematic analysis', *Proceedings of the Annual Hawaii International Conference on System Sciences*, 2019-Janua, pp. 5878–5887.

Williams, G. and Mahmoud, A. (2017) 'Analyzing, classifying, and interpreting emotions in software users' tweets', in *Proceedings - 2017 IEEE/ACM 2nd International Workshop on Emotion Awareness in Software Engineering, SEmotion 2017*. Institute of Electrical and Electronics Engineers Inc., pp. 2–7.

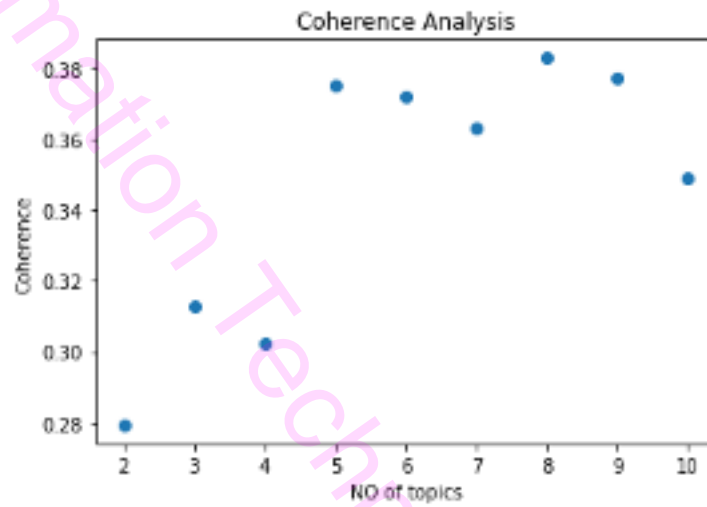
Yu, X., Xu, S. and Ashton, M. (2022) 'Antecedents and outcomes of artificial intelligence adoption and application in the workplace: the socio-technical system theory perspective', *Information Technology and People*. doi: 10.1108/ITP-04-2021-0254.

Zhang, P. (2013) 'The Affective Response Model: A Theoretical Framework of Affective Concepts and Their Relationships in the ICT Context', *MIS Quarterly*, 37(1), pp. 247–274.

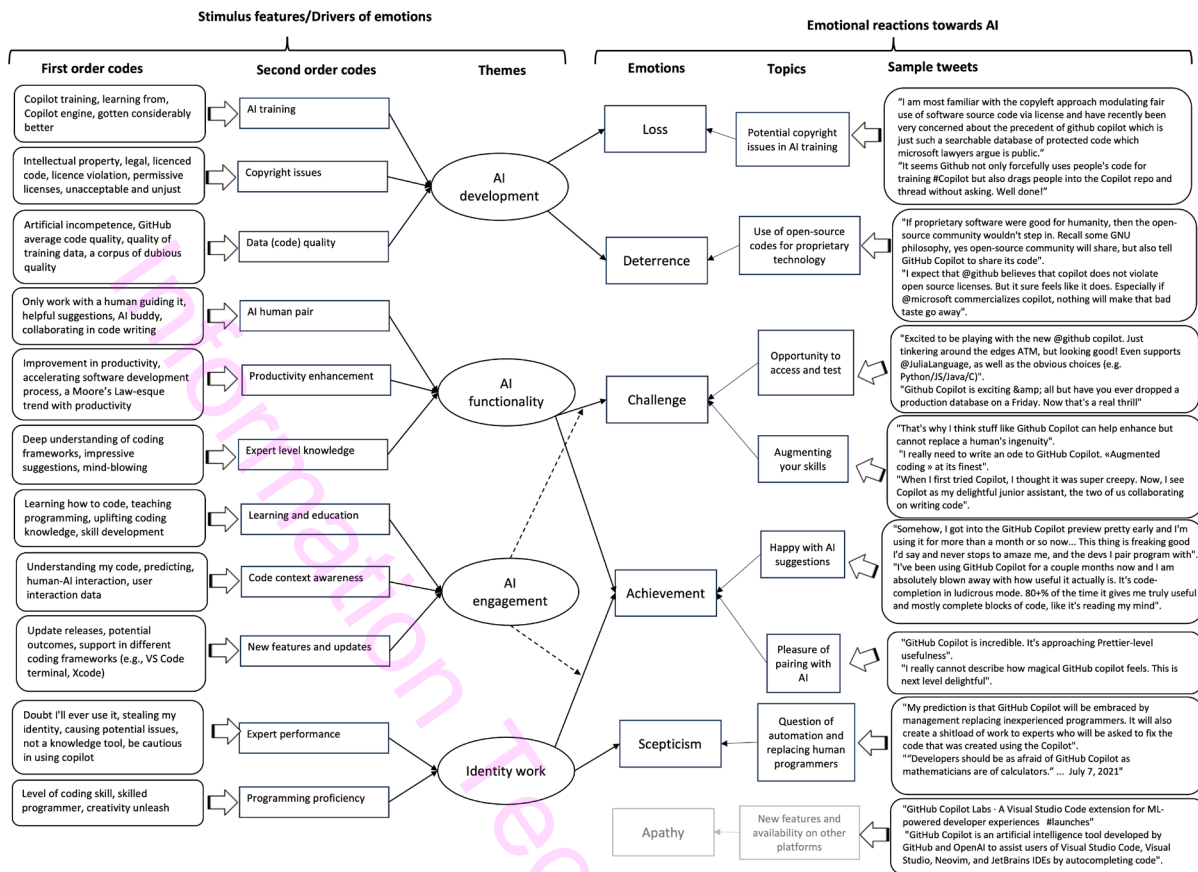
Van Zoonen, W., Verhoeven, J. W. M. and Vliegenthart, R. (2016) 'How employees use Twitter to talk about work: A typology of work-related tweets', *Computers in Human Behavior*, 55, pp. 329–339. doi: 10.1016/j.chb.2015.09.021.



