

# Why and When LLM-Based Assistants Can Go Wrong: Investigating the Effectiveness of Prompt-Based Interactions for Software Help-Seeking

Anjali Khurana  
anjali\_khurana@sfu.ca  
Simon Fraser University  
BC, Canada

Hari Subramonyam  
harihars@stanford.edu  
Stanford University  
USA

Parmit K Chilana  
pchilana@cs.sfu.ca  
Simon Fraser University  
BC, Canada

## ABSTRACT

Large Language Model (LLM) assistants, such as ChatGPT, have emerged as potential alternatives to search methods for helping users navigate complex, feature-rich software. LLMs use vast training data from domain-specific texts, software manuals, and code repositories to mimic human-like interactions, offering tailored assistance, including step-by-step instructions. In this work, we investigated LLM-generated software guidance through a within-subject experiment with 16 participants and follow-up interviews. We compared a baseline LLM assistant with an LLM optimized for particular software contexts, *SoftAIBot*, which also offered guidelines for constructing appropriate prompts. We assessed task completion, perceived accuracy, relevance, and trust. Surprisingly, although *SoftAIBot* outperformed the baseline LLM, our results revealed no significant difference in LLM usage and user perceptions with or without prompt guidelines and the integration of domain context. Most users struggled to understand how the prompt's text related to the LLM's responses and often followed the LLM's suggestions verbatim, even if they were incorrect. This resulted in difficulties when using the LLM's advice for software tasks, leading to low task completion rates. Our detailed analysis also revealed that users remained unaware of inaccuracies in the LLM's responses, indicating a gap between their lack of software expertise and their ability to evaluate the LLM's assistance. With the growing push for designing domain-specific LLM assistants, we emphasize the importance of incorporating explainable, context-aware cues into LLMs to help users understand prompt-based interactions, identify biases, and maximize the utility of LLM assistants.

## CCS CONCEPTS

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

## KEYWORDS

feature-rich software; large language models; prompt-based interactions; help-seeking

## ACM Reference Format:

Anjali Khurana, Hari Subramonyam, and Parmit K Chilana. 2024. Why and When LLM-Based Assistants Can Go Wrong: Investigating the Effectiveness of Prompt-Based Interactions for Software Help-Seeking. In *29th International Conference on Intelligent User Interfaces (IUI '24)*, March 18–21, 2024, Greenville, SC, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3640543.3645200>

## 1 INTRODUCTION

Learning to use feature-rich software applications for tasks such as advanced word processing, data analysis, image manipulation, and video editing can be challenging for end-users. Users currently turn to various software help resources to learn and seek help for such software tasks. For example, they usually begin by querying online search engines using keywords to locate specific resources, such as video and text-based tutorials, forums posts, and blogs and articles [3, 22, 35, 36]. However, online software-help seeking is a complex endeavour, demanding precise queries to pinpoint the most pertinent information that can be directly applied within the application [18, 22].

The recent emergence of Generative AI and pre-trained Large Language Model (LLM)-based assistants like *ChatGPT* [2, 8] offers a novel approach to support end-users' software help-seeking by leveraging these assistant's advanced natural language understanding [52]. For example, LLMs provide the potential for step-by-step instructions and explanations tailored to specific task needs, saving users time and effort searching through online resources. In the past two years, LLMs have demonstrated promising capabilities in assisting users in various domains, including tasks related to programming and software development [50], language generation [2], and question answering [2]. However, the effectiveness of LLM-based assistance and how end-users employ LLMs to seek help for feature-rich applications remain important open questions. Furthermore, there have been calls [29] for a better understanding of the dynamics of user interaction across diverse AI usage scenarios, aiming to avoid false assumptions.

In this work, we investigate how software users make use of LLMs with *prompt-based interactions* for seeking help for feature-rich applications. In particular, we investigate the effectiveness of recently emerging prompt guidelines [42, 44] offered by OpenAI, Microsoft, and others, which, when prepended to the user prompts, can potentially enhance LLM output to provide desired assistance [52]. Furthermore, we also consider the impact of integrating additional domain context, such as software documentation, into LLMs to improve the precision of the LLM output. For example, such techniques are currently being explored in in-application

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*IUI '24, March 18–21, 2024, Greenville, SC, USA*

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-0508-3/24/03  
<https://doi.org/10.1145/3640543.3645200>

LLM assistants, such as Copilot in Microsoft 365 applications [43], Firefly in Adobe applications [1], etc.). For this investigation, we developed *SoftAIBot*, our implementation of a state-of-the-art LLM assistant that integrates prompt guidelines and domain context, and compared it with a Baseline ChatGPT Plus LLM assistant. By analyzing users' interactions, behaviors, and challenges when using these state-of-the-art LLMs for software help-seeking, we aim to help end users harness the full potential of LLM-based assistants for feature-rich software [31]. The research questions guiding this exploration were:

- **RQ1:** How do end users make use of prompt-based interactions when finding software help ?
- **RQ2:** To what extent can SoftAIBot generate accurate and relevant software help for end-users of feature-rich applications?
- **RQ3:** How do end users' mental models of LLMs influence their use of LLM-generated software help?

In this paper, we report on the results from a controlled study and follow-up interviews that illustrate how end-users make use of prompt-based interactions when using LLM-based assistance in concert with tasks that involve visual interactions (e.g., Microsoft PowerPoint) and advanced data analysis/visualizations (e.g., Microsoft Excel). We ran a within-subject experiment with 16 participants from varied backgrounds without expertise in Machine Learning (ML) or Natural Processing Language (NLP). We hypothesized that our implementation of SoftAIBot with ChatGPT Plus (GPT-4) as the underlying LLM, state-of-the-art prompt guidelines, and integrated domain context will generate accurate and relevant assistance for users' prompts and help users finish their software tasks. However, our findings showed that even though SoftAIBot performed better than BaseLine in generating more accurate and relevant LLM output, users could not recognize these differences and struggled in mapping the LLM instructions to the software application, leading to poor task completion. The impact of the prompt text on the quality of the LLM output was not clear to users, and they rather followed LLMs' suggestions blindly, even when the output was inaccurate. Lacking an accurate mental model of LLMs, users tended to over-trust the LLM assistance without much contemplation.

The main contributions of this research are in providing empirical insights that: (1) demonstrate how software users employ new-generation LLM assistants to seek software help, both with and without prompt guidelines as well as integrated software context; (2) illustrate the challenges that software users experience with prompt-based interaction (e.g., crafting prompts, comprehending how prompts bias LLM output, mapping LLM-suggested steps to software, overtrusting output correctness); (3) identify gaps in users' mental models as they try to seek and apply assistance from LLMs to software tasks, which affect both their use of LLMs and perception of software features, regardless of implicit enhancements in the underlying model or explicit prompt guidelines (*SoftAIBot*). We discuss the implications of these findings for designing and implementing LLM help tailored to feature-rich applications, the need to incorporate transparent and responsible LLM assistants, and ways to bridge the disparity between mental models and LLM interfaces, ultimately helping end-users form accurate mental models of LLMs.

## 2 RELATED WORK

This research drew upon insights from prior work on how users learn and seek help for feature-rich software, the emergence of LLMs for task-based assistance, and the use of prompt-based interactions.

### 2.1 Software Help-seeking evolution

HCI research has a rich history of investigating the challenges that users experience when learning and seeking help for complex feature-rich applications [11, 18, 22, 35]. Help-seeking resources and approaches have evolved over the years: from formal documentation and manuals [36, 39] to the use of videos [23, 25], interactive tutorials [35], Google Search, Q&A or FAQ sites, blogs, dedicated forums [22] and even contextual help systems embedded within applications [7, 11, 17, 19, 24].

However, studies have shown that although users have increased access to help information, they often get lost in search results and forum posts and still face difficulty in recognizing effective and relevant [22, 34] resources for completing their software tasks. Users also face numerous issues with articulating search queries using precise keywords [16, 18, 22], often termed as *vocabulary problem* [16]. Another related issue that users face is using the located help in coordination with the application and going back and forth between the two to accomplish their software tasks [22]. Past studies have shown that users tend to find step-by-step guidance within the context of feature-rich applications useful and trustworthy [11, 21, 22]. There have been constant innovations in devising help resources in the form of in-context help and video tutorials [17, 25, 35]. However, recent developments in Generative AI have opened a new outlet of help-seeking for users through LLM-based assistance, which is much more forgiving in letting users describe their queries using natural language [41, 52]. These assistants tend to provide more specific and in-context assistance [52] to users, unlike traditional resources that may be scattered and require precise queries for retrieval. While there is a rich history of work supporting software learnability and easing help-seeking processes [11, 14, 17, 25], it is unclear whether users' learning approaches and strategies apply to LLM-based assistants. We extend this prior work by examining how novices employ LLM assistants for software help-seeking and the types of challenges they experience.

### 2.2 LLM use for Task-Based Assistance

The emergence of powerful LLM assistants, such as ChatGPT, has significantly impacted task-based assistance for a range of domains, including programming, text generation, and text summarization [47]. Recent studies have attempted to understand the use of LLM assistants for programming and software development-related help-seeking [5, 51]. For example, Xu, Vasilescu, & Neubig (2022) [51] investigated the use of LLMs for programming-related tasks and found that users struggle to generate assistance especially for complex queries as users struggled in formulating specific code-related input queries. Another study [48] highlighted that users mostly rely on trial and error while debugging their code using LLM assistance and often do not feel confident about applying the output.

Recent interest in LLMs has inspired initiatives [1, 43] to utilize their capabilities for assisting users with software tasks as well [1, 43]. For example, some experimental work is being explored by integrating LLMs directly into feature-rich applications, such as Copilot in Microsoft 365 [43] and Firefly in Adobe [1]. This process has shown that developers can face new challenges in ensuring accurate and effective use of this new avenue of conversational UX experience [29, 52]. As many of these interfaces are still at a nascent stage, it is unclear how this current practice (i.e., integrating the context of these feature-rich applications) can help novice end-users in seeking accurate and relevant assistance from LLMs. Furthermore, to harness the full potential of these LLM assistants for seeking help for their software tasks, we need more insights on where and how users struggle with these LLMs [31]. Our study contributes new knowledge on how non-AI expert end-users employ LLMs' generated software guidance assistance in accomplishing tasks for feature-rich applications by comparing our own implemented, SoftAIBot, an LLM optimized for particular domain context (e.g., software documentation) with the Baseline ChatGPT.

