
USER INTENT RECOGNITION AND SATISFACTION WITH LARGE LANGUAGE MODELS: A USER STUDY WITH CHATGPT

A PREPRINT

Anna Bodonhelyi

Chair for Human-Centered Technologies for Learning
Technical University of Munich
Munich, Bavaria, Germany
anna.bodonhelyi@tum.de

Efe Bozkir

Chair for Human-Centered Technologies for Learning
Technical University of Munich
Munich, Bavaria, Germany
efe.bozkir@tum.de

Shuo Yang

Chair for Responsible Data Science
Technical University of Munich
Munich, Bavaria, Germany
shuo.yang@tum.de

Enkelejda Kasneci

Chair for Human-Centered Technologies for Learning
Technical University of Munich
Munich, Bavaria, Germany
enkelejda.kasneci@tum.de

Gjergji Kasneci

Chair for Responsible Data Science
Technical University of Munich
Munich, Bavaria, Germany
gjergji.kasneci@tum.de

January 2024

ABSTRACT

The rapid evolution of large language models such as GPT-4 Turbo represents an impactful paradigm shift in digital interaction and content engagement. While these models encode vast amounts of human-generated knowledge and excel in processing diverse data types, recent research shows that they often face the challenge of accurately responding to specific user intents, leading to increased user dissatisfaction. Based on a fine-grained intent taxonomy and intent-based prompt reformulations, we analyze (1) the quality of intent recognition and (2) user satisfaction with answers from intent-based prompt reformulations for two recent ChatGPT models, GPT-3.5 Turbo and GPT-4 Turbo. The results reveal that GPT-4 outperforms GPT-3.5 on the recognition of common intents, but is conversely often outperformed by GPT-3.5 on the recognition of less frequent intents. Moreover, whenever the user intent is correctly recognized, while users are more satisfied with the answers to intent-based reformulations of GPT 4 compared to GPT-3.5, they tend to be more satisfied with the answers of the models to their original prompts compared to the reformulated ones. Finally, the study indicates that users can quickly learn to formulate their prompts more effectively, once they are shown possible reformulation templates.

Keywords large language models, ChatGPT, user study, intent recognition, prompt reformulation, intent taxonomy, human-AI interaction

1 Introduction

Generative AI models, especially those handling language and vision, are rapidly evolving. Models like OpenAI's GPT-4 Turbo [OpenAI et al., 2023, OpenAI, 2023] and Google's Gemini [DeepMind, 2023] are at the forefront of this evolution, impacting how we interact with digital content and services. At the heart of this development are Large

Language Models (LLMs). With human-like text processing, generation, and reasoning capabilities, LLMs have broad applications ranging from creative content generation to complex problem-solving [Brown et al., 2020, Romera-Paredes et al., 2023]. The incorporation of human feedback and reinforcement learning in the training of LLMs, such as those used in ChatGPT, has further improved the alignment of these models with societal norms and goals [Ouyang et al., 2022, Lee et al., 2023]. Similarly, models like Google’s PaLM [Chowdhery et al., 2022, Anil et al., 2023] and Meta’s LLaMA [Touvron et al., 2023a,b], together with open-source variants Vicuna [Zheng et al., 2023] and Alpaca [Taori et al., 2023], represent parallel advances in the field.

LLMs have already had impact in diverse sectors such as healthcare [Chintagunta et al., 2021, Enarvi et al., 2020], finance [Dowling and Lucey, 2023], journalism [Pavlik, 2023], and creative writing [Yuan et al., 2022, Seßler et al., 2023]. Notably, they have also contributed to scientific discoveries [Romera-Paredes et al., 2023], despite serious challenges in areas like mathematical reasoning [Imani et al., 2023], student error detection [Bewersdorff et al., 2023], and mitigating hallucinations in outputs [Ji et al., 2023, Azamfirei et al., 2023, Zhang et al., 2023, Manakul et al., 2023]. Efforts towards enhancing explainability are underway to foster trust and reliability in these models [Wu et al., 2023a].

Building on the foundations of LLMs, multimodal generative AI models have further expanded the scope to encompass visual, auditory, and other sensory data [Alayrac et al., 2022, Zhu et al., 2023, Huang et al., 2023, Ye et al., 2023, Li et al., 2023, Maaz et al., 2023, Driess et al., 2023, Chen et al., 2023, Su et al., 2023, Wu et al., 2023b, Shen et al., 2023]. Hence, the capabilities of LLMs to reason over natural language are essential for many of the recent breakthroughs in generative AI, such as NEX-T-GPT [Wu et al., 2023c], OpenAI’s GPT-4 Vision, GPT-4 Turbo, Google’s Gemini, and Apple’s Ferret [OpenAI et al., 2023, DeepMind, 2023, You et al., 2023]. Such advancements, coupled with other deep learning and explainability techniques [Borisov et al., 2022, Rombach et al., 2022], can potentially revolutionize science [Wong et al., 2023] and society.

Trained on extensive human-generated datasets and various web corpora, LLMs encode a vast amount of knowledge and excel in processing and reasoning tasks over text. Yet, the effectiveness of their responses strongly depends on accurately inferring the user’s intent, which is typically conveyed implicitly in the prompt. Accurately discerning and categorizing user intents in prompts poses considerable challenges for LLMs, which are mainly related to the inherent ambiguity, variability, or the clarity of the used language, cultural and contextual subtleties, as well as evolving user expectations. A recent study [Kim et al., 2023] highlighted a frequent source of user dissatisfaction with ChatGPT: its occasional failure to understand user intentions. The authors further observed that users often lack effective strategies to mitigate this dissatisfaction. Moreover, it was noted that users with a limited understanding of LLMs tend to experience greater dissatisfaction and are less proactive in addressing these challenges.

In this paper, we explore the hypothesis that enhancing the understanding of a user’s specific intent in their prompt can significantly improve the quality of responses generated by Large Language Models (LLMs). This assumption finds substantial backing in the recent advancements in natural language processing methodologies, notably the Chain-of-Thought, Tree-of-Thought, and Graph-of-Thought techniques [Wei et al., 2022, Yao et al., 2023, Besta et al., 2023]. These approaches emphasize the importance of contextual and hierarchical understanding in processing user queries. To systematically approach this hypothesis, we first establish a comprehensive taxonomy of potential user intents. This taxonomy is carefully crafted, taking into account the distinct requirements of different intent types and integrating insights from recent scholarly work on intent categorization [Dang et al., 2022, Shah et al., 2023] and well-established intent categories from information retrieval [Azad and Deepak, 2019].

Our study then proceeds in two critical phases. First, we evaluate recent LLM versions, namely GPT-3.5 and GPT-4, to determine their proficiency in accurately recognizing the intent behind user prompts. We expect more recent LLM versions to outperform older ones on this task. This aspect is crucial, as the correct identification of intent is foundational to providing relevant and accurate responses. The second phase of our study focuses on the ramifications of prompt reformulation. Here, we explore whether accurately recognizing and then reformulating a user’s prompt to better mirror their intended query leads to an enhancement in the quality of the LLM’s responses. This phase is particularly focused on evaluating the response quality from the perspective of user satisfaction and relevance, thereby addressing a significant aspect of human-computer interaction in the realm of natural language understanding. This approach not only provides insights into the current capabilities of LLMs in intent recognition but also sheds light on the potential benefits of prompt reformulation in enhancing user experience.

In summary, our work provides the following contributions:

1. The development of a taxonomy of potential intents in conversational contexts.
2. Assessment of the intent recognition proficiency of GPT-3.5 and GPT-4 based on a user study.
3. Exploration of the impact of prompt reformulation on response quality, focusing on user satisfaction and relevance in conversational context.

4. Insights into LLM capabilities and the potential of prompt reformulation to enhance user experience.
5. An extensive data set including an overview of our user study is published on [GitHub](#) for research purposes.

Section 2 presents existing literature on intent recognition and prompt optimization. We introduce our intent taxonomy in Section 3, laying the groundwork for the user study in Section 4. The results of the study are presented and discussed in Section 5, before concluding in Section 6.

2 Related Work on Intent Recognition in User Prompts and Prompt Optimization

Effectively understanding user prompts is crucial for better interaction with LLMs. This section covers two key aspects of related works in this context: first, intent recognition in user prompts, reviewing methods and classifications developed to understand the user’s intent in the LLM context; and second, prompt optimization, focusing on the latest techniques in prompt engineering to improve user-model interactions. These two research directions are important to emphasize the complex interplay between the user’s intent and the model’s response.

