

# What You Need is What You Get: Theory of Mind for an LLM-Based Code Understanding Assistant

Jonan Richards , Mairieli Wessel 

*Radboud University*

Nijmegen, The Netherlands

jonan.richards@ru.nl, mairieli.wessel@ru.nl

**Abstract**—A growing number of tools have used Large Language Models (LLMs) to support developers’ code understanding. However, developers still face several barriers to using such tools, including challenges in describing their intent in natural language, interpreting the tool outcome, and refining an effective prompt to obtain useful information. In this study, we designed an LLM-based conversational assistant that provides a personalized interaction based on inferred user mental state (e.g., background knowledge and experience). We evaluate the approach in a within-subject study with fourteen novices to capture their perceptions and preferences. Our results provide insights for researchers and tool builders who want to create or improve LLM-based conversational assistants to support novices in code understanding.

**Index Terms**—LLM, LLM-based tools, comprehension support, user perceptions, novice developers

## I. INTRODUCTION

Large Language Models (LLMs) have shown promising results supporting code comprehension tasks and are often perceived as useful by developers [1]–[3]. However, the extent to which users benefit from these tools depends on their background knowledge and experience levels [3], [4]. Developers often experience difficulty in writing prompts that are effective at addressing their information needs [2], [3], particularly novices [5]. Many studies suggest the potential of AI assistants learning from users to create personalized guidance [1], [4], [6]. In fact, successful explanations should fill the gaps in questioners’ background knowledge [7], while using language that is understandable to the questioner [8]. Combined, these findings indicate a need for LLMs to infer and adapt to developers’ needs, intents, knowledge, experience, and preferences.

Humans’ ability to reason about what others are thinking is called Theory of Mind (ToM). LLMs have performed well on a range of ToM tasks [9]–[11]. Whether this is due to LLMs having an innate ToM capacity [12] or their ability to find shallow heuristics to perform well [13] remains an open debate. Still, by reflecting on users’ minds, LLMs can produce insightful information in practical settings [14].

We hypothesize that LLMs can gather a substantial amount of useful ToM information from interactions between a developer trying to understand unfamiliar code, and a conversational agent. This led us to create a conversational agent called **ToMMY** (Theory of Mind: Making explanations Yours). ToMMY comprises a chain of prompts meant to generate insights about a developer’s mental state. It then uses these

insights to provide personalized explanations about code. In this paper, we evaluate how ToMMY compares to a more basic agent. Our work is the first application and evaluation of ToM prompting techniques in LLMs supporting Software Engineering activities. This study focuses on the following research questions:

- RQ1** How do novices interact with LLM-based assistants when trying to understand unfamiliar code?
- RQ2** How does using ToMMY impact novices’ code understanding?
- RQ3** How do novices perceive interacting with LLM-based assistants to understand unfamiliar code?

To answer these research questions, we conducted a within-subject study with 14 novices, including undergraduate and graduate students. Participants interacted with each approach (ToMMY vs. simple conversational agent) individually and comparatively. Our findings reveal that novices exhibited distinct interaction styles based on whether they phrased some questions as hypotheses (or not). While interacting with ToMMY, participants less frequently stated their intent or asked follow-up questions, and slightly more often provided instructions regarding the response format. Also, using ToMMY had distinct impacts on novices’ code understanding depending on their interaction styles. We made our supplementary material available for replication purposes [15].

## II. RELATED WORK

Researchers have investigated how developers use and interact with LLMs for development-related tasks [1], [2], [16]–[18]. Ross et al. [1], [18] suggests adapting programming assistants to the individual strengths and needs of its users toward providing personalized responses. A growing number of tools have also recently brought LLMs into the code editor to support code understanding [3], [19], [20]. Nam et al. [3], for instance, proposed an approach that supports on-demand explanations of code. Complementing previous literature, we focused on personalizing developers’ interactions with conversational LLM-based code understanding assistants using existing Theory of Mind prompting techniques.