### 2.3 Prompt-based interactions

To leverage the potential of LLMs, a lot of the focus in HCI and AI research is turning to *prompt-based interactions* as users generally have to provide input or queries in the form of prompts that are then processed or responded to by a conversational AI system [52]. Recent studies on the usability of prompt-based interactions [5, 41, 51] reveal that prompts have a significant impact on pre-trained language models' ability to produce desired outputs, even though the prompts themselves are simple textual instructions for the task at hand [52]. For example, Advait et. al's (2022) [41] study on the use of LLM-assisted tools for programming tasks revealed that the crucial concern is crafting effective prompts that elevate the probability of an LLM model to generate efficient code. Thus, the big challenge for end-users, especially novices and non-AI experts, is to define the appropriate prompts and learn prompting strategies to get the desired assistance from these LLMs.

Unlike traditional help-seeking mediums that rely on keyword matching, prompt-based interactions within LLMs offer human-like language capabilities [29], which is unique, but can also be unreliable. This unreliability comes from the biases (e.g., hallucinating and non-deterministic output) inherent within prompt-based interactions of LLMs. Considering LLMs are a tremendous leap from traditional help-seeking mediums that most users are familiar with, there have been calls to investigate users' mental models as they interact with LLMs [29]. This becomes necessary when seeking assistance for feature-rich software tasks, where there is an interplay between the mental model of LLM vs the software application [22, 29]. Recent studies have focused on understanding users' prompting strategies and proposing a catalogue of prompting guidelines [42, 44, 50] for allowing users to craft better prompts and seek desired LLM assistance. However, the use of these prompt guidelines in practice and their effectiveness remains unclear. Our study complements the existing research by observing users with prompt-based interactions and assessing the efficacy of prompt-based guidelines and integration of domain context in enhancing LLM assistance for software tasks.

## 3 METHOD: CONTROLLED EXPERIMENT AND FOLLOW-UP INTERVIEWS

We conducted a two-part user study with 16 users that consisted of a controlled experiment and follow-up interviews. Our main goal in this study was to investigate the effectiveness of two recent advancements: 1) prepending prompt guidelines to user prompts to enhance the accuracy of LLM output, as advocated by OpenAI, Microsoft and others to enhance Generative AI tools [42, 44]; and, 2) directly integrating domain context (e.g., software documentation) into LLM assistants (e.g., as demonstrated in Copilot in Microsoft 365 applications [43] and Firefly in Adobe applications [1], etc.) to enhance the relevance and accuracy of LLM output for software-related tasks. For our investigation, we implemented both of these advancements in a new GPT-4-based assistant, which we call *SoftAIBot*, that offers in-context prompt guidelines (See Figure 1) and enhances the LLM output by making it specific to the feature-rich application (See Figure 1.c). For comparison, we also implemented *Baseline ChatGPT* based on the pre-trained state-of-the-art LLM assistants, ChatGPT. (SoftAIBot is explained in more detail below.) The prompt suggestions were not included in the Baseline ChatGPT because it was a control condition for the experiment.

Based on our research questions, we derived the following hypotheses for users seeking software help:

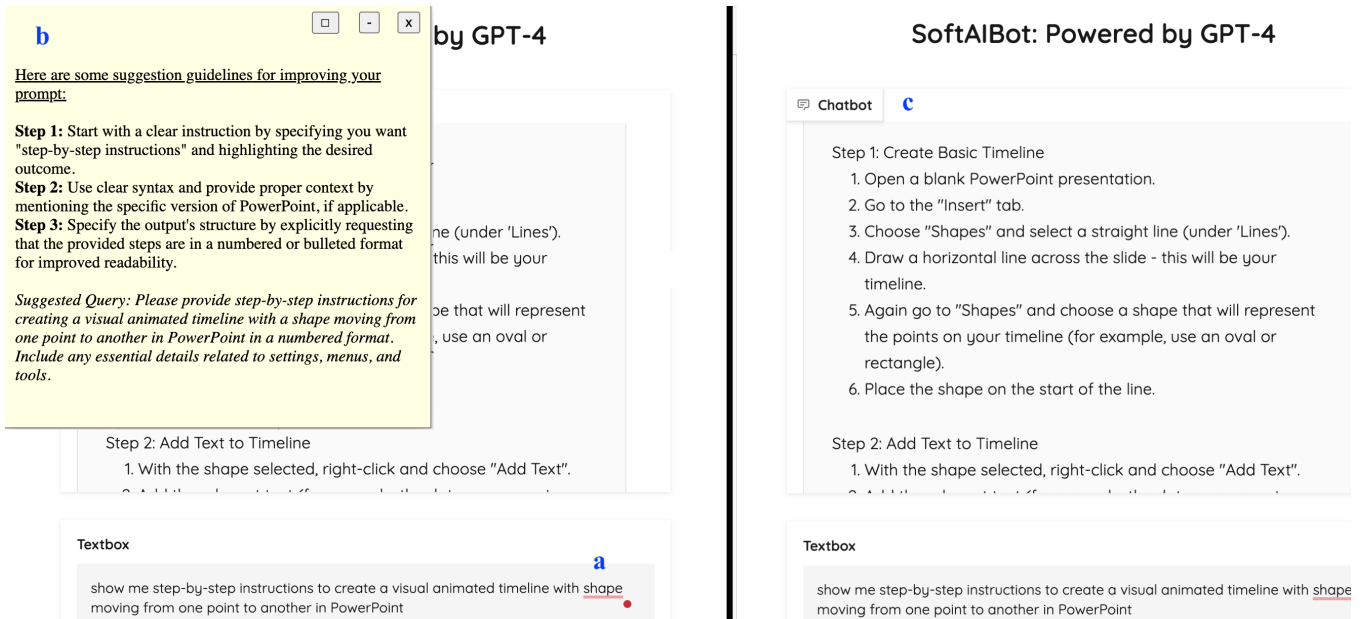
- **H1:** Users will perceive SoftAIBot as being more accurate than Baseline ChatGPT.
- **H2:** Users will perceive SoftAIBot as being more relevant than Baseline ChatGPT.
- **H3:** Users will trust SoftAIBot more than Baseline ChatGPT.
- **H4:** Users will find the output provided by SoftAIBot to be easier to apply in the software application than Baseline ChatGPT.

### 3.1 Participants

We recruited 16 participants (9F|7M) for our study, focusing on non-AI expert users who had little to no prior experience or knowledge of ML or NLP. Our participants came from different backgrounds (CS, Engineering, Business, Arts) and professions (administrative services, business analytics, information designers, client services, students, and researchers). Participants were familiar with LLM-based assistant, ChatGPT (10/16), and a range of traditional chatbots (12/16) such as Siri, and Google Assistant (12/14). They had used ChatGPT before for text generation, text summarization, and programming tasks, but none of them used it for software tasks used in the study before. About half of the participants (7/16) had frequently used PowerPoint and Excel applications and the remaining were occasional users. Our participants covered a range of age groups: 18-24 (25%), 25-34 (62%), 35-44 (13%) and had different levels of education (2 Diploma, 4 Bachelor's, 5 Master's, 5 PhD). We recruited participants mainly from our university's mailing lists and found additional participants through snowball sampling.

### 3.2 Design and Implementation of SoftAIBot and Baseline ChatGPT

In this section, we describe the design and system implementation of our two interventions.



**Figure 1: SoftAIBot integrates domain context via documentation and offers prompt guidelines to construct better prompts: (a) allows users to type in the prompt text and submit it; (b) generates prompt suggestions in-response to a user’s text (in this case, also shows a sample transformed query that users can directly use); (c) formats the response as step-by-step instructions optimized for particular software contexts, in this case PowerPoint. To see the contrast in LLM response, please see Baseline in Figure 2.**

### 3.2.1 SoftAIBot Intervention (GPT-4 with Prompt Guidelines and Software Documentation)

SoftAIBot LLM intervention suggests in-context prompt guidelines for constructing prompts while interacting with GPT-4 as well as generates guidance for particular software contexts.

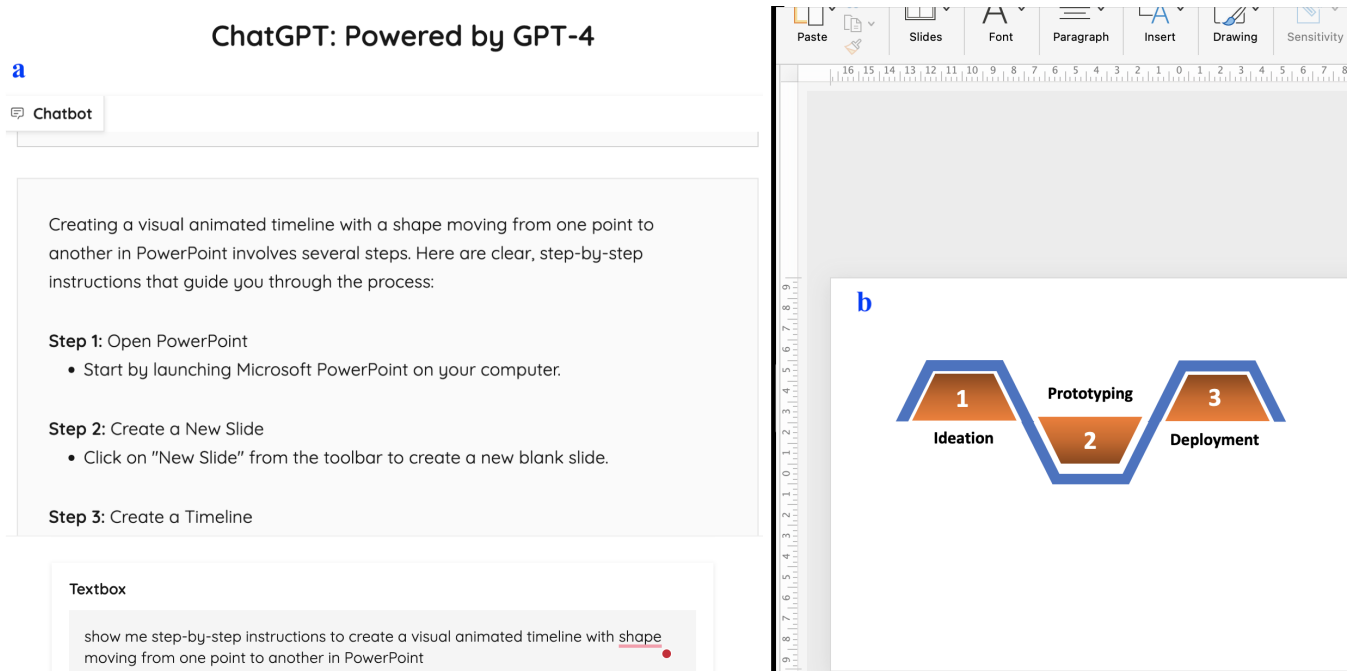
**Automatic Prompt Guidelines:** The SoftAIBot UI interface (See Figure 1) lays out the various user interface components and receives the user’s prompt. Next, this user’s prompt gets transmitted to our custom API developed in Python connected to GPT-4 for generating prompt suggestions in-context to user’s prompt and a sample transformed query based on suggested guidelines that users can directly use. We used the prompt guidelines provided by OpenAI, Microsoft and others [42, 44]. The generated prompt suggestions are displayed as an overlay (as shown in Figure 1.b) on the top left of the SoftAIBot interface via a *Chrome* Extension which triggers with the click of “send” button component. These prompt suggestions appear after every prompt-based interaction initiated by the user. To allow users freedom and control in accessing the prompt suggestions, we included the *minimize*, *maximize*, and *close* options for either using the suggestion or simply closing it anytime during the interaction.