Intent recognition in user prompts: The field of Information Retrieval (IR) has thoroughly examined the concept of intent, leading to various methodologies for understanding and representing it. Search intent classification enables systems to better understand and respond to user objectives. A number of these classifications have been suggested in the literature [Broder, 2002, Jansen et al., 2007], and taxonomies were developed through iterative processes, primarily involving the manual analysis of search log data. Common categories identified from these logs include: *navigational*, *informational*, *transactional*, *browsing*, and *resource-seeking* intents. However, with the advent of LLMs, the types and dynamics of user interactions are evolving, particularly in terms of *content generation* as opposed to traditional search methods. Consequently, there is ongoing research into new intent classifications tailored to represent intents in these specific interaction contexts. Our research builds upon a recent work [Shah et al., 2023], where the authors introduced an LLM-based approach with human-in-the-loop to generate intent taxonomies.

Our work significantly advances [Shah et al., 2023] and introduces a detailed and nuanced taxonomy of user intents specifically tailored to capture the evolving landscape of user interactions with multimodal LLMs.

Prompt optimization An emerging field in natural language processing is *prompt engineering*, which focuses on designing and optimizing prompts to effectively interact with and guide LLMs. Various techniques have been developed in this domain to enhance the quality of interactions [Liu et al., 2023]. Prompt engineering explores adding text or vectors to inputs and outputs of LLMs to streamline interactions without altering the core parameters of the model, offering thus an efficient alternative to traditional fine-tuning in scenarios with limited data [Min et al., 2023]. Such prompting strategies are pivotal in models like the GPT series, emphasizing instructions and demonstrations, and in template-based learning, where examples are integrated into natural text formats [Min et al., 2023]. Furthermore, zero- and few-shot learning have emerged as powerful strategies in prompt engineering. Zero-shot learning involves crafting prompts that enable the model to generate useful responses without prior examples, relying solely on pre-trained knowledge. In contrast, few-shot learning involves providing a few examples within the prompt, guiding the model to understand the task context and desired response format. This technique has shown remarkable effectiveness in adapting models to new tasks with minimal examples. Several works have recently focused on few-shot prompting techniques, particularly to improve the reasoning capabilities of LLMs, such as the Chain-of-Thought [Wei et al., 2022], Tree-of-Thoughts [Yao et al., 2023], or Graph-of-Thoughts prompting [Besta et al., 2023].

Another technique to significantly improve the quality of user interactions and enhance user experiences during interactions with LLMs is prompt optimization through reformulation [Wang et al., 2023]. By automating the process of prompt refinement, such strategies ensure higher accuracy and contextual appropriateness of the responses of LLMs. One fundamental approach in this context is templating, where prompts are structured in a consistent format to elicit specific types of responses. This method relies on careful wording and phrasing to guide the model towards the desired output.

Additionally, using context reiteration and clarification in prompts has proven effective. By repeating or rephrasing key parts of the prompt, the clarity and focus of the model’s responses can be significantly enhanced [Dang et al., 2022]. Lastly, incorporating explicit instructions or constraints within prompts has been employed to direct the model’s responses more precisely, ensuring adherence to specific guidelines or objectives.

Bridging these techniques with the focus of our work, we explore the effect of prompt optimization based on intent recognition from user prompts on user satisfaction.

3 Towards a Comprehensive and Fine-grained Taxonomy of Intents in User Prompts

Taxonomies have historically played an important role in information retrieval and knowledge representation, especially for the systematic organization of large volumes of information, providing a standardized classification across various systems [Medelyan et al., 2013]. A clear hierarchical structure can simplify search and discovery and facilitate understanding of contextual relationships between concepts. Along similar lines, in this section, we introduce a comprehensive and fine-grained taxonomy of user intents to address current limitations of user interactions with LLMs, specifically the user dissatisfaction [Kim et al., 2023] when LLMs fail to accurately interpret user intentions.

3.1 Important Characteristics for an Effective Intent Taxonomy

An effective taxonomy for guiding intent-based interactions with LLMs must have certain key features [Shah et al., 2023]. First, it should provide a **comprehensive intent coverage** that encompasses a wide range of user intents, from factual queries to personal and creative interactions. Such comprehensive coverage is critical to accurately identify the intent behind user queries and enable LLMs to provide relevant and targeted responses. Second, the taxonomy should provide a **clear, precise, and consistent** categorization of intents. This clarity and precision are critical to extracting the most appropriate and accurate responses from LLMs. Finally, it is essential to have the taxonomy that is **versatile and applicable across various applications and use-cases**, ranging from technical and educational contexts to personal and artistic interactions. This versatility is key to expanding the user-centered utility of LLMs in different domains.

3.2 A Fine-Grained Intent Taxonomy

In line with the above characteristics, informed by related work [Shah et al., 2023], and current trends in human-AI interactions [Kumar, 2023, Fishkin, 2023], we have identified eight intent categories, each comprising three subcategories. This structure results in 24 fine-grained intents, as detailed in Table 1.

User interactions with LLMs are multifaceted, reflecting a wide spectrum of needs and purposes. To bring structure into this wide and diverse spectrum of purposes, the authors of a recent work [Shah et al., 2023] have conducted a user study and introduced an initial intent taxonomy specifically for user prompts in LLMs, which comprises five primary categories: Information Retrieval, Problem-Solving, Learning, Content Creation, Leisure, and Other. Building upon this foundation and on recent trends in the categorization of user tasks and chat messages in conversations between humans and LLM-based assistants [Fishkin, 2023, Kumar, 2023], we have derived a more detailed and comprehensive taxonomy, in which the identified intents are in alignment with the desired characteristics presented above. This taxonomy organizes user intents into eight distinct categories: Informational, Problem-Solving, Creative, Educational, Personal Interaction, Technical and Professional, Transactional, and Ethical and Philosophical Intents. Each category includes a range of specific intents, from factual queries and problem-solving assistance to creative ideation, personal advice, technical guidance, and moral inquiries. The key characteristics outlined earlier have guided the choice of detailed and fine-grained intent classes. Hence, our taxonomy is designed not only to improve the precision and relevance of LLM responses, but also to offer deeper insights into the diverse and evolving nature of user needs and interactions with AI in digital settings.

4 User Study for the Analysis of Intent Recognition and Intent-Based Prompt Reformulation

To evaluate the accuracy of intent recognition and the effect of intent-based prompt reformulation on user satisfaction, we designed and conducted a three-phase user study as described in the following subsections. More details can also be found in the Appendix.

4.1 User Study Phase 1: Quality of Intent Recognition by State-of-the-Art LLMs

The first phase of the user study was devoted to assessing the accuracy of intent recognition by comparing two different LLMs, namely GPT-3.5 Turbo (i.e., gpt-3.5-turbo-1106) and GPT-4 Turbo (i.e., gpt-4-1106-preview), in a between-subjects design fashion. For the sake of readability, in the remaining of this paper, we refer to these models as GPT-3.5 and GPT-4, respectively. We selected these models as they currently represent the most widely used LLMs available, both in their free versions and paid subscriptions, across different application domains. This comparison aimed to uncover how advances in model architecture might impact user experience in real-world scenarios. In this part of the user study, the participants were introduced to conversational contexts extracted from a publicly accessible dataset [Lhoest et al., 2021] and were asked to continue these dialogues. The underlying LLM analyzed the user's

Table 1: Taxonomy of Intent Types for User Prompts to MFMs.

Intent Type	Fine-granular Intent
Informational Intent	<ul style="list-style-type: none"> • Factual Queries: Requests for specific facts or data • Explanatory Inquiries: Explanations or clarifications about concepts, events, phenomena • Tutorial Requests: Step-by-step instructions or guidance
Problem-Solving Intent	<ul style="list-style-type: none"> • Troubleshooting Assistance: Diagnose and resolve issues or problems • Decision Support: Assistance in decision-making through insights, comparisons, evaluations • Planning and Organization: Aid in planning events, organizing tasks, or managing projects
Creative Intent	<ul style="list-style-type: none"> • Idea Generation: Inspiration or ideas for creative projects • Content Creation: Help in writing, visually representing, or designing original content • Artistic Exploration: Exploration of artistic styles, techniques, historical art contexts
Educational Intent	<ul style="list-style-type: none"> • Learning Support: Assistance with understanding educational material or concepts • Skill Development: Guidance on developing specific skills or competencies • Curricular Planning: Help in designing or choosing educational curricula or courses
Personal Interaction Intent	<ul style="list-style-type: none"> • Conversational Engagement: Dialogue for entertainment, companionship, interaction • Personal Advice: Advice on personal matters or life decisions • Reflection and Insight: Help in self-reflection, personal growth, or to gain insights into certain behavior or thoughts
Technical and Professional Intent	<ul style="list-style-type: none"> • Technical Guidance: Assistance with technical tasks, coding, or problem-solving in a professional context • Business and Career Advice: Guidance on business, career choices, or professional development • Industry-Specific Inquiries: Requests for information or assistance specific to certain industries or professional fields
Transactional Intent	<ul style="list-style-type: none"> • Service Utilization: Requests to use specific functionalities of the model (e.g., language translation, summarization) • Data Processing: Help in processing, analyzing, or visualizing data • Task Automation: Inquiries about automating tasks or workflows
Ethical and Philosophical Intent	<ul style="list-style-type: none"> • Moral and Ethical Queries: Questions about ethical dilemmas, moral principles, or philosophical theories • Societal and Cultural Inquiry: Exploring societal, cultural, or historical topics • Existential Questions: Delving into existential themes or abstract philosophical questions

prompt to recognize their intent based on the taxonomy introduced in Table 1. The participants were then asked whether they agreed with the intent category identified by the LLM. In cases where participants disagreed with the detected intent, they were asked to select a more suitable category from our intent taxonomy, as detailed in Table 1. This process was repeated ten times for various conversational contexts. An exemplary conversation from the user study is provided in the Appendix 7.1.