## III. TOMMY’S DESIGN

The design of ToMMY was motivated by findings from previous studies of LLMs, which found that predicting and reflecting on mental states improve LLMs’ performance [10],

[14]. Furthermore, perspective-taking is an effective way to enhance the ToM performance of LLMs [11], [21]. Wilf et al. [21] highlights the need to separate mental state inference and question answering into two prompts. The prompt engineering process was conducted through manual evaluation, in addition to a pilot study used to evaluate the last iteration of prompts.

The first of ToMMY’s *three prompts* identifies aspects of the user’s mental state relevant to producing a response, and phrases these as open questions about the user. Not having this step, and directly prompting for the user’s mental state, was found to produce irrelevant mental state data. In the second prompt, we instruct the LLM to take the perspective of the user to answer these questions, based on the conversation history. These answers constitute the inferred mental state. Finally, in the third prompt, the LLM is instructed to respond to the user, considering the mental state.

Each of ToMMY’s prompts includes context for the LLM, containing a code snippet, the conversation history, and the user’s latest input. The basic agent, used as a control condition in the user study, comprises a single prompt containing the same context, and an instruction to respond to the user.

The LLM we used in our evaluation was GPT-3.5, a common choice in research on LLMs for Software Engineering [22]. We set the model’s temperature, a hyper-parameter related to output randomness, to 0. This helped isolate the impact of using different prompts. We limited the history to the 10 most recent interactions, and capped the length of input messages to 1,000 characters, to prevent the prompt from overflowing the LLM’s token limit. A pilot study exposed no issues with this strategy.

#### IV. RESEARCH DESIGN

When preparing and executing this study, we followed the guidelines set by Radboud’s Research Ethics Committee and received approval from that committee to conduct our study.

**Pilot sessions.** We conducted pilot sessions with three novice programmers: one PhD student, one master’s student, and one third-year bachelor’s student. The participants suggested a few minor adjustments, which were incorporated into ToMMY and the experiment instruments. Due to the participant’s fatigue, we time-boxed the code understanding tasks (20 minutes) and quizzes (10 minutes). Data from pilot sessions were discarded.

**Participant recruitment.** We used a convenience sampling strategy to recruit participants for our experiment sessions, inviting students through (i) short presentations at lectures, (ii) study programme-wide messaging groups, and (iii) outreach to personal contacts. The study was not part of any course; all participants were volunteers and signed a consent form before their sessions. The only prerequisite was a low-level programming experience (e.g., having taken a programming-related course or having coded as a hobby). In total, 6 undergraduate and 8 graduate students were recruited. Their experience with programming ranged between low and experienced (with an average 5.7 on a scale from 1 to 10). Participants also reported having low (2), moderate (7), considerable (2), and extensive (1) familiarity with LLMs.

**Experimental sessions.** We conducted a series of synchronous within-subject sessions with participants using both ToMMY (*treatment*) and the basic agent (*control*). Using variance minimization [23], participants were sorted into two groups balanced in terms of participants’ programming experience, LLM familiarity, and background knowledge regarding natural language processing and data analysis. This grouping was used to determine the order in which the two agents were encountered, mitigating order effects.

Before the session, participants received email instructions, a survey with demographic questions, and a consent form. Participants could choose between online or in-person sessions. Each session started with a brief explanation of the research objectives and guidelines. Then, participants were asked to complete two code comprehension tasks (one with each conversational agent). Participants were presented with a different code snippet for each task and tasked with understanding it as thoroughly as possible within a 20-minute time limit.

We induced the participant’s need to interact with the agents by choosing code snippets that require domain knowledge to understand (i.e., data analysis and natural language processing). The two code snippets written in Python were carefully selected from open-source repositories on GitHub and not used for prompting engineering. Code comments were removed to increase the need for understanding, but typos or other implementation choices were left.

After each task, to measure their code understanding we asked participants to complete a few closed-ended questions inspired by previous literature [24] (max. 10 minutes). We also asked them to complete a mid-term survey (right after the first task) and a post-study survey at the end of the experiment to elicit their perceptions of the approaches.

