

Beyond Prompts: Learning from Human Communication for Enhanced AI Intent Alignment

Yoonsu Kim

yoonsu16@kaist.ac.kr
School of Computing, KAIST
Daejeon, Republic of Korea

Seoyoung Kim

youthskim@kaist.ac.kr
School of Computing, KAIST
Daejeon, Republic of Korea

Kihoon Son

kihoon.son@kaist.ac.kr
School of Computing, KAIST
Daejeon, Republic of Korea

Juho Kim

juhokim@kaist.ac.kr
School of Computing, KAIST
Daejeon, Republic of Korea

ABSTRACT

AI intent alignment, ensuring that AI produces outcomes as intended by users, is a critical challenge in human-AI interaction. The emergence of generative AI, including LLMs, has intensified the significance of this problem, as interactions increasingly involve users specifying desired results to AI systems. In order to support better AI intent alignment, we aim to explore human strategies for intent specification in human-human communication. By studying and comparing human-human and human-LLM communication, we identify key strategies that can be applied to the design of AI systems that are more effective at understanding and aligning with user intent. This study aims to advance toward a human-centered AI system by bringing together human communication strategies to the design of AI systems.

CCS CONCEPTS

• Human-centered computing → Empirical studies in HCI.

KEYWORDS

Large Language Models, Chat-based LLM, Human-AI Interaction, AI alignment, Intent Alignment, Human-Human Communication

ACM Reference Format:

Yoonsu Kim, Kihoon Son, Seoyoung Kim, and Juho Kim. 2024. Beyond Prompts: Learning from Human Communication for Enhanced AI Intent Alignment. In *Proceedings of CHI Conference on Human Factors in Computing Systems (CHI '24 Getting Back Together Workshop)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Recent advances in generative AI, particularly large language models (LLMs), have led to a paradigm shift in human-computer interaction. Instead of explicitly instructing machines how to perform tasks, users can now specify their desired outcomes in natural language, empowering a new paradigm of intent-based outcome

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI '24 Getting Back Together Workshop, May 12, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

specification [9]. By specifying desired outcomes in natural language, users can now leverage the capabilities of the AI to perform a wide range of tasks from creative content generation to problem-solving [1, 14].

While this shift has opened up exciting possibilities for various applications of AI, it has also introduced challenges in ensuring that AI systems reliably and correctly understand and align with user's intent [3, 13, 17]. Misalignment with the user's intent significantly impairs the user's experience when interacting with AI, making it difficult for them to achieve their goals. Specifically, a recent study based on user data found that the users most frequently experience dissatisfaction in terms of intent understanding when interacting with LLMs [8]. Also, prior work of AI alignment reported that users have difficulties in knowing how to effectively "prompt" LLM and other generative AI systems to produce the desired outcomes [2, 7, 19].

Users often struggle to deliver their intent clearly when interacting with generative models, whereas this issue is less pronounced in human-human interactions, like in conversation. Therefore, Humans have evolved sophisticated strategies for achieving intent alignment during conversation, such as active listening, clarification questions, and feedback [10, 18]. We believe that these strategies can be adapted to the context of human-AI interaction. By studying these strategies through a comparative study between human-human interaction and human-AI interaction, we aim to identify key principles adaptable to AI system design, enhancing intent alignment in human-AI interaction. Specifically, we initially focus on verbal communication, such as text and voice, as this is the most common form of communication in human-AI interaction, and nonverbal cues can be difficult to interpret and may introduce additional complexity.

This paper presents a study that investigates human-human communication and extracts insights that can be applied to LLM-based interactions. We focus on the following research questions:

- **RQ1.** What are the key verbal strategies that humans use to achieve intent alignment during conversation?
- **RQ2.** What are the benefits and limitations of these strategies?
- **RQ3.** How can these strategies be adapted and implemented in LLM-based user interfaces?

We believe that this study will provide valuable insights into how to improve intent alignment in human-AI interaction, and will

ultimately lead to the development of more effective and user-friendly AI systems.

2 COMPARATIVE STUDY: HUMAN-HUMAN INTERACTION VS HUMAN-AI INTERACTION

Study Design In order to investigate how human-human communication supports intent alignment compared with human-LLM interaction, we designed a comparative study. The study involves two participants in the following two roles: (1) **User**: the user is the primary role who interacts with the LLM and the human assistant to complete a given task, (2) **Assistant**: the assistant is a human participant who has the role of helping the user complete a given task using LLM. This study employs a within-subject design, where the **user** engages in two distinct conditions: **user-LLM** interaction (human-LLM interaction) and **user-assistant** interaction (human-human interaction). In our study, we use GPT-4 [12] as an LLM model, and in the **user-assistant** interaction condition, the **assistant** can utilize the same LLM (GPT-4) to provide appropriate support to the **user**.