4.2 User Study Phase 2: Effect of Intent-based Prompt Reformulation on User Satisfaction

The primary objective of the second phase of our user study was to assess how well the LLMs' responses, both to the original and reformulated user prompts, align with user preferences. Following the first phase, which assesses the

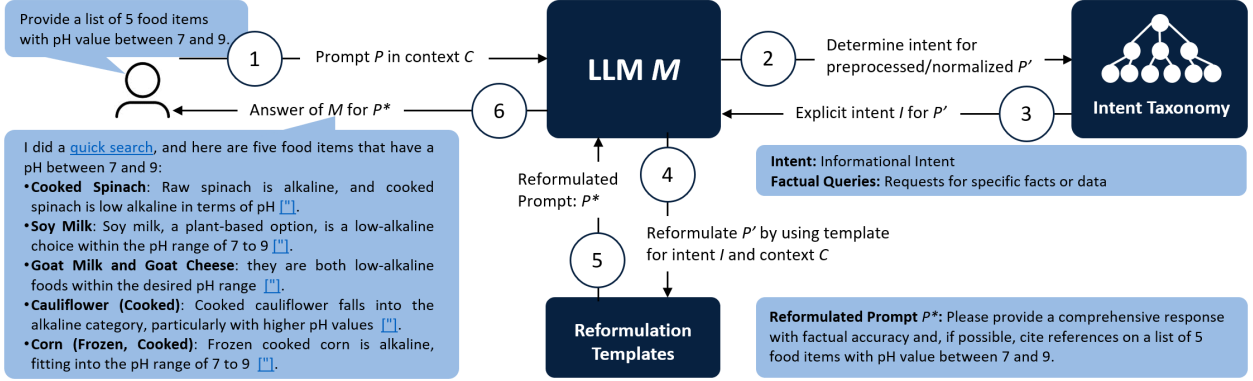


Figure 1: Overview of the proposed prompt reformulation framework, visual summary of Algorithm 1.

intent recognition process, in the second phase of the study, each participant was presented with two distinct responses to their prompt. The underlying LLM generated one of the responses directly from the original user prompt, while the other response was generated based on the reformulated user prompt aligning with the correctly identified intent category, as presented in the following Table 2. To maintain objectivity in participant responses, the participants were blind to the answer options, and the presentation order of responses was randomized. Furthermore, to mitigate bias associated with the length of answers, both responses were kept to approximately the same length. Participants were then asked to choose the more suitable or preferable response. This phase of the user study not only facilitated in-depth analysis of the ability of LLMs to accurately recognize the user’s intent but also provided insights into user preferences regarding AI-generated responses, both in their original and reformulated states. A detailed description of the prompt reformulation based on intent classification is provided in the following subsections.

4.2.1 Prompt Reformulation through Intent Classification

Advanced generative AI systems are trained on vast text corpora and can reason and generate responses across various topics and scenarios. Hence, the main challenge is not including more knowledge into such systems, but rather precisely extracting the most pertinent and accurate responses to user prompts. Recent research on prompt reformulation often argues that conventional direct querying methods often fall short of leveraging the full potential of LLMs, leading to responses that, while correct, may not fully align with the user’s intent or context. Thus, the intent-based reformulation strategy for our user study aims to investigate the link between the user’s intent, the conversational context, and the available LLM knowledge. Figure 1 and Algorithm 1 present our prompt reformulation strategy.

As depicted in Algorithm 1, the framework begins by allowing the LLM (M) to preprocess the user’s prompt (P) to ensure linguistic correctness, followed by the intent classification of the refined prompt (P') using M . This process involves asking M to categorize P' based on predefined intent types from Table 1. The next step is to retrieve a concise summary of the conversational context (C) from M . Subsequently, a suitable template (T) for the identified intent is selected from a set of predefined templates from Table 2. This template is then tailored to incorporate the conversational context C , forming a new, context-enhanced template (T'). The final step involves reformulating the user prompt using T' , resulting in a reformulated prompt that is then used by the LLM M to generate a response. This approach ensures that the user’s original intent is captured more accurately in a realistic, contextually relevant manner.

Algorithm 1: User Prompt Reformulation Framework (LLM M , Prompt P)

Procedure *ReformulateUserPrompt* (P)

```

 $P' \leftarrow \text{PreprocessPrompt}(M, P)$ ;           // Try to express  $P$  in correct English, in case there are
  linguistic issues
 $I \leftarrow \text{ClassifyIntent}(M, P')$ ;       // Ask  $M$  to classify  $P'$  according to one of the intent types in
  Table 1
 $C \leftarrow \text{AnalyzeGetContext}(M)$ ;       // Get from  $M$  a succinct summary of the conversational context  $C$ 
 $T \leftarrow \text{SelectTemplate}(I)$ ;           // Returns template  $T$  from Table 2 for  $I$ 
 $T' \leftarrow \text{IncorporateContext}(C, T)$ ; // I.e., a new template of the form ["In the context" &  $C$  &  $T$ ] is
  created
 $\text{reformulatedPrompt} \leftarrow \text{ReformulatePrompt}(T', P')$ 
return  $\text{GenerateResponse}(M, \text{reformulatedPrompt})$ 

```

4.2.2 The Prompt Reformulation Templates

For the proposed templates, we aim to generate prompt prefixes satisfying the following desiderata:

- **Clear intent representation:** The template should be designed to interpret and respond to the expressed intent, not just the literal prompt.
- **Clear and precise goal description:** The template should clearly and precisely specify the expected response type, level, and detail. The more specific the template, the more targeted the model’s response will be.
- **Flexibility:** The template should be adaptable to a range of prompts of the same intent without requiring substantial modifications for different types of information.
- **Reference indication for correctness support:** If needed, especially for informational queries, the template should include a request for references to ensure the credibility of the response ¹.

Specifically, for each intent type and each corresponding request type from the proposed taxonomy (Table 1), we asked ChatGPT (GPT-4) to provide the most appropriate template according to the above desiderata that could be used as a prompt prefix for the specific query type. In addition, as a post-hoc test to assure the quality of the generated templates, we asked ChatGPT for each template whether it adequately reflects the above desiderata and to improve the template accordingly in case one or more desiderata can be better reflected. Interestingly, an improvement was suggested for some templates. For example, for the intent of ‘learning support’, ChatGPT suggested the following adjustment:

<p>Original: [Provide comprehensive educational support and resources for a deeper understanding of]</p> <p>Assessment: Clear and flexible, but could specify the nature of educational support.</p> <p>Improvement: [Offer educational support through explanations, examples, or resources for a comprehensive understanding of]</p>
--

For other intents, no adjustment was needed. For instance, for ‘skill development’, the answer of ChatGPT was:

<p>Original: [Offer detailed guidance and practical tips for skill enhancement in]</p> <p>Assessment: Specific and adaptable to different skills.</p> <p>Improvement: [None required.]</p>
--

The final suggestions were considered the most appropriate templates to express the related intents. Table 2 gives an overview of all templates generated in this way and employed in the second phase of the user study.

4.3 User Study Phase 3: User’s Understanding of Prompt Reformulation

In the last phase of the user study, we collected demographic data of the participants and focused on their understanding of the prompt reformulation concept. In addition, we provided three prompt reformulation examples and asked them to reformulate the following text “Hey, tell me about Albert Einstein. I need this info ASAP.” To understand participants’ comprehension of the concept of prompt reformulation based on a content-creation intent, we inquired about their likelihood of utilizing the provided prompt reformulation templates based on a 5-point Likert scale, where 1 signifies ‘extremely unlikely’ and 5 denotes ‘extremely likely’. In this stage of the user study, our objective was to assess the users’ willingness to employ the templates after presenting them with potential templates for the given prompt on the same page. Additionally, participants were given the chance to offer further feedback on the study through a free-text response. A word cloud generated from these feedbacks can be found in the Appendix, Figure 8.