**Implicit Domain Context Integration:** To optimize SoftAIBot for a particular software context (as shown in Figure 1.c), we leveraged GPT-4 augmented with corresponding software documentation. When a user submits the prompt, the relevant textual information or pertinent excerpts are extracted from software documentation by using Facebook AI Similarity Search (FAISS) index

and vector search. Our custom API send this software context information along with user’s original prompt, together as a payload, to OpenAI’s GPT-4 8k to generate better tailored responses with reduced hallucinations. This process of using LLM with extracted relevant text is known as Retrieval Augmented Generation (RAG) Vector search [26]. We used open source model BGE-smal-en [32] for producing chunks of the software documentation and transforming it to dense contextual vectors with 384 dimensions. For the search approach used in RAG, we tried different approaches such as text chunking, and embedding models. The chunk size of 512 and BGE-smal-en approach provided us better context and fast retrieval. Next, we searched this vector against the FAISS index to retrieve the most similar text vectors. Based on the indices of these similar vectors, we fetch the corresponding original text data from software documentation. We fed this extracted text from software documentation, along with the user’s query, as a prompt to GPT-4 (See Appendix B for details).

**3.2.2 Baseline ChatGPT intervention (ChatGPT plus).** Our Baseline ChatGPT mimics the existing ChatGPT (See Figure 2.a) plus based on GPT-4, where users can type in their query and LLM provide assistance to users for variety of tasks through out-of-the box multi-turn conversations using OpenAI GPT-4 8k API [2, 8]. To maintain the user’s history and conversations when interacting with the OpenAI GPT-4 API, we designed Python algorithm using chat completion objects [40] (See Appendix B for details).

To make the LLM UIs consistent for the experiment, we simulated both the interventions using Gradio framework [46], an open source framework, for developing intuitive web-based UI interfaces,



**Figure 2: Overview of sample user task with PowerPoint application: (a) Users were asked to look up and use instructions from LLM intervention. In this case, Baseline ChatGPT mimics the existing ChatGPT plus based on GPT-4, where users can type in their prompt in the textbox and LLM provide assistance to users for variety of tasks; (b) Use LLM assistance to develop shown project timeline in Microsoft PowerPoint that is visual and animated.**

making it easier to gather feedback, showcase results, and enable end-users to interact with LLMs during the user study.

### 3.3 Choice of Application and Tasks

To choose tasks and applications, we explored different productivity-related feature-rich applications (e.g., PowerPoint, Excel, Photoshop, Teams, etc.) popular among everyday novice users. After our initial exploration, we selected Microsoft PowerPoint and Excel to cover a range of tasks involving visual interactions, interactions involving application of statistical functions or formula and other visualization related tasks.

To assess the users' help-seeking approaches and observe any potential challenges when using LLM assistants for software help, we selected tasks that would require multiple steps for completion, and would necessitate multi-stage help and prompts (e.g., help within different steps needed for task completion). For example, one of the Excel tasks asked participants to use instructions from the LLM for analyzing and visualizing the predictive analytics of the sales values based on income using linear regression. Similarly, one of the PowerPoint tasks (See Figure 2) asked users use instructions from the LLM to develop a project timeline in Microsoft PowerPoint that is visual and animated.

### 3.4 Study Design and Procedure

We used a within-subject design to minimize the effect of inter-participant variability. To eliminate order effects, we used a Latin Square counterbalancing [38] with 2 LLM conditions (total possible

order= 4) to balance the order in which tasks were presented. During the experiment, each participant completed two tasks with each LLM interventions (4 tasks in total). The participants performed all tasks using one feature-rich application, subsequently transitioning to a second feature-rich application, but the order of the tasks and associated LLM were randomized.

Each study session began by introducing the participant to the LLM assistants, and provided some general tips to interact with the application (e.g., using the prompts). We conducted the study remotely through Zoom and participants were each given a \$15 Amazon gift card in appreciation of their time. Participants were provided instructions to install our LLM interventions via a Chrome extension. Next, participants completed a demographic questionnaire on their background and prior experiences with LLM assistants, chatbots and software applications.

We presented each LLM intervention along with software application to the participant in a random order, one by one. We designed the complicated software tasks so that participants would be able to spend at least 8 minutes for completing each software task regardless of their familiarity and experience with the software application using given LLM intervention. After completing each of the 4 tasks, participants were asked to fill the post-task questionnaire hosted on the *SurveyMonkey* to assess the overall experience seeking assistance from LLM interventions for software tasks along with their perceptions of accuracy, relevancy, ease of use, and trust of the assistance provided by LLM intervention. We encouraged participants to think aloud [33] throughout the session

and reminded participants that the study was seeking to understand how they seek assistance from LLM interventions for their software tasks rather than their performance or ability to master software applications or use the LLM.

Lastly, we conducted follow-up interviews to further probe any difficulties that impacted the use of prompt-based interaction and LLM assistance for software tasks, and any potential gaps in mental models about how LLMs work. Each session lasted approximately one hour and sessions were video and audio-recorded for transcription, and the participants were asked to share their screen through Zoom (only during the usability test).

### 3.5 Data Collection and Analysis

Throughout the study session, we recorded the participant’s screen and audio recorded their interview responses. We captured screen recordings to evaluate two key aspects: how users sought help from LLM assistants (e.g., formulated their prompts) and how they used the LLM output to complete the prescribed software tasks in Excel and PowerPoint.

We used a combination of statistical tests and inductive analysis approach to make sense of the data captured from the user study. We ran Pearson’s Chi-square test for independence with nominal variable “LLM Interventions” (having two levels: SoftAIBot, and Baseline ChatGPT) and ordinal variable (having three collapsed levels: Agree, Neutral and Disagree) to quantitatively determine the significance of the results.

We used an *expert-rating approach* where the experimenter (in consultation with all authors) analyzed the tasks performed by users in the software application and compared all of our metrics against the ground truth for all tasks. The ground truth in our context refers to the correct or optimal sequence of steps for task completion, the ideal application of software features, and the most relevant responses from the LLM to user queries. Our metrics included:

- (1) *Task Completion*: To measure the task completion, we evaluated how many of the software task steps (e.g., sequence of features/ functions) users completed using the LLM help.
- (2) *Task Accuracy*: To measure the participants’ success in applying LLM assistance to complete the study tasks accurately, we evaluated how accurately users identified the approach or functionality (e.g., macro, animation, motion paths, particular statistical function) from the LLM output and then applied the help using the corresponding software menu options and features in the software application.
- (3) *Accuracy and Relevance of LLM Assistance*: We compared queries and LLM response logs against pre-established ground truth to assess: a) how accurately the LLM provided instructions needed for successful software task completion; and, b) how relevant the LLM response was to the users’ input prompt.

Finally, to complement our experimental findings, we corroborated the data with participants’ think-aloud verbalizations and probed into the reasons behind users’ decisions and identify any potential gaps in their mental models about how LLMs work. We used an inductive analysis approach [12] and affinity diagrams [12] along with discussions amongst the research team to categorize

the interview findings and identify key recurring themes. In particular, our coding approach for the inductive analysis considered reasons influencing users’ perception of how prompts impact output, recognition of accurate vs. hallucinated output, and difficulties in applying the LLM assistance to the software task.