Task Selection Given the importance of observing intent alignment processes and the variability in user preferences and needs, selecting tasks that catered to diverse user needs was crucial. To effectively observe the process of intent alignment, we intended to design a task that can motivate the users. Therefore, we asked the study participants about tasks they personally found motivating beforehand. Subsequently, we selected two types of tasks: ones that involve relatively straightforward criteria for defining desired outcomes (e.g., choosing gifts for family/friends) and tasks with more complex criteria (e.g., learning web development skills).

Study Procedure Eight participants (3 male, 5 female, average age 25.8) with prior experience using LLMs were recruited for this study. They were paired into four distinct teams, each consisting of a **user** and an **assistant**. Two teams were assigned to each task among the four teams. During the study sessions, the **user** engaged in conversations with both an LLM and with a human **assistant** for the same task. When conversing with LLM, the **user** used the ChatGPT (model: GPT-4) interface. When conversing with human **assistant**, they had the option to use voice, and a shared document¹ facilitated the conveyance of information that might be challenging to express verbally. The time limit for each task was set at a maximum of 15 minutes. After finishing two tasks, participants underwent semi-structured interviews to provide insights into their experiences. The study sessions lasted approximately one hour and were conducted using the zoom² with cameras turned off to eliminate non-verbal cues.

3 STUDY RESULT: INTENT ALIGNMENT STRATEGY

Study Result With participants' consent, we recorded and transcribed conversations and interviews from the study sessions and analyzed them to identify the strategies employed by individuals to achieve intent alignment. While a thorough analysis is ongoing,

we want to share our initial results about the strategies adopted by human assistants and the strategies exhibited by users during conversations with human assistants, which were not evident in interactions with LLMs. The summarized results are presented in table 1 and 2.

Assistant's Strategies Our initial findings indicate that human assistants better align with user intent by (1) actively requesting information from the user, (2) providing tailored responses based on previous conversation history and user responses, and (3) actively seeking feedback on responses and interactions from the user. We found that human assistants proactively engage with users' inquiries, requesting additional information, asking clarification questions, or inquiring about potential areas for improvement in their responses. It makes users provide more diverse and detailed contextual information, and it helps intent alignment. Moreover, we observed that when the human assistant provides responses to the user, it filters and offers information gradually, reflecting on previous conversations with the user, rather than presenting a large amount of information at once. For example, when the human assistant used LLM to respond to the user, they filtered and provided only the information they deemed necessary for the user based on previous conversations, amidst the many options provided by the LLM. On the other hand, we observed passive behavior in LLM interactions, only responding passively to user prompts and failing to go further from user prompts. Additionally, LLMs tend to present large lists at once in a response, often containing excessive, general, or even irrelevant information. It causes users to perceive that their intentions are not fully reflected in the responses.

User's Strategies Beyond comparing the differences between human assistants and LLMs, we delved into analyzing interaction patterns when users conversed with LLMs compared to human assistants. Through this analysis, when the user interacts with the human assistant, we found that users (1) actively provided feedback in case of dissatisfaction with responses, (2) often interrupted the middle of the assistant's responses to ask questions if there were unknown parts, (3) actively offered additional context while responding to the assistant's questions, and (4) requested answers from the assistant's perspective. If we examine the reasons why users exhibit these different behavioral strategies, it can be attributed to the differences in how human assistants and LLMs provide responses to users. Unlike human assistants, who provide a few filtered suggestions gradually, LLMs generate a large amount of suggestions at once. Therefore, users tend to behave by selecting their preferred options from the many suggestions provided by LLMs, rather than providing feedback on each suggestion to convey their intent more clearly.

A significant factor contributing to users exhibiting different behavioral strategies may be the ability to converse using voice. Unlike text, voice communication is conducive to providing feedback in short turns and displaying reactions while the other party is speaking. Users can interrupt in the middle if there are parts they do not understand and can also provide ambient reactions like "uh-huh" or "okay" intermittently if they understand the human assistant's answer.

