



# Bridging HCI and AI Research for the Evaluation of Conversational SE Assistants

Jonan Richards , Mairieli Wessel   
Radboud University  
Nijmegen, The Netherlands  
jonan.richards@ru.nl, mairieli.wessel@ru.nl

**Abstract**—As Large Language Models (LLMs) are increasingly adopted in software engineering, recently in the form of conversational assistants, ensuring these technologies align with developers’ needs is essential. The limitations of traditional human-centered methods for evaluating LLM-based tools at scale raise the need for automatic evaluation. In this paper, we advocate combining insights from human-computer interaction (HCI) and artificial intelligence (AI) research to enable human-centered automatic evaluation of LLM-based conversational SE assistants. We identify requirements for such evaluation and challenges down the road, working towards a framework that ensures these assistants are designed and deployed in line with user needs.

**Index Terms**—SE, LLM, HCI, conversational agent, evaluation

## I. INTRODUCTION

A variety of bots have been used in software engineering (SE) to automate tasks and increase developers’ productivity [1]. Large Language Models (LLMs) show much promise in powering these automations, increasingly supporting various activities, including code generation, code summarization, bug detection, and requirements analysis [2]. Recent research has demonstrated the potential of integrating multiple tasks in a single SE bot to reduce developers’ cognitive overload [3]. A similar trend is seen in the emergence of LLM-based “assistants” capable of supporting several SE activities simultaneously, both in industry [4], [5] and research [6].

Implementing such a multi-faceted assistant as a conversational agent is a logical choice [6]. LLM-based conversational assistants increase productivity over non-conversational tools by allowing developers to work towards their goals incrementally [7]. Moreover, they can facilitate developers’ learning [8], and natural language provides an intuitive mechanism for users to give feedback [1]. Developers greatly value conversational interaction in LLM-based SE assistants [5], [6].

As the capabilities of bots in SE evolve, constant evaluation is needed to ensure they are designed and deployed in accordance with users’ needs, which tools for SE often fail to do [9]. Human-centered methods from the field of human-computer interaction (HCI) include any method that involves humans, and are necessary for this alignment with users’ needs [10]. An important aspect of using human-centered methods is to design in iterations, while evaluating the tool with experts or users [10]. Evaluation with users is the most direct evidence of the performance of an SE tool [11] and is the gold standard for

conversational agents [12]. However, user studies are costly, time-intensive, and difficult to scale [11].

LLM-based SE tools, and specifically conversational assistants, suffer from a lack of robust evaluation methods [2], [5]. LLMs are highly sensitive to the phrasing of their prompts, or instructions, and the design of an LLM-based tool takes much “prompt engineering” to achieve the desired functionality [13]. As LLMs are non-deterministic and allow unrestricted textual interaction, human-centered methods such as expert analyses of LLMs’ responses [14] and user studies are infeasible for the scale and frequency of evaluation that is required during the design process. Since LLM-based tools for SE are commonly evaluated on benchmark datasets [2], a quick and more systematic alternative is evaluating small changes in prompts on a subset of a benchmark. However, LLMs’ responses to benchmark inputs are traditionally compared to corresponding reference outputs using similarity metrics such as BLEU, ROUGE, and METEOR [2]. These reference-based metrics fail to capture how LLM-based tools perform in practice and do not correspond well to human evaluation of natural language [15].