## 4 RESULTS

### 4.1 Task Completion, Accuracy and Relevance of LLM Assistance

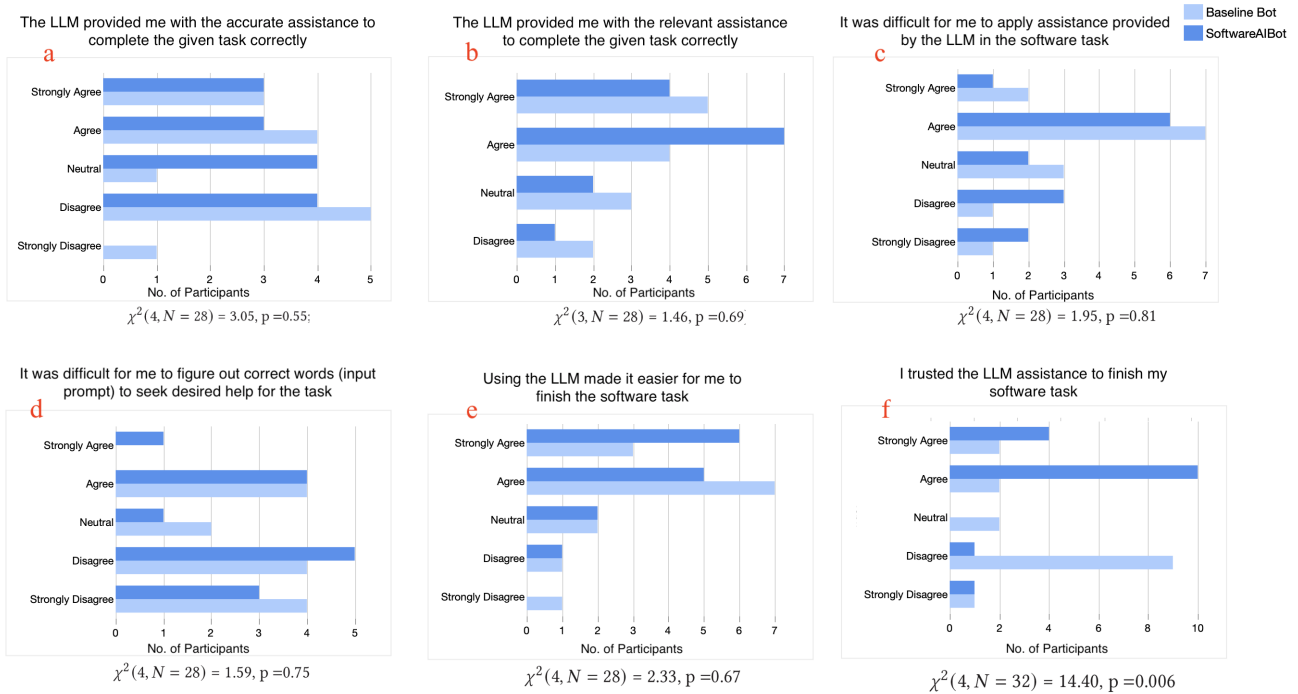
**Expert-rated Accuracy and Relevancy of LLM assistance:** As predicted, we found that the assistance provided by SoftAIBot was more accurate than Baseline ChatGPT for both PowerPoint (SoftAIBot: Mean=64.4%; Baseline ChatGPT: Mean=37.5%) and Excel (SoftAIBot: Mean=65.7%; Baseline ChatGPT: Mean=45%) tasks. These differences between accuracy scores and LLM interventions were significant for PowerPoint ( $t(21.3) = 4.0, p=0.0006$ , two-tailed) and Excel tasks ( $t(23.4) = 3.7, p=0.0011$ , two-tailed). Similarly, we found that the assistance provided by SoftAIBot had higher average relevancy score than Baseline ChatGPT for both PowerPoint (SoftAIBot: Mean=74.7%; Baseline ChatGPT: Mean=44.4%) and Excel (SoftAIBot: Mean=78.2%; Baseline ChatGPT: Mean=55.4%) tasks, with significant differences (PowerPoint:  $t(24.5) = 5.4, p<0.0001$ , two-tailed; Excel:  $t(24.8) = 4.8, p<0.0001$ , two-tailed).

Since the Baseline lacked relevant domain context, it typically failed to offer relevant and accurate steps available within the software (e.g., macros, animations, motion paths, etc.). Instead, it often provided references to features and functionality that did not exist. This phenomenon of the LLM providing information that is somewhat relevant to the user’s query but not accurate to the users’ intent for performing the task has been termed as an *hallucination* [4, 6, 29] in the literature (See Figure 4). On the other hand, while SoftAIBot provided step-by-step instructions on how to implement the required functionality in the software, it also demonstrated instances of hallucination. Furthermore, it did not always provide specific and relevant instructions (for example, where to locate menu functions within the UI).

In subsequent sections, we shed light on users’ performance and qualitative perceptions of both LLMs, highlighting various inconsistencies and misconceptions among users that impacted their use of prompt-based interactions.

**Task Completion and Task Accuracy:** None of the participants were able to completely finish either of the tasks in PowerPoint or Excel, even though each participant made full use of both LLM assistants (*Baseline ChatGPT* and *SoftAIBot*). On average, participants completed 35% of the PowerPoint tasks (maximum= 50% and minimum= 0%) with Baseline ChatGPT and 45% of the tasks (maximum= 60% and minimum= 0%) with SoftAIBot with no significant difference across the two LLM interventions ( $t(28.1) = 1.1, p= 0.28$ , two-tailed). Similarly, participants completed 40% of the Excel tasks (maximum= 55% and minimum= 0%) with Baseline ChatGPT and 55% of the tasks (maximum= 75% and minimum= 0%) with SoftAIBot with no significant difference across the two interventions ( $t(28.7) = 1.9, p=0.07$ , two-tailed).

Amongst the portion of the task completed by each participant, the accuracy scores were low. For the PowerPoint tasks, the average



**Figure 3: Overview of participants' responses to post-task questionnaire. Pearson Chi-Squared test showed no significant difference for each metric across both LLM interventions for completing both Excel and PowerPoint tasks. Despite having low completion rate and low task accuracy, the majority of users perceived that they obtained accurate (a) and relevant (b) assistance from both LLM interventions. Still, the majority of participants (c) found it difficult to apply LLM assistance and instructions to the software application to complete their task; (d) Participants overall did not find it difficult to craft prompts; a few participants did indicate that they struggled to find the correct words for Powerpoint tasks that were more visual and interactive. Although expert ratings showed that users did not finish the task accurately, most users believed that it was easier for them to finish the task using both forms of LLM assistance; (f) The majority of users trusted both LLMs - this was surprising to see because expert ratings showed that both LLMs frequently provided inaccurate assistance.**

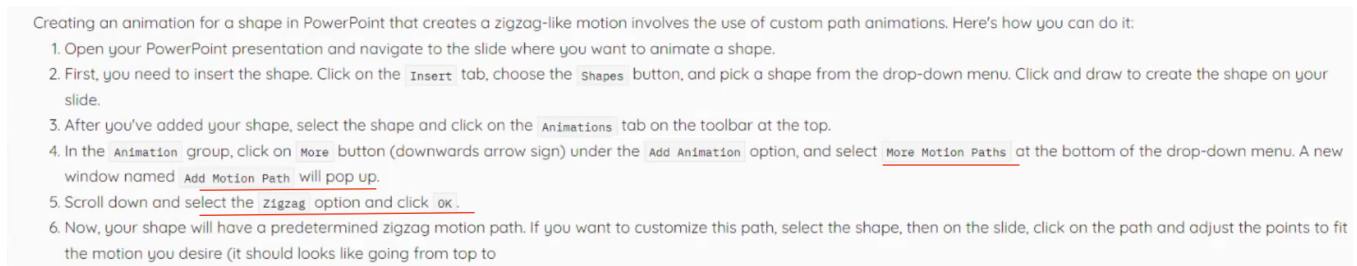
task accuracy score across all participants was 31.3% using SoftAIBot (maximum= 50%, minimum= 0%) and only 17.2% using Baseline ChatGPT (maximum= 50% and minimum= 0%). Although users achieved better task accuracy scores with SoftAIBot in comparison to Baseline, paired-sample t-test showed no significant difference across both interventions ( $t(26.6) = 1.7, p = 0.10$ , two-tailed), and we did not observe any order effects. The trend of non-significant accuracy persisted in Excel tasks as well ( $t(29.9) = 1.9, p = 0.06$ ).

**Users' Perceptions of Accuracy and Relevancy of LLM assistance:** Even though task completion scores were poor across both LLM interventions, in the self-report data, surprisingly, the majority of users perceived assistance from both LLMs to be accurate (SoftAIBot: 12/16 participants; Baseline ChatGPT: 8/16 participants) and relevant (SoftAIBot: 13/16 participants; Baseline ChatGPT: 9/16 participants) in completing the PowerPoint tasks. Pearson Chi-Squared test showed no significant difference in perceived accuracy ( $\chi^2(4, N = 28) = 3.05, p = 0.55$ ) and relevancy of LLM assistance ( $\chi^2(3, N = 28) = 1.46, p = 0.69$ ). Similarly for Excel tasks, there was no significant difference in perceived accuracy ( $\chi^2(3, N = 28) =$

$3.96, p = 0.27$ ) and relevancy ( $\chi^2(2, N = 28) = 0.29, p = 0.86$ ) across both LLM interventions.

In terms of other self-report data, such as users' perceptions of difficulty in applying assistance across both interventions, difficulty in figuring out correct input prompt, and ease of completing the task using LLM assistance, we did not observe any statistical difference across SoftAIBot and Baseline (See Figure 3 for the detail statistical results and test on remaining metrics). The one exception was the perception of trust as, interestingly, we observed that most participants (14/16) trusted SoftAIBot more than the Baseline ChatGPT (4/16) for the PowerPoint tasks and this result was significant ( $\chi^2(4, N = 32) = 14.40, p = 0.006$ ), but for the Excel task, there was no significant difference in users' perceptions of trust.

The key takeaways from our experiment were that while the expert rating showed that SoftAIBot performed better than the Baseline ChatGPT in producing more accurate results, users could not recognize the differences in accuracy and relevancy among both LLMs. Furthermore, having more accurate and relevant LLM output did not impact task completion nor task accuracy across both software applications. Next, we use our qualitative findings



**Figure 4: LLM Hallucination evidence (P15):** In response to P15’s prompt, “I want you to give me instructions on how to animate a shape that rotates from top to lower middle side and then come back up almost like a zigzag.”, Baseline ChatGPT generated the hallucinated response of Zigzag menu option (highlighted in red) which did not even exist in the software application.

to explain factors that impacted user performance and highlight some of the key challenges that users experienced in figuring out appropriate prompts and applying LLM assistance for different software tasks.

## 4.2 Inconsistency in Prompt-based Interaction

Prompt guidelines and prompt engineers usually suggest that breaking down the task and prompting it as a process for LLM can usually lead to desired LLM outputs. However, with the exception of 2 participants in our study who followed this approach and were successful (See example of Participant P02 in Figure 6), all of the other participants were inconsistent and varied in how they constructed prompts, failing to leverage the in-context prompt guidelines.