¹<https://docs.google.com/>

²<https://zoom.us/>

Type	Assistant's strategy
Information Request to User	<ul style="list-style-type: none"> - Requests information directly to user - Asks questions to understand the user's context - Asks questions to understand the user's context - Previous user's response specification question - Asks why the user said something. - Asks clarifying questions about the user's previous response. - Asks for examples
Reflecting Previous Conversation	<ul style="list-style-type: none"> - Maintains previous recommendations while incorporating the user's previous response. - Provides answers reflecting the context of the previous conversation with the user. - Provides information that the user may need based on the context of the conversation.
Feedback Request to User	<ul style="list-style-type: none"> - Asks for the user's opinion when providing a response. - Asks which response the user liked best. - Asks if there is anything that can be improved in the response. - Asks if the user has any additional questions. - Asks if the user is satisfied with the response.

Table 1: The types of strategies adopted by the assistant during the conversation with the user to achieve intent alignment.

Type	User's behavior with human assistant	Comparison to User's behavior with LLM
Feedback on Dissatisfactory Assistant Responses	<ul style="list-style-type: none"> - Evaluates the Assistant's response and provides the context/reason of the evaluation - Informs problems in Assistant response - Corrects misunderstandings in Assistant response - Asks questions about unknown parts in Assistant response. 	When talking to LLM, there are many options that LLM suggests at once, so users tend to move on to the parts they don't like and focus on the good parts of the answers to continue the conversation.
Interrupting Assistant Mid-Response	<ul style="list-style-type: none"> - Interrupts in the middle of the Assistant's response and asks something 	Users and LLM rarely interrupt during the conversation
Provide Additional Context	<ul style="list-style-type: none"> - Provides additional context when answering the Assistant's questions. - Provides additional context on their own after the conversation. 	LLM responds passively to users' prompts and does not actively ask questions. Therefore, the user rarely gives additional context while answering LLM's questions.
Request an Assistant Perspective Answer	<ul style="list-style-type: none"> - Asks for Assistant's opinion - Asks about Assistant's experiential content. 	This could be because the user perceives the human assistant as a person.

Table 2: The types of strategies exhibited by the user during the conversation with the human assistant to aid intent alignment, and comparison of user behaviors during interactions with LLMs.

Another reason may be due to the difference in the active/passive attitudes of human assistants and LLMs. Human assistants demonstrate an active attitude in conversations with users by actively asking questions or requesting feedback. Therefore, users tend to adopt a similar active communication style with human assistants or reveal their intent more clearly by responding to the assistant's questions.

4 DISCUSSION

Building upon our study findings, we aim to delve into a comprehensive discussion regarding the strategies humans employ to achieve intent alignment during human-human conversations, their respective advantages and limitations, and how these strategies can be integrated into human-AI interaction, specifically into LLM-based user interfaces.

Firstly, **incorporating proactive engagement features** within AI interfaces can enhance intent alignment. Recently, many systems based on LLMs are seeking to introduce proactive features. These

systems primarily provide users with prompt suggestions (e.g., Notion [11], Grammarly [5]). Subramonyam et al. also discussed that proactively suggesting ideas for prompting users is a useful design pattern that helps users obtain the intended answers [16]. However, by analyzing the behavioral strategies of human assistants in our study, we found that, beyond suggesting prompts for users, actively seeking information from users and requesting feedback on responses contributes to intent alignment. Therefore, designing AI interfaces that mimic this proactive behavior, such as prompting users for additional context or feedback, could improve alignment between user intents and AI responses.

Secondly, the findings highlighted **the importance of tailoring responses** to the user's context and preferences. Human assistants adjusted their responses based on previous conversations and the user's specific needs, leading to more aligned interactions. To achieve more aligned human-AI interaction with individual users' context and intent, AI interfaces should be designed to dynamically adapt responses based on user history. We can get insights from

recent studies [4, 15] that have designed natural language interfaces based on the Grice's Cooperative principles [6] by considering the Recipient Design approach [16]. However, it is necessary to discuss how such design approaches can be applied to interfaces for generative AI, including LLMs.

Furthermore, **emulating the advantages of voice interfaces** in interactions with LLMs can facilitate intent alignment. Voice interfaces offer several benefits, including natural and fluid communication, real-time feedback, and short-turn conversations, allowing for quick exchanges of information. While conversing with LLMs primarily through text, participants found it advantageous to use voice communication with human assistants. Incorporating flexible communication channels into AI interfaces can accommodate diverse user preferences and enable more effective communication of complex information.

5 CONCLUSION

In conclusion, our study explores key strategies for enhancing intent alignment in human-AI interactions, compared with human-human interaction. We identified the intent alignment strategy that each user and human assistant takes in their conversation. Human assistants demonstrate proactive engagement, actively seeking information and feedback and reflecting previous conversation with users. While users adapt by offering feedback and interrupting for clarification. Our discussion underscores the importance of incorporating these strategies into AI interface design. These discussions can serve as a foundation for designing more effective and user-friendly AI interfaces, addressing challenges in user intent understanding and alignment with LLMs.