4.4 Participant Recruitment

In our user study utilizing both GPT-3.5 and GPT-4, we engaged two distinct groups of participants through the Prolific platform. This choice was made considering Prolific’s reputation for yielding higher quality data [Peer et al., 2022]. We ensured that participants were 18 years old or older, fluent in English and opted for a gender-balanced sample pool. Once participants were recruited on Prolific, they were forwarded to Qualtrics, where our user study was implemented, and upon the completion of the study, they were redirected to Prolific for compensation. Participants were compensated according to €15/hour rate for participation in our studies. Because of the difference in the necessary time for the

¹This last point is optional because the model’s ability to provide actual references might depend on its training and access to retrieval systems on external data.

Table 2: Reformulation Templates for Detailed Intent Types.

Detailed Prompt Intent	Reformulation Template
Factual Queries	[Please provide a comprehensive response with factual accuracy and, if possible, cite references on]
Explanatory Inquiries	[Elucidate the concept of, ensuring a detailed explanation of key aspects and implications adaptable to various topics in]
Tutorial Requests	[Provide a step-by-step tutorial or instruction on how to effectively]
Troubleshooting Assistance	[Assist in identifying and solving the problem related to, considering possible solutions for]
Decision Support	[Offer an evaluation and comparison of advantages and drawbacks for various options regarding]
Planning and Organization	[Outline a structured plan with key steps and considerations for efficiently organizing]
Idea Generation	[Suggest innovative ideas or creative approaches, adaptable to different contexts for]
Content Creation	[Assist in creating engaging content, such as articles, videos, etc., focused on]
Artistic Exploration	[Explore and discuss various artistic approaches and styles suitable for]
Learning Support	[Offer educational support through explanations, examples, or resources for a comprehensive understanding of]
Skill Development	[Offer detailed guidance and practical tips for skill enhancement in]
Curricular Planning	[Assist in selecting or developing a curriculum, focusing on subjects, levels, and educational goals for]
Conversational Engagement	[Engage in an interactive and thoughtful conversation about]
Personal Advice	[Provide thoughtful and considerate personal advice regarding]
Reflection and Insight	[Encourage self-reflection and offer insights, adaptable to various personal or professional topics on]
Technical Guidance	[Offer in-depth technical guidance and support for issues related to]
Business and Career Advice	[Provide strategic guidance and advice, adaptable to various business and career paths on]
Industry-Specific Inquiries	[Present detailed and industry-specific insights and information about]
Service Utilization	[Instruct me in a detailed way and step-by-step on how to use]
Data Processing	[Assist me in processing and analyzing]
Task Automation	[Provide specific guidance on automating]
Moral and Ethical Queries	[Engage in a thoughtful examination and discussion of the moral and ethical aspects of]
Societal and Cultural Inquiry	[Investigate and discuss the societal and cultural dimensions of]
Existential Questions	[Delve into and discuss philosophical perspectives and viewpoints on]

API calls to the two language models, the time frames to complete the study ranged between 40 and 50 minutes, with GPT-3.5 as the underlying LLM and GPT-4, respectively. The data was collected anonymously, and each participant provided digital informed consent before starting the study. Participants had the option to withdraw their consent or leave the study at any point, without the need for any further explanation. In total, we recruited $n_{GPT-3.5} = 124$ ($M_{age} = 30.9, SD = age = 11.2$) for the GPT-3.5 and $n_{GPT-4} = 116$ ($M_{age} = 28.3, SD = age = 8.0$) for the GPT-4 study, respectively. We filtered the participants out if there was a reported age inconsistency between Prolific and Qualtrics and a mismatch in gender, which led us to $n_{GPT-3.5} = 120$ and $n_{GPT-4} = 114$ participants, respectively. Upholding ethical standards, we collected the user data anonymously and obtained written consent from each participant to cover the storage, usage, and potential sharing of their conversational data for research purposes.

We also excluded participants who finished the study in an exceptionally short amount of time. Considering the API response time for GPT-4 was nearly double that of GPT-3.5, and following an analysis of outliers based on the total duration of the study, we established time thresholds of 20 minutes for GPT-3.5 and 35 minutes for GPT-4. These thresholds also align with our pilot testing, where we determined that the ideal time range for each main page, corresponding to the models, is approximately between 1 and 2 minutes, leading to an expected time range between 25-45 for the whole study per participant. If the participants were to spend less time within this interval, they might not have spent sufficient time to fully read and comprehend the GPT answers, making their choice less reliable. Three additional users had to be excluded from the further analysis based on their answer quality, namely giving at least four inputs with only the space character or asking the same questions such as ‘Is that true?’ throughout the survey. After these relatively conservative exclusions, there were $n_{GPT-3.5} = 116$ and $n_{GPT-4} = 95$ participants for GPT-3.5 and GPT-4, respectively.

5 Results and Discussion

This section evaluates the intent recognition capabilities of GPT-3.5 and GPT-4, analyzing their accuracy and the impact of prompt reformulation on user satisfaction. We explore both quantitative results and qualitative user feedback, discuss our key findings, and highlight future research directions.

5.1 Quality of Intent Recognition by State-of-the-Art LLMs

As we asked the participants to subjectively evaluate the intent recognition results provided by the GPT models, we first assessed whether participants’ self-reported intent categories align with those provided by the GPT models. This is necessary as the intent recognition process is essential for the subsequent prompt reformulation. To this end, we employ descriptive statistics and accuracy measures and report confusion matrices for the performance of intent recognition of both GPT-3.5 and GPT-4.

Table 3: Category-wise accuracy values for both models. Due to the low frequency of some unusual intent categories in the dataset, the recognition rate of 0.00% is a rather conservative performance estimate for both models. (See Fig. 2 for details.)

	Factual Queries	Explanatory Inquiries	Tutorial Requests	Troubleshooting Assistance	Decision Support	Planning and Organization	Idea Generation	Content Creation	Artistic Exploration	Learning Support	Skill Development	Curricular Planning
GPT-3.5	75.50	66.34	59.26	57.14	97.56	91.67	84.00	36.67	0.00	36.84	66.67	50.00
GPT-4	96.02	89.50	98.39	80.95	98.63	77.78	85.71	90.48	100	0.00	55.56	0.00
	Conversational Engagement	Personal Advice	Reflection and Insight	Technical Guidance	Business and Career Advice	Industry-Specific Inquiries	Service Utilization	Data Processing	Task Automation	Moral and Ethical Queries	Societal and Cultural Inquiry	Existential Questions
GPT-3.5	71.43	93.68	82.61	80.00	0.00	52.38	80.00	66.67	90.16	90.43	46.15	0.00
GPT-4	92.86	88.24	100	70.00	66.67	25.00	90.91	0.00	100	95.92	81.25	77.27

More specifically, based on each conversational context and the subsequent user query, we asked the respective LLM model to categorize the user intent into the predefined 24 fine-grained intent categories. With this approach, we achieved accuracy values of 75.28% and 89.64% using GPT-3.5 and GPT-4, respectively. In addition, F_1 scores of 74.28% and 88.84% were obtained using GPT-3.5 and GPT-4, respectively. Relevant confusion matrices are depicted in Fig. 2. According to these results, while GPT-3.5 struggled with the classification of factual queries and explanatory inquiries, GPT-4 achieved superior performance by 20.5% and 23.17% in the accuracy corresponding to the above-mentioned categories, respectively. The increase of the model performance in the content creation category is also significant, as GPT-4 reached an increased accuracy of 53.81%. Both models struggled however with the recognition of the ‘learning support’ intent. Considering all categories, GPT-4 achieved better accuracy in 16 out of 24 intent categories (Table 3).

5.2 Effect of Intent-based Prompt Reformulation on User Satisfaction

In the second phase of the user study, we analyzed user preferences for responses generated by GPT-3.5 and GPT-4 with and without leveraging prompt reformulation aligned with the previously determined intent category. The collected dataset is available on [GitHub](#).

For each of the 24 detailed intent categories, we conducted separate analyses to determine the number of participants who favored either the original GPT responses or the GPT responses to reformulated user prompts for both models independently. To this end, we utilize paired samples t-tests with an alpha level of 0.05 separately for GPT-3.5 and GPT-4. Furthermore, we employed paired samples t-tests, maintaining the same alpha level, to assess if GPT-3.5 and GPT-4 show comparable performance.