**Constructing Prompts as Search Queries:** Most participants (10/16) started with a generic “how to do the [task]” prompt, such as (e.g., “how to animate in Microsoft PowerPoint” (P14), “how to do correlation in Microsoft excel” (P04)). Most users were translating their mental model from other query-based systems, relying on similar queries they would issue on Google: “I think it [LLM] relies on the keywords that I am giving...at the beginning I just formulated a very vague question because it’s easier to get started with the vague question and then I can refine it as I go.” (P08)

Even after users experimented with different phrasings, they could not understand why the LLMs were producing nearly identical responses. These participants did not have an accurate mental model of how LLMs work and did not appear to recognize the impact of the prompt text on the quality of the LLM output: “I think it works same as a search engine. It has a back end and it takes your question through tons of data...it tries to give you an answer with all of that data that it has in the back end..it’s so quick that it goes through it within nanoseconds.” (P15)

Some participants (4/16), without even interpreting the task, just copy-pasted the entire task instructions along with some data sample (e.g., for Excel tasks) in hopes that the LLM would simplify the task and provide some instructions. However, both of these approaches were not that successful (as shown in Figure 5), as users had to ask follow-up questions on figuring out the correct steps. For example, when P07 (Figure 5) prompted the LLM using the keywords interpreted from the task, they struggled in getting an relevant response and went through several rounds of clarifications with the LLM. Only 2/16 participants who broke down the task and

drafted prompts as a process for SoftAIBot LLM obtained desired LLM outputs (See example of P02 in Figure 6)

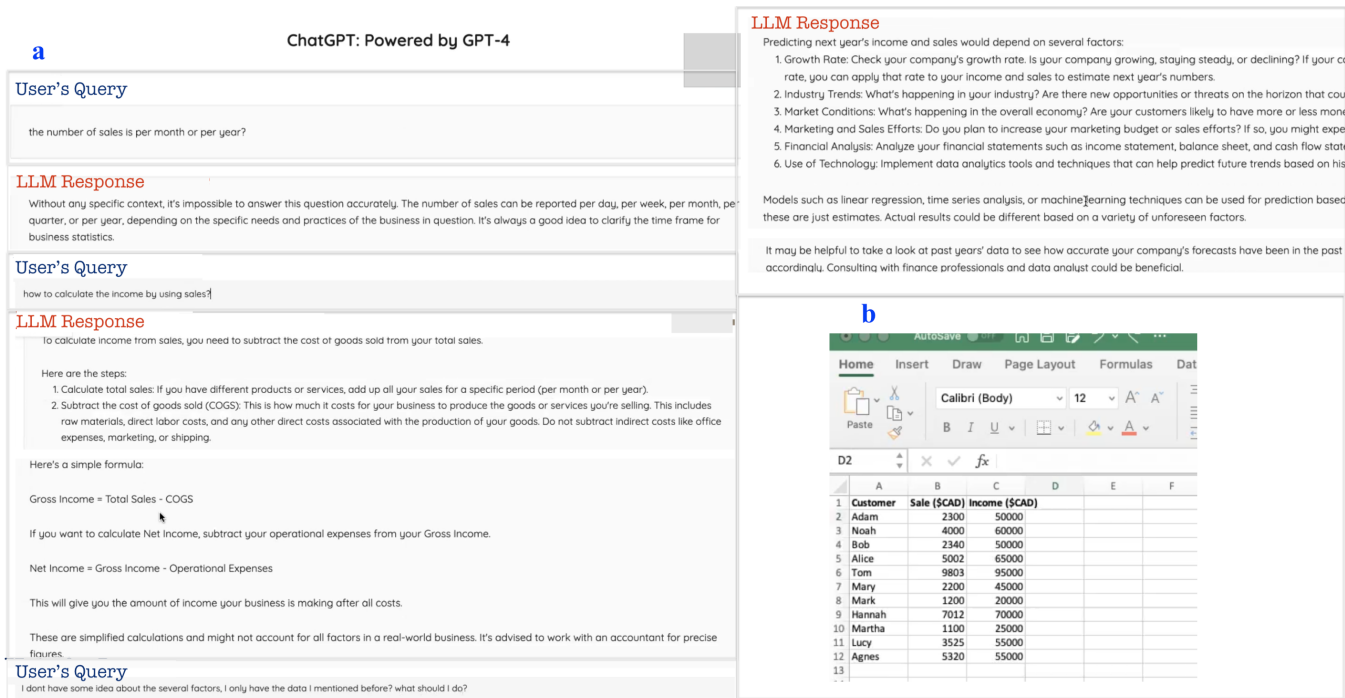
**Using Trial and Error Due to the Vocabulary Problem:** The majority of participants (11/16) also struggled in crafting prompts because they did not know how to express their task intent using software-specific terminology (often termed as the vocabulary problem [16]). This was especially prevalent during the PowerPoint tasks that involved references to different visual and interactive elements and participants frequently engaged in long trial-and-error episodes. Participants found it challenging to articulate their intended actions accurately and frequently blamed themselves: “I did not know how to describe those visual graphics in PowerPoint...I said words like flip and move but I am not sure if it was right for these kind of tasks. It took me a lot of time to try to understand which menu to select and not being sure if my prompt was ok or not... If the problem was my prompt or my system or the problem was the solution provided, I was not sure which one of them was making mistake.” (P05)

**Ignoring Prompt Guidelines:** Contrary to our hypothesis, only a handful of participants (5/16) employed the in-context prompt guidelines provided by SoftAIBot within the context of their queries and the majority simply ignored them. Participants expressed that the prompt guidelines were “not necessary” and “not useful” as it was faster for them to iterate on their own prompts. This behaviour was similar to the phenomenon commonly referred to as the “active user paradox” [10]. In fact, about half of the participants (7/16) were confident that they already possess the knowledge and experience required for generating prompts: “I would not say it is hurtful to have it [prompt guidelines] but is not necessary. I knew what to search for. I do not think prompt guidelines would have helped...” (P06) Some participants (4/16) noted that they were confident about crafting their own prompts because the LLMs were able to accept “any input” and generate corresponding output. If needed, they can revisit the generated response and assess their prompt alignment with their intended outcome for further improvement: “It is not difficult to figure out words because it [LLM] was accepting anything I typed. I can go back and verify whether that is what I need.” (P11)

## 4.3 User Perception of LLM Assistance

Some of the surprising results from our experiment were that users did not recognize the differences in accuracy and relevance between





**Figure 5: Unsuccessful Prompting by using keyword-based approach with Baseline ChatGPT: (a) P07 prompted the LLM using the keywords interpreted from the task, and struggled in getting an relevant response and went through several rounds of clarifications with the LLM; (b) Eventually, user could not even get started and failed to perform the task on the software application (P07).**

the two LLMs nor were they able to leverage the assistance to complete the software tasks accurately. Below, we discuss some factors that shaped users' perceptions and use of both LLMs.

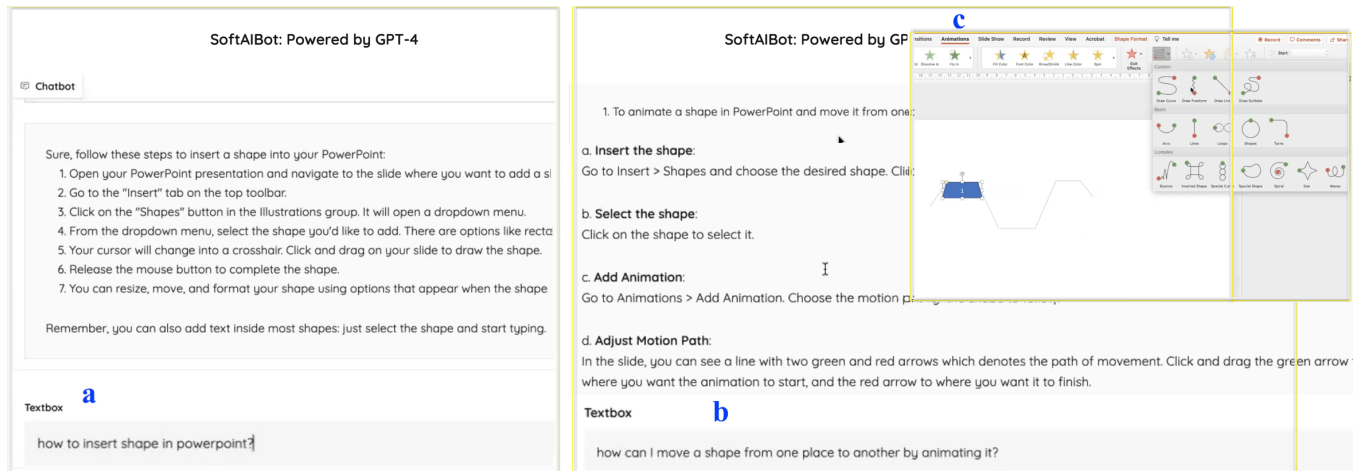
**LLMs' Ability to Generate Response Created an Impression of Credibility:** Most participants (13/16) perceived that both LLMs produced relevant responses because in contrast to systems like Siri and Google, which do not always provide response for every user's query, both LLM consistently generated a reasonable response matching the user's prompt. Thus, most participants formed an impression that the LLM outputs were credible: "... it [LLM] gave me what was relevant to my query...like some steps to find those options. This bot [LLM] was more detailed even than Siri. When you ask Siri, it's not giving what exactly I am looking for and keeps on giving me some different options...which is unnecessary about the topic. But in this one [LLM], what you put that's what you're getting. So the output is like, 90 to 95% near to what you just asked...so that made me trust it." (P12)

**Difficulty in Applying LLM Assistance to the Software Application:** In addition to the tensions in forming an accurate mental model of both LLMs (as described in Section 4.2), our participants who were infrequent users of Powerpoint and Excel also struggled because they lacked an accurate mental model of these applications. We observed that participants were quick to blame themselves for not being able to apply the instructions given by LLM due to lack of their familiarity with software: "For some reason, I got what it [LLM] tells me but...maybe it's giving me correct step,

but I am not able to apply to Excel because I'm not a regular user. I need to figure out [myself] how to use this formula; where can I put my formula and how to apply it." In comparison to SoftAIBot, users felt that the Baseline ChatGPT generated more generic or vague responses which users struggled to apply to software tasks. In fact, 9/16 participants began to doubt the credibility of Baseline ChatGPT: "Unless you apply the steps, you do not know whether that [Baseline ChatGPT] works or not." (P06)