**Web-based evaluation interface.** To enable the user study, we developed a custom-built web-based tool. This tool included a chat interface (Figure 1-2) for interacting with the conversational agents, alongside a code snippet (Figure 1-4) to mimic the experience of being in a code editor. It also displayed instructions and allowed participants to navigate the study steps (i.e., tasks, quiz questions, and evaluation forms). Participants could track their progress in the left-side menu (Figure 1-1). A timer for each step was also displayed (Figure 1-3).

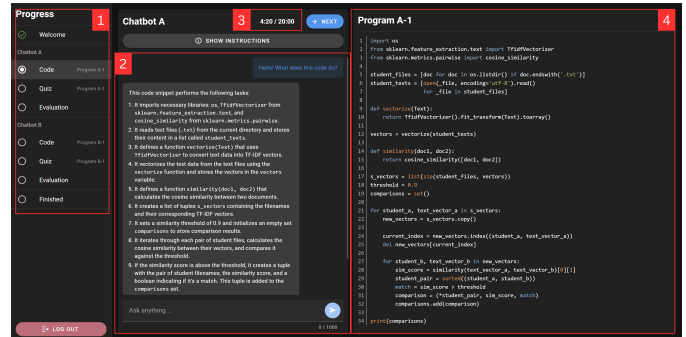


Fig. 1: Web-based interface designed for the experiment

**Data analysis.** We analyzed participants’ interactions with the agents (RQ1) using *open* card sorting [25]. First, the first two authors independently analyzed the inputs for the first two participants (cards) and applied codes, sorting them into meaningful groups. A discussion meeting followed this step until a consensus was reached on the categorization. The first author then coded the remainder of the participant interaction data. The researchers analyzed the categories to refine the classification and group-related codes into higher-level themes.

To answer RQ2, we used linear regression to assess each approach’s impact, considering the effect of participants’ programming experience, background experience, and familiarity with LLM tools. We chose these three metrics as our independent variables and a variable to indicate whether a participant used our approach for the task. To measure code understanding, we considered the number of correct answers on the quizzes and the time participants took to complete the task. We fit two separate regression models to these two dependent variables. After observing distinct patterns during the qualitative analysis of participants’ interactions, we performed additional analysis for code understanding. We fit two more regression models on the quiz scores: one for participants who interacted by phrasing hypotheses and the other for participants who did not.

In RQ3, we assessed the perceived usefulness, ease of use, and cognitive load of working with each agent by applying the Technology Acceptance Model (TAM) [26] and Task Load Index (NASA-TLX) [27], respectively. Additionally, we quantitatively analyzed the open feedback collected after completing each task. This was processed using card sorting, similar to RQ1.

## V. FINDINGS

### A. Interaction (RQ1)

Participants did not interact more frequently, or for a longer time, with either agent (Table I). This was determined using two-sided paired t-tests (Shapiro-Wilk tests showed normality of differences for these variables,  $p = 0.88$  and  $0.43$  respectively).

	Control		ToMMY		<i>p</i> -value
	mean	std	mean	std	
# of interactions	7.21	4.56	7.14	4.70	0.96
Interaction time (s)	739.00	363.66	717.14	319.38	0.72

TABLE I: Average number of interactions and interaction time per participant for each task, broken down by model.

In our qualitative analysis of the participants’ inputs, we grouped the codes into 5 categories, which are described in the following sections.

**Question.** All but four of the 201 recorded interactions contained a question, and only one input contained two questions. These question were phrased either as an open question (“What does ... do?”), a hypothesis (“Does ... do ...?”), an instruction (“Explain ...”), or implicitly (e.g. by only providing a line of code for the agent to explain). We clustered

participants based on if they phrased at least one question as a hypothesis, and found that 6 out of 14 participants did. We found that this cluster had significantly more interactions and spent a significantly longer time interacting with the agents than the participants who did not use hypotheses (using two-sided Mann-Whitney U tests, Figure 2). The difference in programming experience and LLM familiarity between these clusters of participants was negligible.