Our analysis only included instances where GPT correctly identified the user intent in the first phase. Furthermore, to ensure a fair and consistent analysis, we implemented a preliminary filtering step on the collected samples. This process

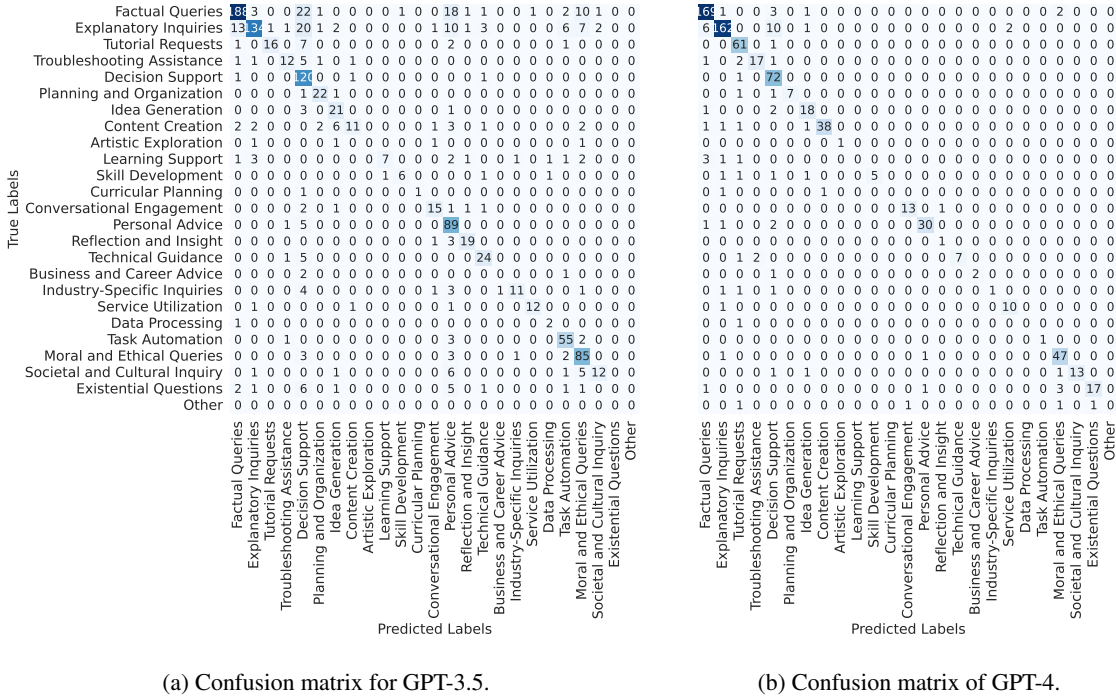


Figure 2: Details about the exact intent classification, represented in confusion matrices.

involved comparing the length of GPT’s responses to the original user prompts with those to the reformulated prompts. More specifically, We only included responses in our evaluation where the length difference between the two sets of answers was within a 10% margin in terms of character count. This approach was adopted to mitigate any potential bias that might arise from variations in answer length.

By analyzing the answer preferences of the users, we compared the number of samples category-wise (Fig. 3). Interestingly, participants showed a preference for the answers to the reformulated prompts generated by GPT-4 over those by GPT-3.5. However, as indicated in Fig. 3, in both GPT-3.5 and GPT-4 cases, users generally favored the original model answers, with a preference rate of 56.61% for GPT-3.5 and 53.5% for GPT-4, respectively. For certain categories, users showed a preference for the answer to the reformulated prompts. Specifically, in the context of GPT-3.5, this trend was observed in categories like ‘tutorial requests’ and ‘learning support’. Meanwhile, with GPT-4, users favored the model answers to their reformulated prompts in areas such as ‘troubleshooting assistance’, ‘idea generation’, ‘skill development’, as well as ‘moral and ethical queries’, and ‘societal and cultural inquiries’. This observation underscores that the higher the advancement level of the GPT model, the better the quality of the answers to the reformulated prompts. Notably, in the ‘factual queries’ category, our template is designed to prompt the model to include references in its responses. Upon examining the quantity of references present in the responses to the reformulated prompts, GPT-4 demonstrated superior performance compared to GPT-3.5. However, it is noteworthy that only 25.53% and 40.24% of the responses actually incorporated references, either in citation or link format. This discrepancy in reference inclusion has significant implications for user-model interactions. It suggests that while advanced models like GPT-4 show promise in enriching responses with references, there is a clear need for further refinement to consistently meet user expectations for detailed and sourced information. In evaluating user preferences for the inclusion of references, we found that participants preferred the responses to reformulated queries produced by GPT-3.5 in 58.33% of the cases. In contrast, users favored the GPT-4 answers to reformulated prompts in only 35.29% of instances.

For a more comprehensive statistical analysis of the collected preference data, we further implemented paired t-tests. This approach was appropriate given that the aggregated preferences (and the differences in preference) conforms to a normal distribution and meets other essential assumptions. According to the p-values obtained from this analysis, there is no substantial mean difference between the responses to original prompts and those to reformulated prompts for GPT-4, as well as in the comparison between the GPT-3.5 and GPT-4 datasets. However, a significant difference emerges when comparing the original responses to those generated in response to reformulated prompts for GPT-3.5.

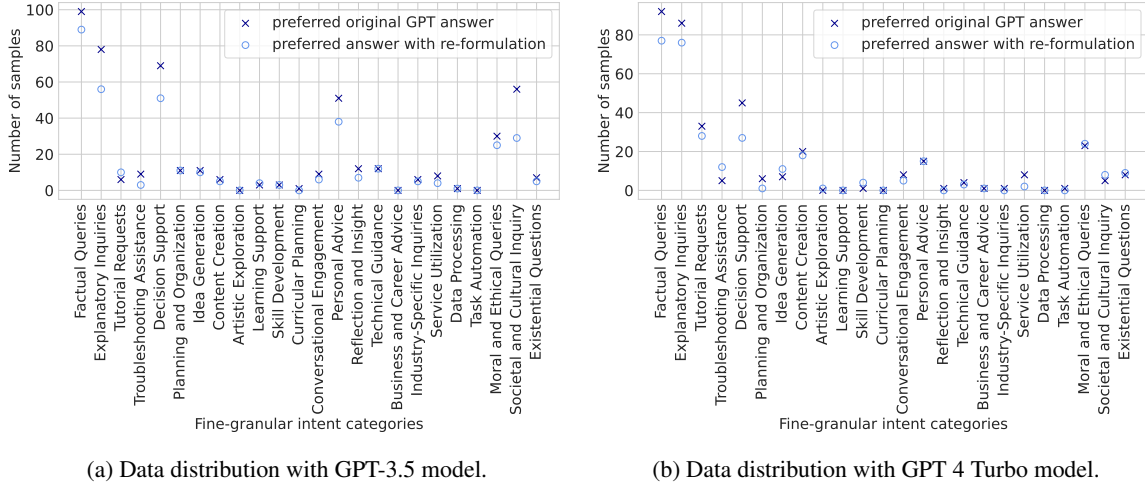


Figure 3: Data distribution corresponding to the fine-granular intent categories.

Table 4: Since the data follows a normal distribution, paired t-tests were utilized.

	t -statistic	p -value
GPT-3.5 Turbo: original and re-prompted	2.9821	0.0067
GPT-4 Turbo: original and re-prompted	1.7250	0.0979
GPT-3.5 Turbo - GPT-4 Turbo	1.2506	0.2173

5.3 User’s Understanding of Prompt Reformulation

After analyzing the two main parts of the study, we measured user satisfaction by asking the following two questions:

1. In the Comparison phase, the new answers were produced by the same large language model, but a prompt reformulation was used. Would you try reformulations if the templates were available to you?
2. Here are some examples on how to use the reformulation templates:
 - Factual Queries: Please provide a comprehensive response with factual accuracy and, if possible, cite references on *-your preferred topic-*
 - Content Creation: Assist in creating engaging content, such as articles, videos, etc., focused on *-your preferred topic-*
 - Existential Questions: Delve into and discuss philosophical perspectives and viewpoints on *-your preferred topic-*

Would you try out these reformulations to make your conversation with a large language model more effective?

The average scores were 4.0 (Likely) for GPT-3.5 and 3.8 (Likely) for GPT-4, respectively, suggesting that users are open to learning how to formulate their prompts more effectively. Interestingly, when the users were asked to apply reformulation to a given prompt, 46.55% and 54.74% of users corresponding to GPT-3.5 and GPT-4, respectively, used the keywords from the correct template from the previously provided three different templates in their reformulated prompts.

5.4 Discussion and Limitations

The results of our user study highlight several key insights into the evolving capabilities and user interactions with state-of-the-art large language models, particularly GPT-3.5 and GPT-4.

- **Improved Accuracy in GPT-4:** The considerable improvement in intent recognition accuracy from GPT-3.5 to GPT-4 indicates a substantial advancement in the model’s ability to understand and categorize user intent. This is particularly notable in categories such as ‘factual queries’, ‘explanatory inquiries’, and ‘content creation’, where GPT-4 demonstrates a significant lead over its predecessor.