In cases where the LLM output was actually relevant and accurate, our participants still struggled to locate menu options and apply LLM instructions. They explained that it was due to the lack of visual cues or guidance that are typically provided in tutorial videos and other forms of visual help: "...the instructions were here, but sometimes it gets difficult to use...If I get an image or graphical help along with screenshots of where to look for particular option, I could have made it slightly faster." (P11) Due to this difficulty in locating and applying LLM assistance, few participants (6/16) came up with different theories on whether LLM output did not considered their system version, as commented by P04: "It [LLM] did not provide instructions for the version of the software I'm using. Sometimes it's little bit difficult to apply the assistance, sometimes I would not find the button he [LLM] asked me to look for." (P04)

When the LLMs hallucinated or produced an incorrect set of instructions, we observed that more than half of the participants (10/16) could not map the (wrongly generated) LLM output to the intended features and menu functions in the application. As a result,



**Figure 6: Example of successful prompting by Participant P02. They started by breaking down the task, beginning with asking steps to (a) implement shape, followed by (b) animations in PowerPoint task. User prompted SoftAIBot as a step-by-step process asking steps for each functionality at one time (P02).**

they formed an incorrect mental model of the software application’s user interface. Instead of having awareness of this LLM’s bias of hallucination, these users felt burdened to make LLM instructions work: “For someone who is very new to PowerPoint...I could not find like the stuff, the menu names [ZigZag] that it was mentioning. It was not there in my application. I am not sure why all the burden of [finding] came to me.”(P15) Only a few participants (3/16) were able to recognize that the LLM does not always give factual information. They compared LLMs with Stack Overflow and Google search which they considered to be more credible than LLMs: “...if it’s a Stack Overflow, I would know that it’s just one person’s comment so I have ways to verify how trustworthy that instruction is...with Google, I have that much control over the source of information. But with LLM, I have no way to verify that at first sight. I have to follow it and, and decide on my own if it works or not. I prefer to be able to verify the credibility of a solution before actually going through the steps and putting more time to it.” (P07)

#### 4.4 Coherent LLM Output Leads to Blind Faith

**The presentation of LLM output fosters trust:** The most surprising finding from our study was that after receiving an output from the LLM, most users blindly followed the provided steps without a critical evaluation of the output’s veracity. For example, as illustrated in Figure 4, even when Baseline ChatGPT produced the hallucinated response of the *Zigzag menu* option which did not even exist in the software application, P15 still demonstrated unwavering trust of this LLM. The boundary between right and wrong for the participants while seeking LLM assistance, even in cases where it might yield incorrect results, was blurred as both LLMs consistently produced output that is coherent and in plain English, which enhances its perceived credibility. Especially for SoftAIBot, most participants (14/16) assumed that LLM is correct just because it generated a well-formatted response in response to their query: “I trust the system because once I get the match answer from him

[LLM]...makes me feel he [LLM] is helping and he is better than me.” (P03) Another user who followed SoftAIBot’s step-by-step instructions for PowerPoint’s visual tasks commented: “I was able to trust it [SoftAIBot] because I liked the way it gave these steps. SoftAIBot is more specific and gave me step-by-step sort of directions...it was user friendly and easy to follow. I could find all the steps that it [SoftAIBot] was referring to. Because this is what you need when you ask AI for help. You look for baby steps.” (P14)

## 5 DISCUSSION

We have contributed insights into how novices employ new-generation LLM assistants to seek software help, highlighting many of the challenges that users face while crafting prompts, comprehending how prompts bias LLM output, and mapping the LLM-suggested steps to software. Our key findings suggest that even though SoftAIBot outperformed the Baseline ChatGPT in providing relevant and accurate software-related assistance, there were no significant differences in users’ task completion rates or task accuracy scores between the two conditions. Notably, as opposed to our hypotheses H1-H2, there was no difference in users’ perceptions of accuracy and relevance for both LLMs, and users mostly failed to recognize instances where the models provided incorrect answers, including hallucinations. Our qualitative findings further shed light on our research questions and reveal a lack of awareness among participants regarding LLMs’ biases and limitations. Participants attributed their inability to complete a task and locate suggested features in the application to their personal shortcomings rather than recognizing instances of LLM hallucinations offering nonexistent options. Additionally, users misunderstood the prompt text’s influence on LLM output, likening prompts to traditional search engine keywords. User perception of LLM responses as contextually relevant stemmed from a bias towards the query context. Importantly, our study highlights the risks of undue trust in LLMs, as users frequently exhibited unwarranted confidence in LLM-generated responses due to their human-like

nature and consistent, contextual relevance, distinguishing them from traditional chatbots or virtual assistants like Siri.

The implications of our research extend beyond the immediate findings and have far-reaching significance for the broader IUI research community. Our observations highlight the need for end-users to exercise caution and critical thinking when relying on LLMs for software-related assistance. In an era where LLMs are increasingly integrated into various facets of daily life, from virtual assistants to content generation tools, understanding user perceptions and misconceptions about these systems is imperative. By shedding light on the lack of awareness regarding LLM biases and hallucinations, our study calls for a fundamental reevaluation of the way we design, deploy, and educate users about AI-powered assistants.

We now reflect on our key insights and highlight opportunities for designing LLM assistance for feature-rich software tasks while promoting transparent, responsible LLM interfaces to enhance user understanding and mental model formation. Our findings will be valuable for IUI and HCI researchers, interface designers, developers and others working on LLM-powered assistants.

### 5.1 Integrating LLM help into feature-rich applications

Our results demonstrate the value that LLM-based assistants, such as ChatGPT, can provide in generating relevant software-related assistance within a single platform. Unlike traditional help-seeking resources and chatbots (e.g., Google search, blogs, Siri, etc.) where users have to assimilate help content through multiple outlets, our participants appreciated receiving relevant detailed instructions by typing in a prompt. Having said that, our findings resonates with the speculations of other researchers [20] that Baseline GPT-4 is not meant to provide assistance for all types of tasks. Our SoftAIBot, that was optimized for particular feature-rich software guidance context, performed better and generated more relevant and accurate step-by-step software assistance. In our approach, we employed Retrieval Augmented Generation (RAG) on standard software documentation. Future developments could involve *instruction tuning* [53], which includes pairing more specific instructions with the software-specific steps and correlating this with expected output. The onus should transition from users to software developers and customer support to create such instructional pairs for ensuring them to be crafted in a manner that allows general-purpose LLMs to be fine-tuned [30] for generating user-centered and optimized software guidance.

While SoftAIBot generated relevant and accurate software assistance compared to Baseline ChatGPT, there were obvious limitations as users were not able to leverage this information to complete the software tasks accurately. In particular, users found the textual LLM output to be limiting compared to other visual-based help-seeking mediums. For example, software instructions on YouTube allow users to follow procedural steps “as is” without needing to verify and locate specific features. However, with LLMs, users had difficulty in mapping the LLM output to features in the software, especially in cases where the LLM was hallucinating and referred to non-existent features. Video-based help-seeking mediums should not be dismissed as instructional tools. Instead, to mitigate the issue

of locating the exact instructions that users experience with videos, there is an opportunity for technologies like ChatGPT to aid in video summarization tasks and to extract more relevant snippets from videos [15], enhancing the ease of locating specific information.

### 5.2 Transparent and Responsible Interface Design of LLMs

At a more fundamental level, our study raises some caution: while developers and researchers are investing in improving AI models and how LLMs can provide context-specific guidance, users may not always perceive these enhancements as substantial improvements in accuracy and relevance. Recent literature has already raised concerns about users’ over-reliance on AI systems, such as in the context of AI-based maze-solving tasks [49]. Although the landscape of user behaviors and mental models is more multifaceted with LLMs, our study demonstrates a similar phenomena of over-trust with LLMs. Furthermore, we extend prior works by revealing the nuances in users’ mental models of the LLMs triggered by the inherent biases of LLMs, leading to overtrust and users’ failure to recognize erroneous or hallucinated output: “*Because it is AI, how can it be wrong? I am going to stop using my brain as I literally gave it gibberish and still it works. I [will] doubt myself before doubting AI.*” (P15) Such overtrust in LLM assistants can be dangerous, especially for novice users who do not have familiarity with the underlying powerful AI technology. These findings from our study underscore the complexity of user interactions with AI and highlight the need for more transparency in addressing users’ expectations and misconceptions. This can be as critical as advancing the underlying AI technologies and it is essential to design interfaces that are more transparent and responsible [45]. There is need to consider more innovative user-centered solutions for mitigating bias and enhancing transparency in AI systems, thereby contributing to the responsible and ethical development of AI technologies.

To enhance the transparency of LLMs among end-users, one approach could be to embed interpretability within these systems. By articulating *why* and *how* LLMs derive specific recommendations, users can gain perspectives into the underlying mechanisms and trust the provided instructions with a higher degree of certainty. One possible direction is through the illustration of confidence percentage scores for each LLM instruction as these scores have shown to enhance the perception of transparency and trust [21]. Other potential direction is through explainability techniques [37] (e.g., visual example-based explanations [9, 21]) that can indicate why the system did what it did and verify an AI’s recommendation [13] by demonstrating the similarities between users’ intent and examples in the training set [9, 21, 27, 28]. Training datasets of these LLMs needs to be designed such that confidence scores or visual examples are part of the dataset to enhance transparency within the LLM technologies. Such innovative advancements in LLM development not only pave the way for enhanced user interaction but also ensure that the model’s suggestions are verifiable and reliable.

### 5.3 Bridging the Gap Between Mental Models and LLM Interfaces

With the rapid pace of innovations in Generative AI and LLMs, there is need for HCI research to focus on understanding user

perceptions and users' mental models of LLM-based assistants for software help-seeking. Our study complements existing works highlighting usability issues in crafting prompts [5, 41, 51] and provides initial insights into the gaps in end-users' mental model when using prompt-based interactions in context of software help-seeking. The affordances of LLMs can be misleading as they are designed to be walk-up-and-use and support natural language interaction. However, similar to insights from recent work on non-ML expert designers prototyping ML apps [52], we also found that crafting effective prompts is cumbersome, especially for non-AI experts. Our study contributes new knowledge: non-AI expert end-users of ML/LLM applications must adapt their existing mental models from traditional help-seeking mediums to understand the new interface of LLMs. To bridge the gap between users' mental models and LLM user interfaces, there is an urgent need to leverage strategies such as think-aloud studies to further understand nuances in users' mental models of LLMs [33]. Although we did not see significant individual differences in our sample, it may be worth investigating how different sub-groups of software users might benefit from LLM help seeking. Our study demonstrates that we cannot assume users will intuitively grasp the capabilities and limitations of LLMs. There is a clear need for comprehensive user training and education and clear communication about how Generative AI systems operate.