**Figure 1.** The research process



**Figure 2.** Coherence analysis



**Figure 3.** The connection between topic modelling (first stage) and qualitative analysis (second stage)



**Figure 4.** Screenshot of a Twitter bio from our sample

6,408 contributions in 2021

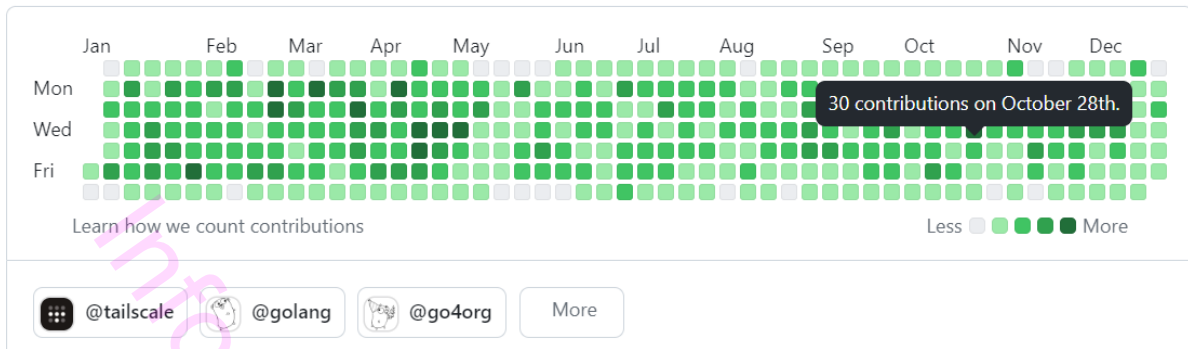


Figure 5. Screenshot of a developer contributions on GitHub

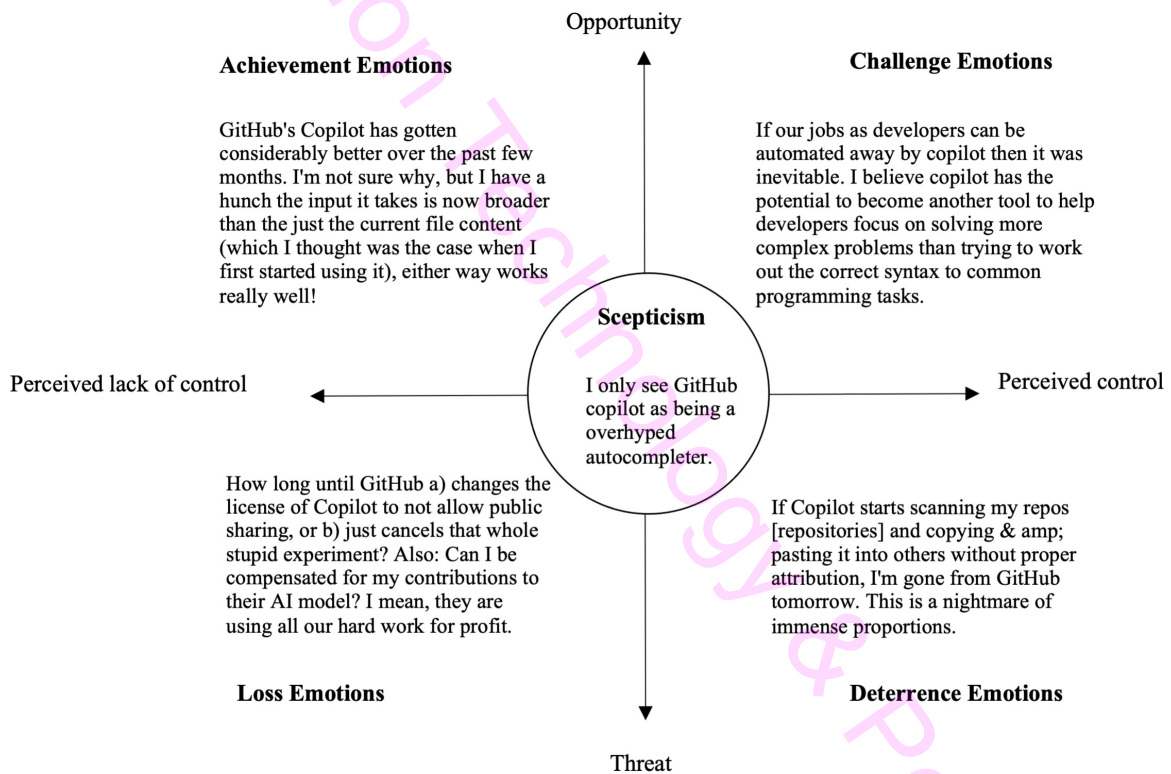
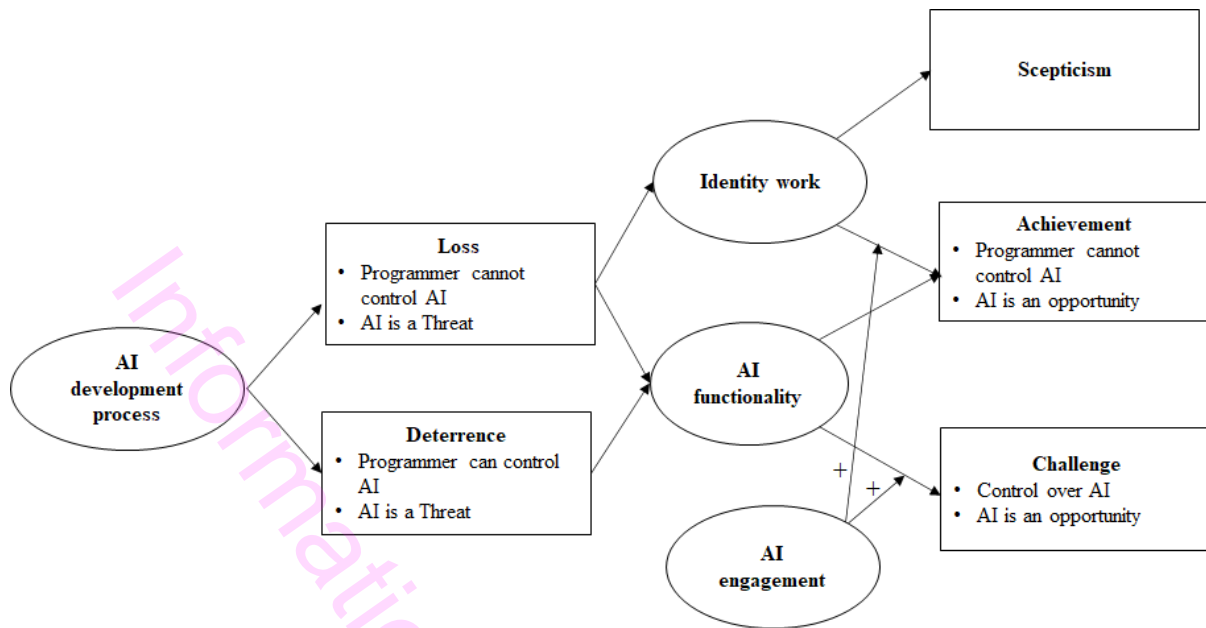


Figure 6. Categories of emotional responses towards Copilot, adapted from Beaudry and Pinsonneault (2010)