- **Consistent Challenges in Certain Categories:** Despite these advancements, both models exhibit consistent difficulties in certain intent categories that occur less frequently, such as ‘curricular planning’ or ‘learning support’. For these kinds of categories, GPT-3.5 often outperforms GPT-4. This suggests inherent challenges in these areas that future developments of LLMs might focus on.
- **Unbalanced Data Distribution:** Our study indicates that ChatGPT users are typically driven by informational intents rather than creative, exploratory, or planning intents such as ‘artistic exploration’ or ‘curricular planning’, regardless of the nature of the preceding chat history.
- **Varied User Preferences and Subjectivity:** The user preference for answers to reformulated prompts of GPT-4 over GPT-3.5, and the overall preference for model answers to original prompts, also highlights the subjectivity in user satisfaction. It suggests that while accuracy and intent recognition are important, they are not the sole determinants of user satisfaction, highlighting thus the necessity for further research in this area.
- **The Curse of Persuasive LLM Answers:** Users interacting with GPT models on OpenAI’s platform are cautioned: “ChatGPT can make mistakes. Consider checking important information”. Accordingly, for factual queries, our template encouraged models to include references to aid user fact-checking. Our findings reveal that participants favored answers to their original prompts even if they did not contain references when using the more advanced GPT-4 model, while for the earlier GPT-3.5 model, the answers to reformulated prompts were preferred. This suggests that advanced models like GPT-4 can satisfactorily and convincingly answer user prompts, even without references.

Implications for LLM Development: Our results demonstrate the rapid pace of improvement in LLMs, suggesting a trajectory towards even more specialized and accurate models. The discrepancies in performance across different intent categories could guide future developments. In addition, the mixed user preferences regarding answers to their original prompts versus those resulting from intent-based reformulations highlight the importance of educating users on effective prompt formulation and a more objective analysis of the provided answers. In this context, our study indicates a user willingness to learn and adapt to this technology.

Limitations and Future Research Directions: The observation that users often favored original answers raises questions about the ability of users to discern the accuracy of LLM outputs. This is particularly crucial given the risk of misinformation, as noted in cases where GPT models could provide incorrect information. Furthermore, the uneven data distribution among different intent categories and the noted biases (e.g., in cases of explanatory approaches in GPT-4) suggest a need for more balanced and controlled datasets in future studies. In addition, the consistent difficulty in certain intent categories like ‘learning support’ points to the need for a deeper examination and possible refinement of these categories, potentially splitting them into more nuanced sub-categories.

6 Conclusion

In this work, we analyzed the capabilities of GPT-3.5 Turbo and GPT-4 Turbo to recognize user intents and the effect of intent-based prompt reformulation on user satisfaction. Our study revealed that while ChatGPT models are considerably improving at recognizing user intents, they still struggle to recognize unusual intents. Interestingly, for some unusual intent categories, GPT-3.5 outperforms GPT-4 in terms of recognition rates. Furthermore, the study showed that for both models, whenever the intent is correctly recognized, the users still tend to prefer the answers to their original prompts to those from the intent-based reformulations. This finding is surprising and in contradiction to current literature on prompt reformulation. Despite their potential benefits for unusual intent categories, strategies akin to chain-of-thought-like prompt reformulation seem to become less impactful with model improvements, as indicated by the results of this study.

References

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun

- Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Pokorny Michael, Michelle Pokrass, Vitchyr Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, March 2023. doi:[10.48550/ARXIV.2303.08774](https://doi.org/10.48550/ARXIV.2303.08774).
- OpenAI. Gpt-4v(ision) system card. Technical report, OpenAI, 9 2023.
- Google DeepMind. *Welcome to the Gemini era*, 2023.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.
- Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. Mathematical discoveries from program search with large language models. *Nature*, pages 1–3, 2023. doi:<https://doi.org/10.1038/s41586-023-06924-6>.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023. doi:<https://doi.org/10.48550/arXiv.2309.00267>.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei

- Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways, 2022.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepey, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. Palm 2 technical report, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023b.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, June 2023. doi:[10.48550/ARXIV.2306.05685](https://doi.org/10.48550/ARXIV.2306.05685).
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Bharath Chintagunta, Namit Katariya, Xavier Amatriain, and Anitha Kannan. Medically aware gpt-3 as a data generator for medical dialogue summarization. In *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*. Association for Computational Linguistics, 2021. doi:[10.18653/v1/2021.nlpmc-1.9](https://doi.org/10.18653/v1/2021.nlpmc-1.9).
- Seppo Enarvi, Marilisa Amoia, Miguel Del-Agua Teba, Brian Delaney, Frank Diehl, Stefan Hahn, Kristina Harris, Liam McGrath, Yue Pan, Joel Pinto, Luca Rubini, Miguel Ruiz, Gagandeep Singh, Fabian Stemmer, Weiyi Sun, Paul Vozila, Thomas Lin, and Ranjani Ramamurthy. Generating medical reports from patient-doctor conversations using sequence-to-sequence models. In *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*. Association for Computational Linguistics, 2020. doi:[10.18653/v1/2020.nlpmc-1.4](https://doi.org/10.18653/v1/2020.nlpmc-1.4).
- Michael Dowling and Brian Lucey. Chatgpt for (finance) research: The bananarama conjecture. *Finance Research Letters*, 53:103662, May 2023. ISSN 1544-6123. doi:[10.1016/j.frl.2023.103662](https://doi.org/10.1016/j.frl.2023.103662).
- John V. Pavlik. Collaborating with chatgpt: Considering the implications of generative artificial intelligence for journalism and media education. *Journalism & Mass Communication Educator*, 78(1):84–93, January 2023. ISSN 2161-4326. doi:[10.1177/10776958221149577](https://doi.org/10.1177/10776958221149577).
- Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. Wordcraft: Story writing with large language models. In *27th International Conference on Intelligent User Interfaces, IUI '22*. ACM, March 2022. doi:[10.1145/3490099.3511105](https://doi.org/10.1145/3490099.3511105).

- Kathrin Seßler, Tao Xiang, Lukas Bogenrieder, and Enkelejda Kasneci. *PEER: Empowering Writing with Large Language Models*, pages 755–761. Springer Nature Switzerland, 2023. ISBN 9783031426827. doi:[10.1007/978-3-031-42682-7_73](https://doi.org/10.1007/978-3-031-42682-7_73).
- Shima Imani, Liang Du, and Harsh Shrivastava. Mathprompter: Mathematical reasoning using large language models, 2023.
- Arne Bewersdorff, Kathrin Seßler, Armin Baur, Enkelejda Kasneci, and Claudia Nerdel. Assessing student errors in experimentation using artificial intelligence and large language models: A comparative study with human raters. *Computers and Education: Artificial Intelligence*, 5:100177, 2023.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, March 2023. ISSN 1557-7341. doi:[10.1145/3571730](https://doi.org/10.1145/3571730).
- Razvan Azamfirei, Sapna R Kudchadkar, and James Fackler. Large language models and the perils of their hallucinations. *Critical Care*, 27(1):1–2, 2023. doi:<https://doi.org/10.1186/s13054-023-04393-x>.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*, 2023. doi:<https://doi.org/10.48550/arXiv.2309.01219>.
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*, 2023. doi:<https://doi.org/10.48550/arXiv.2303.08896>.
- Yongchao Wu, Aron Henriksson, Martin Duneld, and Jalal Nouri. *Towards Improving the Reliability and Transparency of ChatGPT for Educational Question Answering*, pages 475–488. Springer Nature Switzerland, 2023a. ISBN 9783031426827. doi:[10.1007/978-3-031-42682-7_32](https://doi.org/10.1007/978-3-031-42682-7_32).
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning, 2022.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, April 2023. doi:[10.48550/ARXIV.2304.10592](https://doi.org/10.48550/ARXIV.2304.10592).
- Rongjie Huang, Mingze Li, Dongchao Yang, Jiatong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, Yi Ren, Zhou Zhao, and Shinji Watanabe. Audiogpt: Understanding and generating speech, music, sound, and talking head. *arXiv preprint arXiv:2304.12995*, April 2023. doi:[10.48550/ARXIV.2304.12995](https://doi.org/10.48550/ARXIV.2304.12995).
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, April 2023. doi:[10.48550/ARXIV.2304.14178](https://doi.org/10.48550/ARXIV.2304.14178).
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhui Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, May 2023. doi:[10.48550/ARXIV.2305.06355](https://doi.org/10.48550/ARXIV.2305.06355).
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, June 2023. doi:[10.48550/ARXIV.2306.05424](https://doi.org/10.48550/ARXIV.2306.05424).
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. Palm-e: An embodied multimodal language model, 2023.
- Feilong Chen, Minglun Han, Haozhi Zhao, Qingyang Zhang, Jing Shi, Shuang Xu, and Bo Xu. X-llm: Bootstrapping advanced large language models by treating multi-modalities as foreign languages. *arXiv preprint arXiv:2305.04160*, May 2023. doi:[10.48550/ARXIV.2305.04160](https://doi.org/10.48550/ARXIV.2305.04160).
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, May 2023. doi:[10.48550/ARXIV.2305.16355](https://doi.org/10.48550/ARXIV.2305.16355).

- Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, March 2023b. doi:[10.48550/ARXIV.2303.04671](https://doi.org/10.48550/ARXIV.2303.04671).
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. *arXiv preprint arXiv:2303.17580*, March 2023. doi:[10.48550/ARXIV.2303.17580](https://doi.org/10.48550/ARXIV.2303.17580).
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*, September 2023c. doi:[10.48550/ARXIV.2309.05519](https://doi.org/10.48550/ARXIV.2309.05519).
- Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023. doi:<https://doi.org/10.48550/arXiv.2310.07704>.
- Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. doi:[10.1109/TNNLS.2022.3229161](https://doi.org/10.1109/TNNLS.2022.3229161).
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. URL https://openaccess.thecvf.com/content/CVPR2022/html/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.html.
- Felix Wong, Erica J Zheng, Jacqueline A Valeri, Nina M Donghia, Melis N Anahtar, Satotaka Omori, Alicia Li, Andres Cubillos-Ruiz, Aarti Krishnan, Wengong Jin, et al. Discovery of a structural class of antibiotics with explainable deep learning. *Nature*, pages 1–9, 2023. doi:<https://doi.org/10.1038/s41586-023-06887-8>.
- Yoonsu Kim, Jueon Lee, Seoyoung Kim, Jaehyuk Park, and Juho Kim. Understanding users’ dissatisfaction with chatgpt responses: Types, resolving tactics, and the effect of knowledge level. *arXiv preprint arXiv:2311.07434*, 2023. doi:<https://doi.org/10.48550/arXiv.2311.07434>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*, 2023. doi:<https://doi.org/10.48550/arXiv.2305.10601>.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michal Podstawski, Hubert Niewiadomski, Piotr Nyczyk, et al. Graph of thoughts: Solving elaborate problems with large language models. *arXiv preprint arXiv:2308.09687*, 2023. doi:<https://doi.org/10.48550/arXiv.2308.09687>.
- Hai Dang, Lukas Mecke, Florian Lehmann, Sven Goller, and Daniel Buschek. How to prompt? opportunities and challenges of zero-and few-shot learning for human-ai interaction in creative applications of generative models. *arXiv preprint arXiv:2209.01390*, 2022. doi:<https://doi.org/10.48550/arXiv.2209.01390>.
- Chirag Shah, Ryen W White, Reid Andersen, Georg Buscher, Scott Counts, Sarkar Snigdha Sarathi Das, Ali Montazer, Sathish Manivannan, Jennifer Neville, Xiaochuan Ni, et al. Using large language models to generate, validate, and apply user intent taxonomies. *arXiv preprint arXiv:2309.13063*, 2023. doi:<https://doi.org/10.48550/arXiv.2309.13063>.
- Hiteshwar Kumar Azad and Akshay Deepak. Query expansion techniques for information retrieval: a survey. *Information Processing & Management*, 56(5):1698–1735, 2019. doi:<https://doi.org/10.1016/j.ipm.2019.05.009>.
- Andrei Broder. A taxonomy of web search. In *ACM Sigir forum*, volume 36, pages 3–10. ACM New York, NY, USA, 2002. doi:<https://doi.org/10.1145/792550.792552>.
- Bernard J Jansen, Danielle L Booth, and Amanda Spink. Determining the user intent of web search engine queries. In *Proceedings of the 16th international conference on World Wide Web*, pages 1149–1150, 2007. doi:<https://doi.org/10.1145/1242572.1242739>.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35, 2023. doi:<https://doi.org/10.1145/3560815>.
- Bonan Min, Hayley Ross, Elior Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys*, 56(2):1–40, 2023. doi:<https://doi.org/10.1145/3605943>.

- Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Haotian Luo, Jiayou Zhang, Nebojsa Jojic, Eric P Xing, and Zhiting Hu. Promptagent: Strategic planning with language models enables expert-level prompt optimization. *arXiv preprint arXiv:2310.16427*, 2023. doi:<https://doi.org/10.48550/arXiv.2310.16427>.
- Olena Medelyan, Ian H Witten, Anna Divoli, and Jeen Broekstra. Automatic construction of lexicons, taxonomies, ontologies, and other knowledge structures. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 3(4):257–279, 2013. doi:<https://doi.org/10.1002/widm.1097>.
- Varun Kumar. What are people asking to chatgpt? an insight into 10 thousand chat messages between humans and ai assistants. <https://varunon9.medium.com/what-are-people-asking-to-chatgpt-f5f324a6cc27>, 2023. Retrieved 2023-12-29.
- Rand Fishkin. We analyzed millions of chatgpt user sessions: Visits are down 29% since may, programming assistance is 30% of use. <https://sparktoro.com/blog/we-analyzed-millions-of-chatgpt-user-sessions-visits-are-down-29-since-may-programming-assistance-is-30-of-use/>, 2023. Retrieved 2023-12-29.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. Datasets: A community library for natural language processing. *arXiv preprint arXiv:2109.02846*, 2021. doi:<https://doi.org/10.48550/arXiv.2109.02846>.
- Eyal Peer, David Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer. Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, 54(4):1643–1662, 2022. doi:[10.3758/s13428-021-01694-3](https://doi.org/10.3758/s13428-021-01694-3).

7 Appendix

7.1 Study design

The users have time to read the randomly selected chat history out of 240 previously selected conversations. Each chat history was previously classified, and 10 samples were saved corresponding to each fine-granular intent categories, to provide diverse answers across the categories. The ‘Next page’ button was only enabled once the GPT answer and the class prediction in the backend was finished (Fig 4).

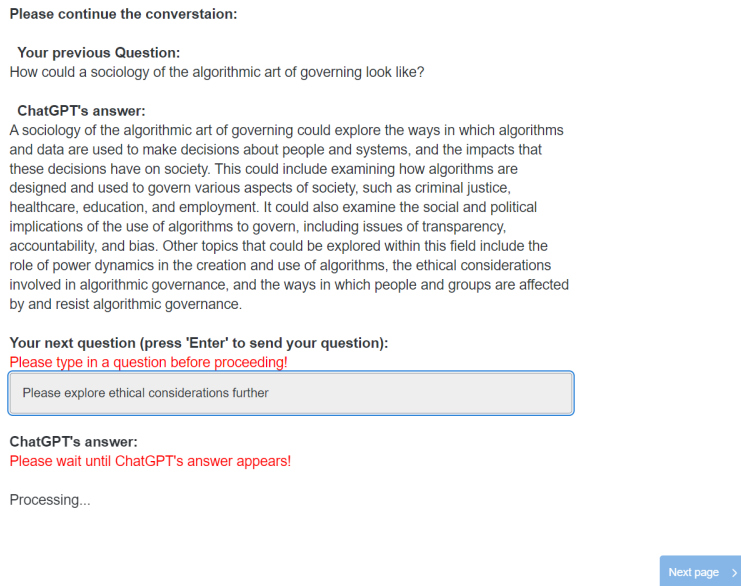


Figure 4: Example of the first part of the study

In the next page, the participants were asked whether they agree with the predicted intent (Fig 5). Their previous prompt and a summary table about the fine-grained intent categories were provided to them. If they chose the radio button ‘Yes, I agree.’ they were directed to the next page with a new chat history. In case they did not agree with the predicted intent category, they were directed to a page, where they could select a more suitable intent with the help of radio buttons. Only one intent category could be selected (Fig 6).

After 10 different chat histories were shown and the intent classification was completed, the participants had a chance to choose between the original and re-prompted GPT answers (Fig 7).

8 User Feedback

We collected feedbacks about the study in a free text format, from where the collected word cloud is visualized in Fig 8. In many cases, the participants used this field to say that they enjoyed our study and are curious about the results. Some also mentioned, that they experienced longer waiting times, which could be caused by poor internet connection.

8.1 Data distribution based on demographics data

We analyzed the collected data samples based on employment status and age, which are demonstrated in Fig 9, 10. We collected data mostly from participants who were working full-time, working part-time, or students at the time of the study. The majority of the participants below the age of 40. We observed that the ratio between the preferred answers came closer, and the category-wise evaluation shows that in specific categories, users preferred the answers generated with re-prompting.

The same phenomenon can be observed, when we measured the understanding based on the free text reformulation input, where the participants were asked to reformulate the sentence ‘Hey, tell me about Albert Einstein. I need info ASAP’. We filtered out the users, whose answer did not contain the main verbs from the previously provided three templates, namely: ‘provide’, ‘assist’, and ‘delve’. The contrast is even more significant for the GPT-4 model. From

*

Do you agree that your intent in this conversation ("Please explore ethical considerations further") **fits into the category:**
22. Moral and Ethical Queries

Yes, I agree.

No, I do not agree.

Intent Type	Fine-granular Intent
Informational Intent	1. Factual Queries: Requests for specific facts or data
	2. Explanatory Inquiries: Explanations or clarifications about concepts, events, phenomena
	3. Tutorial Requests: Step-by-step instructions or guidance
Problem-Solving Intent	4. Troubleshooting Assistance: Diagnose and resolve issues or problems
	5. Decision Support: Assistance in decision-making through insights, comparisons, evaluations
	6. Planning and Organization: Aid in planning events, organizing tasks, or managing projects

Figure 5: Example of intent prediction. For the participants, the complete fine-grained intent categories were provided.