With software help-seeking, the challenge for end-users lies not only in flawed mental models of LLMs but also in the absence of a clear understanding of the underlying software application. Our research complements the emerging work in this space, being novel in documenting the challenges users encounter with LLMs for software help tasks, including their mental models and overtrust, which impacts both their utilization of LLMs and perception of software features, regardless of LLM optimization or explicit prompt guidelines (SoftAIBot). Users found it difficult to understand, map, and apply LLM instructions to software features. To address this, enhancing the UX design of LLM interfaces by highlighting relevant software UI sections during onboarding can improve user interaction, particularly for feature-rich software [21, 27, 28]. Exploring the interplay between users' understanding of LLMs and the underlying software presents new human-AI design possibilities. Overall, our findings show that LLMs optimized for generating software specific guidance (e.g., Copilot [43]) and embedded inside feature-rich applications, could be promising for learning and using complex features. Once these systems become available, future studies can compare our findings from SoftAIBot with such systems and further investigate the level of guidance and automation that may be appropriate for LLMs generating software guidance.

## 6 LIMITATIONS

In this paper, we experimented with two LLM-based assistants to explore how end-users make use of them for software tasks related to feature-rich applications. While our findings shed new light on users struggles in employing LLMs for software help-seeking, some caution should be used in interpreting our results. For instance, our findings could be constrained by the specific applications utilized during the experimentation. Whether our findings would generalize beyond the state-of-the-art LLM implementations used in the study should be investigated in future work. Given the rapid evolution and

variability among modern LLMs, the outcomes may be constrained to currently available LLMs. Although we recruited end-users who were non-AI experts, we did not control for other individual differences, such as expertise or familiarity with underlying feature-rich software applications. Future studies should conduct experiments and more qualitative studies with larger and varying populations to better capture these individual differences.

## 7 CONCLUSIONS

In this research, we investigated the effectiveness of the LLM-generated software guidance and prompt guidelines, by comparing the Baseline LLM assistant with our own implemented SoftAIBot in providing accurate and relevant assistance for software tasks. Our results highlight the challenges users faced in following LLM assistance and point to instances of users attributing flaws to themselves instead of recognizing LLM biases. Our study highlights the pressing need for interdisciplinary collaboration among researchers, designers, developers, and educators to bridge the gap between user expectations and AI realities. By addressing these challenges head-on, we can foster a future where AI systems are not only more powerful but also more comprehensible and accountable to their users, ultimately facilitating human-AI interaction across a range of domains.

## ACKNOWLEDGMENTS

We thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for funding this research.

## REFERENCES

- [1] Adobe. 2023. Adobe unveils Firefly, a family of new creative generative ai. <https://news.adobe.com/news/news-details/2023/Adobe-Unveils-Firefly-a-Family-of-new-Creative-Generative-AI/default.aspx>
- [2] Open AI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>
- [3] Oscar D. Andrade, Nathaniel Bean, and David G. Novick. 2009. The Macro-Structure of Use of Help. In *Proceedings of the 27th ACM International Conference on Design of Communication* (Bloomington, Indiana, USA) (SIGDOC '09). Association for Computing Machinery, New York, NY, USA, 143–150. <https://doi.org/10.1145/1621995.1622022>
- [4] Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023* (2023).
- [5] Shraddha Barke, Michael B. James, and Nadia Polikarpova. 2023. Grounded Copilot: How Programmers Interact with Code-Generating Models. *Proc. ACM Program. Lang.* 7, OOPSLA1, Article 78 (apr 2023), 27 pages. <https://doi.org/10.1145/3586030>
- [6] Ali Borji. 2023. A categorical archive of chatgpt failures. *arXiv preprint arXiv:2302.03494* (2023).
- [7] Joel Brandt, Mira Dontcheva, Marcos Weskamp, and Scott R. Klemmer. 2010. Example-Centric Programming: Integrating Web Search into the Development Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 513–522. <https://doi.org/10.1145/1753326.1753402>
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac12f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac12f64a-Paper.pdf)
- [9] Carrie J. Cai, Jonas Jongejan, and Jess Holbrook. 2019. The Effects of Example-Based Explanations in a Machine Learning Interface. In *Proceedings of the 24th*

- International Conference on Intelligent User Interfaces* (Marina del Ray, California) (IUI '19). Association for Computing Machinery, New York, NY, USA, 258–262. <https://doi.org/10.1145/3301275.3302289>
- [10] John M. Carroll and Mary Beth Rosson. 1987. *Paradox of the Active User*. MIT Press, Cambridge, MA, USA, 80–111.
- [11] Parmit K. Chilana, Amy J. Ko, and Jacob O. Wobbrock. 2012. LemonAid: Selection-Based Crowdsourced Contextual Help for Web Applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 1549–1558. <https://doi.org/10.1145/2207676.2208620>
- [12] Juliet M Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology* 13, 1 (1990), 3–21.
- [13] Raymond Fok and Daniel S Weld. 2023. In Search of Verifiability: Explanations Rarely Enable Complementary Performance in AI-Advised Decision Making. *arXiv preprint arXiv:2305.07722* (2023).
- [14] Adam Fournery, Ben Lafreniere, Parmit Chilana, and Michael Terry. 2014. InterTwine: Creating Interapplication Information Scent to Support Coordinated Use of Software. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 429–438. <https://doi.org/10.1145/2642918.2647420>
- [15] C. Ailie Fraser, Julia M. Markel, N. James Basa, Mira Dontcheva, and Scott Klemmer. 2020. ReMap: Lowering the Barrier to Help-Seeking with Multimodal Search. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 979–986. <https://doi.org/10.1145/3379337.3415592>
- [16] G. W. Furnas, T. K. Landauer, L. M. Gomez, and S. T. Dumais. 1987. The Vocabulary Problem in Human-System Communication. *Commun. ACM* 30, 11 (nov 1987), 964–971. <https://doi.org/10.1145/32206.32212>
- [17] Tovi Grossman and George Fitzmaurice. 2010. ToolClips: An Investigation of Contextual Video Assistance for Functionality Understanding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 1515–1524. <https://doi.org/10.1145/1753326.1753552>
- [18] Tovi Grossman, George Fitzmaurice, and Ramtin Attar. 2009. A Survey of Software Learnability: Metrics, Methodologies and Guidelines. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 649–658. <https://doi.org/10.1145/1518701.1518803>
- [19] Björn Hartmann, Daniel MacDougall, Joel Brandt, and Scott R. Klemmer. 2010. What Would Other Programmers Do: Suggesting Solutions to Error Messages. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 1019–1028. <https://doi.org/10.1145/1753326.1753478>
- [20] Samantha Murphy Kelly. 2023. 5 jaw-dropping things GPT-4 can do that chatgpt couldn't | CNN business. <https://www.cnn.com/2023/03/16/tech/gpt-4-use-cases/index.html>
- [21] Anjali Khurana, Parsa Alamzadeh, and Parmit K. Chilana. 2021. ChatrEx: Designing Explainable Chatbot Interfaces for Enhancing Usefulness, Transparency, and Trust. In *2021 IEEE Symposium on Visual Languages and Human-Centric Computing* (VL/HCC), 1–11. <https://doi.org/10.1109/VL/HCC51201.2021.9576440>
- [22] Kimia Kiani, George Cui, Andrea Bunt, Joanna McGrenere, and Parmit K. Chilana. 2019. Beyond "One-Size-Fits-All": Understanding the Diversity in How Software Newcomers Discover and Make Use of Help Resources. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3290605.3300570>
- [23] Juho Kim, Phu Tran Nguyen, Sarah Weir, Philip J. Guo, Robert C. Miller, and Krzysztof Z. Gajos. 2014. Crowdsourcing Step-by-Step Information Extraction to Enhance Existing How-to Videos. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 4017–4026. <https://doi.org/10.1145/2556288.2556986>
- [24] Benjamin Lafreniere, Parmit K. Chilana, Adam Fournery, and Michael A. Terry. 2015. These Aren't the Commands You're Looking For: Addressing False Feedforward in Feature-Rich Software. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (Charlotte, NC, USA) (UIST '15). Association for Computing Machinery, New York, NY, USA, 619–628. <https://doi.org/10.1145/2807442.2807482>
- [25] Benjamin Lafreniere, Tovi Grossman, and George Fitzmaurice. 2013. Community Enhanced Tutorials: Improving Tutorials with Multiple Demonstrations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). Association for Computing Machinery, New York, NY, USA, 1779–1788. <https://doi.org/10.1145/2470654.2466235>
- [26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) (NIPS'20). Curran Associates Inc., Red Hook, NY, USA, Article 793, 16 pages.
- [27] Toby Jia-Jun Li, Jingya Chen, Haijun Xia, Tom M. Mitchell, and Brad A. Myers. 2020. Multi-Modal Repairs of Conversational Breakdowns in Task-Oriented Dialogs. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 1094–1107. <https://doi.org/10.1145/3379337.3415820>
- [28] Toby Jia-Jun Li, Igor Labutov, Xiaohan Nancy Li, Xiaoyi Zhang, Wenzhe Shi, Wanling Ding, Tom M. Mitchell, and Brad A. Myers. 2018. APPINITE: A Multi-Modal Interface for Specifying Data Descriptions in Programming by Demonstration Using Natural Language Instructions. In *2018 IEEE Symposium on Visual Languages and Human-Centric Computing* (VL/HCC), 105–114. <https://doi.org/10.1109/VLHCC.2018.8506506>
- [29] Q Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *arXiv preprint arXiv:2306.01941* (2023).
- [30] Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. 2022. Few-Shot Parameter-Efficient Fine-Tuning is Better and Cheaper than In-Context Learning. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 1950–1965. [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/0cde695b83bd186c1fd456302888454c-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/0cde695b83bd186c1fd456302888454c-Paper-Conference.pdf)
- [31] HARRY MCCRACKEN. 2023. Microsoft's Satya Nadella is winning Big Tech's AI War. here's how. <https://www.fastcompany.com/90931084/satya-nadella-microsoft-ai-frontrunner>
- [32] Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. MTEB: Massive text embedding benchmark. *arXiv preprint arXiv:2210.07316* (2022).
- [33] Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Basic books.
- [34] David G. Novick, Oscar D. Andrade, and Nathaniel Bean. 2009. The Micro-Structure of Use of Help. In *Proceedings of the 27th ACM International Conference on Design of Communication* (Bloomington, Indiana, USA) (SIGDOC '09). Association for Computing Machinery, New York, NY, USA, 97–104. <https://doi.org/10.1145/1621995.1622014>
- [35] David G. Novick, Oscar D. Andrade, Nathaniel Bean, and Edith Elizalde. 2008. Help-Based Tutorials. In *Proceedings of the 26th Annual ACM International Conference on Design of Communication* (Lisbon, Portugal) (SIGDOC '08). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/1456536.1456538>
- [36] David G. Novick and Karen Ward. 2006. Why Don't People Read the Manual?. In *Proceedings of the 24th Annual ACM International Conference on Design of Communication* (Myrtle Beach, SC, USA) (SIGDOC '06). Association for Computing Machinery, New York, NY, USA, 11–18. <https://doi.org/10.1145/1166324.1166329>
- [37] Samir Passi and Mihaela Vorvoreanu. 2022. *Overreliance on AI: Literature Review*. Technical Report MSR-TR-2022-12. Microsoft. <https://www.microsoft.com/en-us/research/publication/overreliance-on-ai-literature-review/>
- [38] Michael L Raulin and Anthony M Graziano. 2019. Quasi-experiments and correlational studies. In *Companion Encyclopedia of Psychology*. Routledge, 1122–1141.
- [39] Marc Rettig. 1991. Nobody Reads Documentation. *Commun. ACM* 34, 7 (Jul 1991), 19–24. <https://doi.org/10.1145/105783.105788>
- [40] Ted Sanders. 2023. How to format inputs to CHATGPT models. [https://cookbook.openai.com/examples/how\\_to\\_format\\_inputs\\_to\\_chatgpt\\_models](https://cookbook.openai.com/examples/how_to_format_inputs_to_chatgpt_models)
- [41] Advait Sarkar, Andrew D Gordon, Carina Negreanu, Christian Poelitz, Sruti Srinivasa Ragavan, and Ben Zorn. 2022. What is it like to program with artificial intelligence? *arXiv preprint arXiv:2208.06213* (2022).
- [42] Jessica Shieh. 2023. Best practices for prompt engineering with openai API: Openai help center. <https://help.openai.com/en/articles/6654000-best-practices-for-prompt-engineering-with-openai-api>
- [43] Jared Spataro. 2023. Introducing Microsoft 365 copilot – your copilot for work. <https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/>
- [44] Suhriddpalsule. 2023. Prompt engineering techniques with Azure Openai - Azure Openai Service. <https://learn.microsoft.com/en-us/azure/ai-services/openai/concepts/advanced-prompt-engineering?pivots=programming-language-chat-completions>
- [45] Jiao Sun, Q. Vera Liao, Michael Muller, Mayank Agarwal, Stephanie Houde, Kartik Talamadupula, and Justin D. Weisz. 2022. Investigating Explainability of Generative AI for Code through Scenario-Based Design. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) (IUI '22). Association for Computing Machinery, New York, NY, USA, 212–228. <https://doi.org/10.1145/3490099.3511119>
- [46] Gradio Team. 2023. <https://www.gradio.app/>
- [47] Haoye Tian, Weiqi Lu, Tsz On Li, Xunzhu Tang, Shing-Chi Cheung, Jacques Klein, and Tegawendé F Bissyandé. 2023. Is ChatGPT the Ultimate Programming Assistant—How far is it? *arXiv preprint arXiv:2304.11938* (2023).