**Figure 7.** Stimulus characteristics and emotional change

Topic #	Keywords (with different weights in the word vector of the topic)	Labelling topics	% of Tweets corpus
1	"code" + "licens" + "train" + "generat" + "copi" + "copyright" + "work" + "data" + "past" + "say"	Potential copyright issues in AI training	10.33
2	"cod" + access" + "open" + "daysofcod" + "sourc" + "program" + "python" + "microsoft" + "javascript" + "wait"	Use of open-source codes for proprietary technology	11.70
3	"work" + "know" + "preview" + "today" + "thank" + "mind" + "product" + "tri" + "technic" + "access" + "invite"	Opportunity to access and test	14.60
4	"programm" + "develop" + "program" + "code" + "pair" + "openai" + "softwar" + "autom" + "codex" + "engin"	Question of automation and replacing human programmers	8.80
5	"code" + "write" + "suggest" + "time" + "function" + "complet" + "comment" + "good" + "help" + "test" + "like"	Happy with AI suggestions	19.23
6	"code" + "like" + "write" + "think" + "help" + "look" + "develop" + "feel" + "tool" + "learn"	Augmenting your skills	15.34
7	"develop" + "code" + "vscode" + "visual" + "studio" + "githubcopilot" + "extens" + "instal" + "avail"),	New features and availability on other platforms	8.47
8	"exampl" + "post" + "amaz" + "stuff" + "review" + "project" + "need" + "hear" + "interest" + "solve" + "pair"	Pleasure of pairing with AI	7.53

**Table 1.** The outcome of topic modelling and labelling topics



First order codes	Second order codes	Themes
Copilot training, learning from, Copilot engine, gotten considerably better	AI training	AI development
Intellectual property, legal, licenced code, licence violation, permissive licenses, copying and pasting, being held responsible, unacceptable and unjust, legal and ethical considerations	Copyright issues	
Artificial incompetence, data privacy, quality of open-source code, GitHub average code quality, secure/insecure code, quality of training data, a corpus of dubious quality	Data (code) quality	
Only work with a human guiding it, autocompletion, helpful suggestions, documentation, AI buddy, collaborating in code writing, auto fill code, junior assistant	AI human pair	AI functionality
Improvement in productivity, accelerating software development process, a Moore's Law-esque trend with productivity	Productivity enhancement	
Deep understanding of coding frameworks, impressive suggestions, mind-blowing, customising codes, interactive recommendations	Expert level knowledge	
Doubt I'll ever use it, stealing my identity, causing potential issues, not a knowledge tool, be cautious in using copilot	Expert performance	Identity work
Level of coding skill, skilled programmer, creativity unleash	Programming proficiency	
Learning how to code, teaching programming, uplifting coding knowledge, skill development	Learning and education	AI engagement
Understanding my code, predicting, human-AI interaction, user interaction data	Code context awareness	
Update releases, potential outcomes, support in different coding frameworks (e.g., VS Code terminal, Xcode), new applications	New features and updates	

**Table 2.** The coding process in second stage of analysis

Definition	Emotion	% of Tweets corpus
This group was angry and shocked because they believed that GitHub violated copyright law by training the AI programmer on publicly available codes and individuals' works on its platform. Loss emotion demonstrates a combination of anger, lack of control over the outcomes and perceiving the AI programmer as a threat	Loss	13%

<p>This group perceived the AI programmer training as unethical because Microsoft used open-source code to train GitHub Copilot. Deterrence emotion is a combination of feeling control over the outcomes but perceiving the AI programmer as a threat</p>	Deterrence	11%
<p>This group expressed their excitement and happiness over the technical aspects and abilities of the new technology, and they liked to use it. Challenge emotion combines excitement, feeling control over the consequences of using the technology and perceiving it as an opportunity. They also felt the AI programmer could help and play a part in educating programmers as a learning tool (e.g., explaining the logic behind the code by writing comments)</p>	Challenge	30%
<p>This group questioned the ability and capability of the AI programmer as to how a tool trained by an OpenAI Codex Engine could replace human software programmers</p>	Scepticism	9%
<p>This group tested, wrote comments, and completed functions within code using the AI programmer's suggestions. They were happy with the result primarily due to saving time in coding. However, they felt a lack of control over the AI programmer. Achievement emotion encompasses happiness, a lack of control and perceiving the new technology as an opportunity</p>	Achievement	28%
<p>This group only shared news about the development of GitHub Copilot: they did not demonstrate any emotions about it</p>	Apathy	9%

**Table 3.** Topic modelling results

Technology & People