	No hypotheses		Hypotheses		<i>p</i> -value
	mean	std	mean	std	
Programming exp.	4.81	1.94	4.67	1.83	0.95
LLM familiarity	2.88	0.83	3.17	0.98	0.67
# of interactions	10.25	6.86	19.83	5.49	0.02*
Interaction time (s)	1072.25	490.63	1968.00	443.95	0.01**

\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

TABLE II: Participants’ background experience and interaction statistics, broken down by participants who did not (left) and did (right) use hypotheses.

**Question target.** This can be either a fine-grained code feature (such as syntax, a variable, a line of code or multiple lines of code); the entire snippet; an external code feature (such as an external function call or a library); or a programming concept. Figure 2 shows that participants who did not phrase questions as hypotheses tended to ask more questions about the entire snippet, compared to participants who did employ hypotheses. Instead, the latter more often targeted their questions at fine-grained code features. There were no discernible patterns regarding which agent was used.




Fig. 2: Distribution of question targets, broken down by participants who did not (left) and did (right) use hypotheses.

**Question intent.** Most questions did not specify what aspect of a code feature they wanted the agent to explain (e.g. in “What does ... do?”). Sometimes, however, the intended aspect was explicitly provided. We observed high-level, abstract aspects such as the purpose, rationale, or outcome of code, and low-level aspects such as the value of an expression, the effect of modifying a variable, or a trace of multiple lines of code. There was a slight difference in the proportion of interactions that contained an explicit intent between the two agents: 30% for the control agent and 25% for ToMMY. This difference was also observed between participants using and not using hypotheses: 29% and 26%, respectively.

**Conversation.** Some of the inputs exhibited additional conversational features, such as following up on a question posed to

the agent’s response or polite phrasing (e.g. “Can you explain ...?”). There was a noticeable difference in the prevalence of follow-up questions: 15% of interactions for participants who did use hypotheses and 10% for those who did not, and 15% for the control approach compared to 11% for ToMMY.

**Instruction.** In some instances, participants provided instructions on how the response should be formatted, such as asking for a detailed, short, step-by-step, or summarized response. This happened slightly more often for ToMMY than for the control approach: 5% compared to 2% of interactions, respectively. There was a negligible difference when looking at the participants grouped by their use (or not) of hypotheses.

### B. Code understanding (RQ2)

	Quiz score $N=28, R^2_{CS}=0.15$			Quiz time (s) $N=28, R^2_{CS}=0.44$		
	coeff.	std	p-value	coeff.	std	p-value
Intercept	1.19	0.27	0.00***	593.15	110.16	0.00***
Uses ToMMY	0.05	0.10	0.61	-28.94	43.24	0.50
Programm. exp.	0.10	0.03	0.00**	-39.38	12.40	0.00**
Domain exp.	-0.01	0.03	0.71	-9.42	8.96	0.29
LLM familiarity	-0.09	0.04	0.02*	36.06	25.91	0.16

\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

TABLE III: Regression models showing the effect of several variables on quiz scores and completion times. The number of observations (2 per participant) is indicated with  $N$ .

Table III shows that the only factors found significantly to impact code understanding were programming experience and familiarity with LLMs. The positive coefficient for the scores model and negative coefficient for the time model shows that greater programming experience enables participants to score higher and finish the quizzes faster. Curiously, the reverse is true for familiarity with LLMs, although insignificantly so for quiz completion time.