- [48] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI EA '22*). Association for Computing Machinery, New York, NY, USA, Article 332, 7 pages. <https://doi.org/10.1145/3491101.3519665>
- [49] Helena Vasconcelos, Matthew Jörke, Madeleine Grunde-McLaughlin, Tobias Gerstenberg, Michael S. Bernstein, and Ranjay Krishna. 2023. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 129 (apr 2023), 38 pages. <https://doi.org/10.1145/3579605>
- [50] Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382* (2023).
- [51] Frank F. Xu, Bogdan Vasilescu, and Graham Neubig. 2022. In-IDE Code Generation from Natural Language: Promise and Challenges. *ACM Trans. Softw. Eng. Methodol.* 31, 2, Article 29 (mar 2022), 47 pages. <https://doi.org/10.1145/3487569>
- [52] J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (*CHI '23*). Association for Computing Machinery, New York, NY, USA, Article 437, 21 pages. <https://doi.org/10.1145/3544548.3581388>
- [53] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792* (2023).

## A STUDY MATERIALS AND TASKS

We provide detailed information on the study questionnaires, interview and tasks asked to the participants during the experiment.

### A.1 Post-task questionnaire

We asked participants to fill the post-task questionnaire to evaluate each LLM interventions after performing the software tasks. Participants were presented with a Likert scale consisting of five response options, prompting them to express their level of agreement or disagreement in response to the presented questions.

- The LLM provided me with the accurate assistance to complete the given task correctly.
- Using the LLM made it easier for me to finish the software task.
- It was difficult for me to apply assistance provided by the LLM in the software task.
- It was difficult for me to figure out correct words (input prompt) to seek desired help for the task.
- The assistance provided by the LLM was relevant to what I asked in the prompt.
- I trusted the LLM assistance to finish my software task.

### A.2 Follow-up interview

During the follow-up interview, we asked following questions to the participants to better understand the difficulties they faced while seeking LLM assistance for software tasks.

- Describe your experience in seeking the assistance from the LLM assistants for performing the tasks?
- When you normally work with a new software application or unfamiliar features, how do you typically get help?
  - Have you used LLM assistants before such as ChatGPT, Github Copilot, BARD, etc? If yes, can you compare your experience with these assistants to seek help for software tasks during the study vs other tasks you have performed before?
- Now that you have seen and interacted with 2 different types of LLM assistants, how would you rank them in terms of following metrics in providing required assistance
  - Relevance: such that Rank 1 indicates the LLM assistant which provided strongly relevant help for your prompts and tasks and Rank 2 indicates not at all relevant? Why?
  - Accuracy: Which LLM provided you with accurate assistance for your tasks. Please rank them where Rank 1 will be highly accurate and Rank 2 is least accurate? Why?
  - Explainability: Which LLM helped you improve your next query or prompt to obtain the desired assistance for your task. Please rank them where Rank 1 will be highly explainable and Rank 2 is least explainable? Why?
- How do you think the LLM assistant interface worked (i.e., what it did or did not do with respect to your tasks) in LLMs used in the study to accomplish your tasks? (*You can create a sketch or describe in steps using your task example*)
- Please describe what worked well while seeking from (a) Baseline ChatGPT (b) SoftAIBot - any LLM output/messages, prompt guidelines, etc., that were helpful for you in obtaining the desired help?

- Can you describe concerns while using the LLM assistants (Baseline ChatGPT, SoftAIBot) to improve the prompt and obtain desired help- what did not work for you while using the LLM assistants for your tasks?

### A.3 Description of Study Tasks

We describe the actual software tasks the users had to conduct seeking assistance from LLM intervention.

- **Microsoft PowerPoint tasks:**
  - Participants were asked to look up and use instructions from LLM intervention to create an interactive quiz (named as “EasyQuiz”) in Microsoft PowerPoint to create questions where users can type in their response.
  - Participants were asked to look up and use instructions from LLM intervention to develop a project timeline in Microsoft PowerPoint that is visual and animated and looks like the reference animated project timeline (Participants were provided with the reference project timeline).
- **Microsoft Excel tasks:** (For both the tasks, participants were provided with the corresponding dataset for analysis)
  - Participants were asked to look up and use instructions from LLM intervention to find the predictive analytics of the sales values from each customer by considering the income of customer in Microsoft Excel. Based on this analysis, predict the estimated sales of a new customer who has an income of \$40k?
  - Participants were asked to look up and use instructions from LLM intervention to analyze and visualize employees’ performance trends and how performance correlate with the average attendance (i.e., absent hours).

## B TECHNICAL DETAILS OF LLM INTERVENTION

This section lists technical details of prompts and parameters used in OpenAI APIs calls in both LLM interventions.

### B.1 Parameters used in OpenAI ChatCompletion objects:

For both the LLM interventions, we used the same parameter values except the ‘messages’ parameter which had varied values.

- **Model: gpt-4-0613; Temperature: 1; Max Tokens: 400; Top-p (nucleus sampling): 1; Frequency Penalty: 0; Presence Penalty: 0;**  
The temperature is set to 1, introducing randomness to the model’s output. The max tokens parameter limits the response length to 400 tokens. Top-p is set to 1, indicating that the model considers the entire probability mass for sampling. Frequency and presence penalties are both set to 0, indicating no penalty for frequent or rare tokens.
- The “messages” parameter accepts a list comprising initial system message and subsequent user interactions that are dynamically appended to maintain context.  
**messages: [Empty list];** For Baseline ChatGPT, messages parameter starts with an empty message list, allowing users to initiate conversations without a predefined LLM role.

**messages: Role: System; Content: “you are an expert in PowerPoint”**; For SoftAIBot, the system message explicitly defines it as an expert in Microsoft PowerPoint. This provides a contextual foundation for subsequent interactions, emphasizing the LLM’s proficiency in addressing queries related to PowerPoint.

## B.2 RAG Prompt for SoftAIBot

The RAG prompt is constructed to guide SoftAIBot in the knowledge retrieval process:

```
{"role": "user", "content": f"you have the\nfollowing information: \n {top_chunks\n}"} , {"role": "user", "content": f"now\nhelp me in finding the answer to this\nquestion: {user_query}"}
```

The first user prompt provides SoftAIBot with information from the top-retrieved chunks (`top_chunks`) during the knowledge retrieval process from the software documentation. This ensures that the LLM is aware of the context and relevant details extracted from the software documentation. The second user prompt instructs SoftAIBot to assist in finding the answer to a query asked by the user in the chat. This step triggers the LLM to generate a response based on the acquired knowledge, facilitating a seamless interaction between the user and SoftAIBot.

The code repository for both LLM interventions is available upon request.