If you did not agree with the category, please give a better fit (choose the correct class):

Informational Intent

1. Factual Queries: Requests for specific facts or data
2. Explanatory Inquiries: Explanations or clarifications about concepts, events, phenomena
3. Tutorial Requests: Step-by-step instructions or guidance

Problem-Solving Intent

4. Troubleshooting Assistance: Diagnose and resolve issues or problems
5. Decision Support: Assistance in decision-making through insights, comparisons, evaluations
6. Planning and Organization: Aid in planning events, organizing tasks, or managing projects

Creative Intent

Figure 6: Example of intent correction, if the participants did not agree with the predicted intent. The complete list of fine-grained intent was provided.

these results, we can say that in some cases, users preferred the re-prompted GPT answers, when they understood the main driving force of the study.

8.2 Evaluation based on user understanding

We analyzed the collected data based on the understanding of the survey of the participants. Firstly, we filtered the participants, who answered with ‘Likely’ or ‘Extremely likely’ to at least one of the questions asking whether they would use prompt reformulation in the future.

*Which answer to your previous question "Please explore ethical considerations further" do you prefer?

Please wait until BOTH answers appear!

ChatGPT's answer:

Answer A

Exploring the ethical considerations of algorithmic governance is paramount, as it involves moral responsibilities and decisions that significantly affect individuals and communities. Ethical considerations include the following aspects: 1. **Fairness and Bias:** Algorithms can perpetuate and amplify existing social biases if they are trained on biased data sets. Ensuring fairness involves scrutinizing the data and the decision-making processes to prevent discrimination against any individual or group. 2. **Transparency:** The "black box" nature of many algorithms can obscure how decisions are made. Ethical governance demands that the decision-making process be transparent enough for stakeholders to understand and trust the outcomes. 3. **Accountability:** When algorithms make incorrect or harmful decisions, it is crucial to have mechanisms in place to hold the responsible parties accountable. This includes developers, operators, and organizations that deploy these algorithms. 4. **Privacy:** Algorithms often require personal data to function. Protecting this data and respecting individuals' privacy is a fundamental ethical obligation. This involves considering the extent and type of data collected and ensuring it is used responsibly. 5. **Autonomy:** Algorithmic decisions can influence or even dictate human behavior. Respecting individual autonomy means being cautious about how algorithms might unduly influence or constrain human choices. 6. **Informed Consent:** When personal data is used, individuals should have a clear understanding of how their information will be employed and should consent to its use freely and with full knowledge of the implications. 7. **Justice:** The impact of algorithmic governance should be equitable, and not disproportionately benefit or harm any particular group. This includes considering the long-term societal impacts and striving for outcomes that promote social justice. 8. **Impact on Society and Culture:** Algorithms can shape cultural norms and social interactions. Ethical considerations should include the ways these technologies might change society and whether these changes align with societal values and goals. In conclusion, the ethical considerations in algorithmic governance are multifaceted and deeply interconnected with the principles of social justice, respect for human rights, and the democratic process. A thoughtful examination of these considerations is necessary to guide the development, deployment, and regulation of algorithms, ensuring that they serve the public good and enhance, rather than undermine, the fabric of society.

Answer B

Ethical considerations in the algorithmic art of governing are multifaceted and revolve around several key issues: 1. **Consent and Privacy:** As algorithms often process personal data to make decisions, it's crucial to consider the extent to which individuals have consented to their data being used in this way. Ethical governance must address the balance between the benefits of data-driven decisions and the right to privacy. 2. **Transparency and Explainability:** Algorithms can be complex and opaque, making it difficult for people to understand how decisions are made. There's an ethical imperative for transparency so that individuals can comprehend the basis of decisions that affect them and challenge them if necessary. 3. **Accountability:** When algorithms make mistakes or cause harm, it's important to have clear lines of accountability. Ethical considerations include determining who is responsible for the outcomes of an algorithmic decision—the designers, the implementers, or the algorithms themselves—and how they can be held to account. 4. **Bias and Discrimination:** Algorithms can perpetuate and amplify social biases if they're trained on biased data sets. Ethical governance requires active measures to prevent discrimination and ensure that algorithmic decisions are fair and equitable to all social groups. 5. **Autonomy and Human Dignity:** Algorithms should be designed in ways that respect human dignity and autonomy, rather than reducing individuals to mere data points. This includes considering the effect of algorithmic decision-making on people's ability to make their own choices and control their own lives. 6. **Social and Economic Impacts:** The societal consequences of algorithmic governance, such as the potential for job displacement or the widening of social inequalities, must be ethically managed. This involves considering how benefits and burdens are distributed and whether the algorithms serve the public good. 7. **Long-term Effects:** The long-term implications of algorithmic governance on social structures and institutions are an ethical concern. Policies and oversight mechanisms must be in place to ensure that the use of algorithms does not erode fundamental social values and norms over time. Addressing these ethical considerations requires a multidisciplinary approach, involving not only sociologists but also ethicists, computer scientists, legal scholars, and policymakers. The goal is to create a framework for algorithmic governance that is just, responsible, and aligned with the broader values of society.

Which answer to your previous question do you prefer?

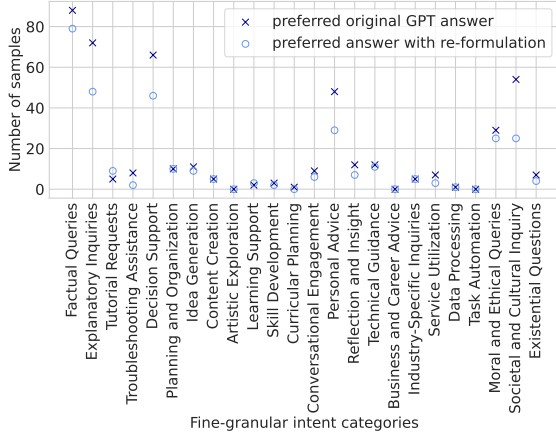
Answer A

Answer B

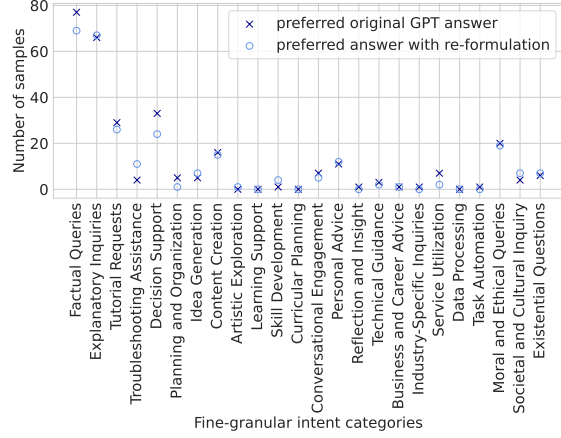
[< Previous Page](#)

[Next page >](#)

Figure 7: Example of the second part of the study

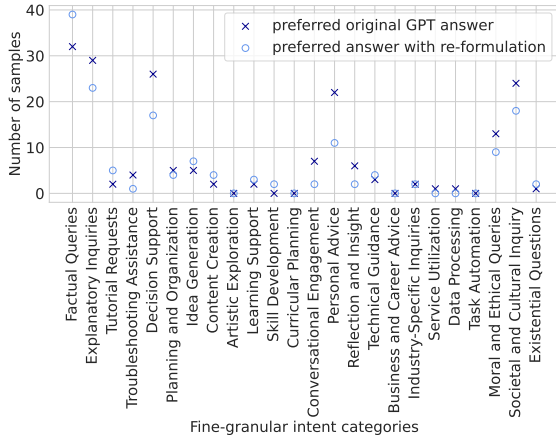


(a) GPT-3.5.

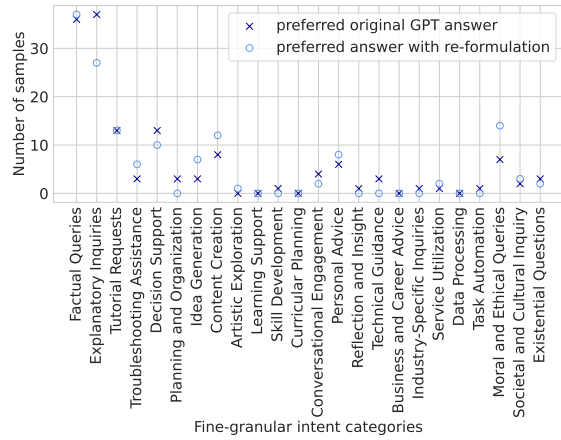


(b) GPT-4.

Figure 11: Results with participants who answered with ‘Likely’ or ‘Extremely likely’ to our questions measuring whether they would use re-formulations.



(a) GPT-3.5.



(b) GPT-4.

Figure 12: Results with participants, who understood templating.