	No hypotheses $N=16, R^2_{CS}=0.26$			Hypotheses $N=12, R^2_{CS}=0.19$		
	coeff.	std	p-value	coeff.	std	p-value
Intercept	0.93	0.27	0.00***	1.66	0.36	0.00***
Uses ToMMY	0.34	0.11	0.00**	-0.28	0.16	0.07
Programm. exp.	0.12	0.04	0.00**	0.06	0.03	0.05
Domain exp.	-0.05	0.03	0.08	0.01	0.03	0.80
LLM familiarity	-0.06	0.04	0.17	-0.15	0.09	0.09

\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

TABLE IV: Regression models showing the effect of several variables on quiz scores, wherein participants are grouped by their interaction style. The number of observations (2 per participant) is indicated with  $N$ .

No significant effects of using ToMMY are found when taking all participants into account. However, when fitting separate regression models to participants based on their interaction style, we do observe an effect (Table IV). Using ToMMY contributes greatly to quiz scores for participants who did not phrase questions as hypotheses, but negatively impacts the scores of participants who did. The latter result is not significant, although it approaches the significance threshold of  $p < 0.05$ .

### C. Perceptions (RQ3)

**Metrics.** No significant differences in TAM and TLX scores were found between the two approaches (using two-sided Wilcoxon signed-rank tests, Table V). What stands out, however, is a lower mean but large standard deviation for the “effort” item of the TLX scores for ToMMY ( $M=2.77, SD=4.46$ ) compared to the control agent ( $M=4.00, SD=1.68$ ). This may suggest that ToMMY reduced perceived effort compared to the control approach, but this effect greatly varied per participant.

	Control		ToMMY		p-value
	mean	std	mean	std	
Perceived usefulness	6.18	0.76	6.09	0.76	0.66
Perceived ease of use	6.27	0.73	6.05	0.97	0.54
TLX score (avg.)	0.45	1.69	0.38	2.58	0.74

TABLE V: Comparison of TAM and TLX scores.

**Feedback.** Both agents were seen as helpful and knowledgeable. Participants indicated that they liked the way their questions were answered and related back to the code. ToMMY was reported to understand the participants’ questions well and produce clear answers ( $n=3$  and  $4$ , respectively). This sentiment was also shared, although slightly less frequently, for the control approach ( $n=2$  and  $3$ , respectively). Furthermore, P11 and P14 noted that they had to rephrase their question or provide more details to the control agent after it did not understand their question on the first try. P6 and P7 reported not always understanding the control agents’ responses, causing P6 to ask for clarification several times.

Both the control approach and ToMMY were found to respond thoroughly, and in some cases with too much detail. Looking at Figure 3, ToMMY seems to produce slightly shorter responses. However, when questions are phrased as hypotheses, ToMMY much more consistently replies with a shorter answer than the control approach. This suggests that ToMMY is better at recognizing when users just want to validate whether their interpretation of the code snippet is correct. Curiously, P1 found that the control approach appeared to “*better fit the length of the answer to the question*”, which does not align with our interpretation of Figure 3. Another counter-intuitive finding is P6 perceiving the control agent to produce shorter responses, while P4 perceived the opposite. Occasionally, the control agent did not answer as extensively as participants would have liked (P6, P10), a problem that was not encountered with ToMMY.

P4 noticed that ToMMY often seemed to respond with single paragraphs, whereas the control usually formatted its responses by emphasizing certain words. While the participant perceived ToMMY’s answers to be much shorter, they preferred the control agent’s longer, formatted responses. This was echoed by P13, reporting that ToMMY alternated between formatting its responses as lists and as single paragraphs, which they perceived as inconsistent. Two more participants (P8, P11) commented positively on the format of the control approach’s responses. Additionally, P3 and P4 noted that ToMMY tended to end its answers with “*if you have more*






Fig. 3: Distribution of response length, broken down by model and whether or not the question was phrased as a hypothesis.

questions, feel free to ask”, which they experienced to be repetitive and useless.

## VI. DISCUSSION

Several participants interacted with the agents by phrasing hypotheses, often aimed at fine-grained code features, and asking many follow-up questions. This aligns with Brook’s top-down model of code comprehension [28], which poses that programmers iteratively generate, test, and refine hypotheses by scanning the code. The other participants asked only open questions, sometimes phrased as instructions, and relatively frequently aimed these at the entire snippet. This corresponds to the bottom-up model of code comprehension [29], stating that programmers recursively answer open “what” questions [30] about chunks of code and combine the resulting knowledge until the entire program is understood. These interaction styles were not associated with differences in background experience.