ACKNOWLEDGMENTS

This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)). This work was also supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2021-0-01347, Video Interaction Technologies Using Object-Oriented Video Modeling).

REFERENCES

- [1] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [2] Minsuk Chang, Stefania Druga, Alexander J. Fiannaca, Pedro Vergani, Chinmay Kulkarni, Carrie J Cai, and Michael Terry. 2023. The Prompt Artists. In *Proceedings of the 15th Conference on Creativity and Cognition* (. Virtual Event, USA.) (C&C '23). Association for Computing Machinery, New York, NY, USA, 75–87. <https://doi.org/10.1145/3591196.3593515>
- [3] Brian Christian. 2020. *The alignment problem: Machine learning and human values*. WW Norton & Company.
- [4] Judit Dombi, Tetyana Sydorenko, and Veronika Timpe-Laughlin. 2022. Common ground, cooperation, and recipient design in human-computer interactions. *Journal of Pragmatics* 193 (2022), 4–20.
- [5] Grammarly. [n. d.]. AI Writing Assistance | Grammarly. <https://www.grammarly.com/ai>. (Accessed on 03/02/2024).
- [6] HP Grice. 1975. Logic and Conversation. *Syntax and Semantics* 3 (1975), 43–58.
- [7] Ellen Jiang, Edwin Toh, Alejandra Molina, Kristen Olson, Claire Kayacik, Aaron Donsbach, Carrie J Cai, and Michael Terry. 2022. Discovering the Syntax and Strategies of Natural Language Programming with Generative Language Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>New Orleans</city>, <state>LA</state>, <country>USA</country>, </conf-loc>) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 386, 19 pages. <https://doi.org/10.1145/3491102.3501870>
- [8] Yoonsu Kim, Jueon Lee, Seoyoung Kim, Jaehyuk Park, and Juho Kim. 2023. Understanding Users' Dissatisfaction with ChatGPT Responses: Types, Resolving Tactics, and the Effect of Knowledge Level. *arXiv preprint arXiv:2311.07434* (2023).
- [9] Jakob Nielsen. 2023. AI: First New UI Paradigm in 60 Years. <https://www.nngroup.com/articles/ai-paradigm/>. (Accessed on 02/18/2024).
- [10] Andy Norman. 2016. Why we reason: intention-alignment and the genesis of human rationality. *Biology & Philosophy* 31 (2016), 685–704. <https://doi.org/10.1007/s10539-016-9532-4>
- [11] Notion. [n. d.]. Notion AI. <https://www.notion.so/help/guides/category/ai>. (Accessed on 03/02/2024).
- [12] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [13] Pedro A. Ortega, Vishal Maini, and the DeepMind Safety Research. 2018. Building safe artificial intelligence: specification, robustness, and assurance | by DeepMind Safety Research | Medium. <https://deepmindsafetyresearch.medium.com/building-safe-artificial-intelligence-52f5f75058f1>. (Accessed on 02/19/2024).
- [14] Bernardino Romera-Paredes, Mohammadamin Barekatian, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. 2023. Mathematical discoveries from program search with large language models. *Nature* (2023), 1–3.
- [15] Vidya Setlur and Melanie Tory. 2022. How do you Converse with an Analytical Chatbot? Revisiting Gricean Maxims for Designing Analytical Conversational Behavior. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>New Orleans</city>, <state>LA</state>, <country>USA</country>, </conf-loc>) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 29, 17 pages. <https://doi.org/10.1145/3491102.3501972>
- [16] Hari Subramonyam, Roy Pea, Christopher Pondoc, Maneesh Agrawala, and Colleen Seifert. 2024. Bridging the Gulf of Envisioning: Cognitive Challenges in Prompt Based Interactions with LLMs. <https://doi.org/10.1145/3613904.3642754>
- [17] Michael Terry, Chinmay Kulkarni, Martin Wattenberg, Lucas Dixon, and Meredith Ringel Morris. 2023. AI Alignment in the Design of Interactive AI: Specification Alignment, Process Alignment, and Evaluation Support. *arXiv preprint arXiv:2311.00710* (2023).
- [18] Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. 2005. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences* 28, 5 (2005), 675–691. <https://doi.org/10.1017/S0140525X05000129>
- [19] 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* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 437, 21 pages. <https://doi.org/10.1145/3544548.3581388>