Participants more often reported ToMMY to understand their question and produce clear answers than the control agent. Additionally, some participants commented negatively about the control agent on these aspects. Combined with participants less often explicitly stating their intent or asking follow-up questions to ToMMY, this suggests that ToMMY was able to better recognize participants’ intent and adapt responses to their background experience. ToMMY also adjusted the length of its responses to the type of question more than the control agent did. However, this was not always perceived as such by participants. We suspect this is due to ToMMY’s responses often being formatted as a single paragraph, which might introduce usability issues seen in the literature [16].

No differences in perceived usefulness, ease of use, and cognitive load were found between the two agents. Using ToMMY had a positive impact on quiz scores for participants with a bottom-up comprehension style, and a negative effect for participants exhibiting a top-down comprehension style. The latter group of participants more frequently asked targeted questions with explicit intent, which may have diminished the benefit of ToMMY’s mental state inference. For all of these metrics, a negative impact of ToMMY’s response format may have obscured other positive effects.

Our results imply LLMs can personalize responses independently on some aspects, but may need guidance on others. This guidance could entail explicitly prompting the LLM for certain

elements of mental state (“What preferences does the user have regarding the response format?”), or even providing explicit instructions (“If the user asks a yes-no question, provide a minimal response.”).

Future research include retaining and building on knowledge inferred in earlier interactions with ToMMY, to provide more robust and accurate mental inference. Retrieval-Augmented Generation [31] has shown promising results for storing user’s mental data [14], and may be suitable for conversational programming assistants as well. Additionally, ToMMY being situated in a social, conversational environment allows for a more robust evaluation of new ToM prompting techniques than commonly-used basic tasks [32]. Also, designing an interface for users to inspect and provide feedback on the inferred mental data would bring ToMMY closer to allowing a “Mutual Theory of Mind” [33]. The resulting increase in transparency and fairness towards the user would also better align the agent with existing human-AI interaction guidelines [34].

## VII. LIMITATIONS

Prompt engineering and validation are notably difficult. We validated our prompt design choices through manual observations and a small-scale pilot study. However, small changes to ToMMY’s prompts may lead to unexpected side effects and ultimately different results. Moreover, our sample was composed of students. Although they are novices [35], we acknowledge that additional research is necessary to consider the perspective of practitioners experienced with LLMs. Also, within-subject studies are vulnerable to learning effects, fatigue, and other order effects. Varying the order in which participants interacted with the agents counteracted this. While the code snippets were carefully selected and modified to have similar complexity, using other snippets may produce different outcomes. Since we leverage qualitative research methods to categorize the open-ended questions asked in our surveys and the participant’s interactions with the agents, we may have introduced categorization bias. To mitigate this bias, we conducted this process in pairs and carefully discussed categorization among the authors.

## VIII. CONCLUSION

In this paper, we took the first steps towards personalizing developers’ interactions with a code understanding assistant by proposing ToMMY, an assistant that can infer developers’ mental states and adapt its responses. Although some novice programmers reported a good experience when using ToMMY, more research is still needed to understand how to structure its content. Interestingly, we also found that novices’ interaction styles with the agent can significantly impact ToMMY’s effectiveness. Researchers and tool builders can leverage our approach to better adjust to users’ needs, ensuring less cognitive load from users when interacting with the tool.

**Acknowledgments.** We thank the novice programmers who spent their time participating in our research.

## REFERENCES

- [1] S. I. Ross, F. Martinez, S. Houde, M. Muller, and J. D. Weisz, "The Programmer's Assistant: Conversational Interaction with a Large Language Model for Software Development," in *Proceedings of the 28th International Conference on Intelligent User Interfaces*. Sydney NSW Australia: ACM, Mar. 2023, pp. 491–514.
- [2] R. Khojah, M. Mohamad, P. Leitner, and F. G. De Oliveira Neto, "Beyond Code Generation: An Observational Study of ChatGPT Usage in Software Engineering Practice," *Proceedings of the ACM on Software Engineering*, vol. 1, no. FSE, pp. 1819–1840, Jul. 2024.
- [3] D. Nam, A. Macvean, V. Hellendoorn, B. Vasilescu, and B. Myers, "Using an LLM to Help With Code Understanding," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. Lisbon Portugal: ACM, Apr. 2024, pp. 1–13.
- [4] G. Pinto, C. De Souza, T. Rocha, I. Steinmacher, A. Souza, and E. Monteiro, "Developer Experiences with a Contextualized AI Coding Assistant: Usability, Expectations, and Outcomes," in *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering - Software Engineering for AI*. Lisbon Portugal: ACM, Apr. 2024, pp. 81–91.
- [5] S. Nguyen, H. M. Babe, Y. Zi, A. Guha, C. J. Anderson, and M. Q. Feldman, "How Beginning Programmers and Code LLMs (Mis)read Each Other," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Honolulu HI USA: ACM, May 2024, pp. 1–26.
- [6] A. E. Hassan, G. A. Oliva, D. Lin, B. Chen, Z. Ming, and Jiang, "Rethinking Software Engineering in the Foundation Model Era: From Task-Driven AI Copilots to Goal-Driven AI Pair Programmers," arXiv:2404.10225, Apr. 2024.
- [7] J. Faye, "Explanation Explained," *Synthese*, vol. 120, no. 1, pp. 61–75, 1999.
- [8] D. Walton, "A new dialectical theory of explanation," *Philosophical Explorations*, vol. 7, no. 1, pp. 71–89, Mar. 2004.
- [9] S. R. Moghaddam and C. J. Honey, "Boosting Theory-of-Mind Performance in Large Language Models via Prompting," arXiv:2304.11490, Apr. 2023.
- [10] P. Zhou, A. Madaan, S. P. Potharaju, A. Gupta, K. R. McKee, A. Holtzman, J. Pujara, X. Ren, S. Mishra, A. Nematzadeh, S. Upadhyay, and M. Faruqi, "How FaR Are Large Language Models From Agents with Theory-of-Mind?" arXiv:2310.03051, Oct. 2023.
- [11] Q. Zhu, L. Chong, M. Yang, and J. Luo, "Reading Users' Minds from What They Say: An Investigation into LLM-based Empathic Mental Inference," arXiv:2403.13301, Mar. 2024.
- [12] M. Jamali, Z. M. Williams, and J. Cai, "Unveiling Theory of Mind in Large Language Models: A Parallel to Single Neurons in the Human Brain," arXiv:2309.01660, Sep. 2023.
- [13] N. Shapira, G. Zwirn, and Y. Goldberg, "How Well Do Large Language Models Perform on Faux Pas Tests?" in *Findings of the Association for Computational Linguistics: ACL 2023*, A. Rogers, J. Boyd-Graber, and N. Okazaki, Eds. Toronto, Canada: Association for Computational Linguistics, Jul. 2023, pp. 10 438–10 451.
- [14] C. Leer, V. Trost, and V. Voruganti, "Violation of Expectation via Metacognitive Prompting Reduces Theory of Mind Prediction Error in Large Language Models," arXiv:2310.06983, Oct. 2023.
- [15] J. Richards and M. Wessel, "Replication package for "What You Need is What You Get: Theory of Mind for an LLM-Based Code Comprehension Assistant"," Aug. 2024. [Online]. Available: <https://doi.org/10.5281/zenodo.13270637>
- [16] J. T. Liang, C. Yang, and B. A. Myers, "A large-scale survey on the usability of ai programming assistants: Successes and challenges," in *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, 2024, pp. 1–13.
- [17] H. Mozannar, G. Bansal, A. Fourney, and E. Horvitz, "Reading Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Honolulu HI USA: ACM, May 2024, pp. 1–16.
- [18] S. I. Ross, S. Houde, F. Martinez, M. Muller, and J. D. Weisz, "The Programmer's Assistant User Experience," in *28th International Conference on Intelligent User Interfaces*. Sydney NSW Australia: ACM, Mar. 2023, pp. 102–104.
- [19] M. Liffiton, B. E. Sheese, J. Savelka, and P. Denny, "Codehelp: Using large language models with guardrails for scalable support in programming classes," in *Proceedings of the 23rd Koli Calling International Conference on Computing Education Research*, 2023, pp. 1–11.
- [20] L. Yan, A. Hwang, Z. Wu, and A. Head, "Ivive: Lightweight anchored explanations of just-generated code," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024.
- [21] A. Wilf, S. S. Lee, P. P. Liang, and L.-P. Morency, "Think Twice: Perspective-Taking Improves Large Language Models' Theory-of-Mind Capabilities," arXiv:2311.10227, Nov. 2023.
- [22] X. Hou, Y. Zhao, Y. Liu, Z. Yang, K. Wang, L. Li, X. Luo, D. Lo, J. Grundy, and H. Wang, "Large Language Models for Software Engineering: A Systematic Literature Review," arXiv:2308.10620, Apr. 2024.
- [23] F. Sella, G. Raz, and R. Cohen Kadosh, "When randomisation is not good enough: Matching groups in intervention studies," *Psychonomic Bulletin & Review*, vol. 28, no. 6, pp. 2085–2093, Dec. 2021.
- [24] C. Izu, C. Schulte, A. Aggarwal, Q. Cutts, R. Duran, M. Gutica, B. Heinemann, E. Kraemer, V. Lonati, C. Mirolo, and R. Weeda, "Fostering Program Comprehension in Novice Programmers - Learning Activities and Learning Trajectories," in *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education*. Aberdeen Scotland UK: ACM, Dec. 2019, pp. 27–52.
- [25] T. Zimmermann, "Card-sorting: From text to themes," in *Perspectives on Data Science for Software Engineering*. Elsevier, 2016, pp. 137–141.
- [26] F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS quarterly*, pp. 319–340, 1989.
- [27] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.
- [28] R. Brooks, "Towards a theory of the comprehension of computer programs," *International Journal of Man-Machine Studies*, vol. 18, no. 6, pp. 543–554, Jun. 1983.
- [29] B. Shneiderman and R. Mayer, "Syntactic/semantic interactions in programmer behavior: A model and experimental results," *International Journal of Computer & Information Sciences*, vol. 8, no. 3, pp. 219–238, Jun. 1979.
- [30] S. Letovsky, "Cognitive processes in program comprehension," *Journal of Systems and Software*, vol. 7, no. 4, pp. 325–339, Dec. 1987.
- [31] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 9459–9474.
- [32] Z. Ma, J. Sansom, R. Peng, and J. Chai, "Towards A Holistic Landscape of Situated Theory of Mind in Large Language Models," *Findings of Empirical Methods in Natural Language Processing*, Nov. 2023.
- [33] Q. Wang and A. K. Goel, "Mutual Theory of Mind for Human-AI Communication," arXiv:2210.03842, May 2024.
- [34] S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collis, J. Suh, S. Iqbal, P. N. Bennett, K. Inkm, J. Teevan, R. Kikin-Gil, and E. Horvitz, "Guidelines for Human-AI Interaction," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Glasgow Scotland UK: ACM, May 2019, pp. 1–13.
- [35] I. Steinmacher, T. U. Conte, C. Treude, and M. A. Gerosa, "Overcoming open source project entry barriers with a portal for newcomers," in *Proceedings of the 38th International Conference on Software Engineering*, ser. ICSE '16. New York, NY, USA: Association for Computing Machinery, May 2016, pp. 273–284.