The high cost of user studies has led to the emergence of an alternative evaluation method in the field of HCI. LLM-based simulated users have been used to qualitatively evaluate interfaces [16], conversational agents [17], and also have the potential to be used for automatic, synthesized user studies of SE tools [18]. The limitations of traditional benchmarks for evaluating natural language generation and the infeasibility of performing user studies to guide prompt engineering have led to a similar development in the field of artificial intelligence (AI). In LLM-as-a-Judge approaches [19], an LLM is employed as a “judge”, replacing real humans to assess LLM-generated texts automatically. LLM-as-a-Judge has increasingly been used by practitioners to evaluate domain-specific LLM-based tools [4], and may be helpful in various SE tasks as well [20]. Although both discussed methods leverage LLMs to synthesize research data, they reflect the fields from which they have emerged. Simulated users have been employed in HCI research to generate artificial interaction data and qualitative insights, whereas LLM-as-a-Judge approaches are used as quantitative benchmarks in AI research. While simulated users have the potential to produce quantitative metrics and LLM-as-a-Judge approaches may offer qualitative insights, these areas are under-explored.

In this paper, we advocate combining these two methods from different fields to enable human-centered automatic evaluation of conversational LLM-based assistants for SE. We identify several requirements for such evaluation, discuss LLM-as-a-Judge approaches and simulated users, and propose combining insights from both methods into an approach we envision will address all requirements. Finally, we identify challenges to address when implementing the proposed approach. These findings were derived from an extensive literature search. Our contribution is these insights towards realizing scalable and comprehensive automatic evaluation of conversational LLM-based developer assistance, meant to not replace but complement manual evaluation in the human-centered design process.

## II. REQUIREMENTS

In this section, we describe several requirements for the automatic evaluation of conversational SE assistants.

### R1. Explore realistic conversations

Typically, benchmarks for large-scale evaluation of conversational agents come in either of two forms: single-turn question and answer pairs, or multi-turn dialogues [12]. However, single-turn interactions do not capture conversational dynamics spanning multiple turns, leaving a large part of the interaction space unexplored. Both multi-turn and single-turn dialogue datasets typically provide reference responses, ignoring the wide range of responses that may be valid and the many ways in which a conversation might unfold [12]. Additionally, evaluation of generated text based on similarity to a reference does not correspond well to how humans judge natural language [15]. Figure 1 demonstrates how reference-based evaluation of SE assistants’ responses may lead to flawed insights. Thorough evaluation of conversational SE assistants, therefore, requires reference-free exploration of realistic interactions over multiple turns.

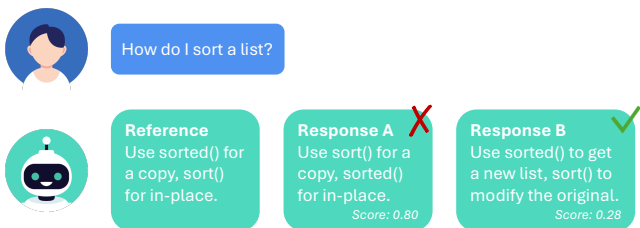


Fig. 1. Flaws in reference-based evaluation: incorrect response A scores high on BLEU-4, while correct response B scores low due to phrasing.

### R2. Support diversity

One individual’s needs may not fully correspond to another’s. Problem-solving tools often best support the information processing styles and motivation of young, male users, while other users’ perspectives are neglected [21]. This can lead to the occurrence of “inclusivity bugs”, which are issues with the usability of a tool that are triggered only for specific groups of users [22]. LLMs are especially vulnerable to these

inclusivity bugs, as their inherent unpredictability may cause disruptions for one user but not for others [23]. Instances of inclusivity bugs have been encountered in user studies of conversational LLM-based SE assistants, being triggered based on experience level [24], gender [8], or interaction style [25]. Prompt engineering can be used to put “guardrails” on interactions between LLMs and users to remove biases and ensure fairness [26]. First, however, conversational SE assistants need to be evaluated from the perspective of a diverse, representative user base to reliably detect inclusivity bugs.

### R3. Provide quantitative metrics

Incremental changes to LLMs’ prompts and parameters are enabled by constant evaluation and comparison [13]. Quantitative measures are invaluable for this benchmarking process, allowing for quick assessment of whether a change improved or reduced a tool’s performance [14]. However, the sensitivity of LLMs to its input may cause evaluation to reflect superficial aspects of the prompt’s phrasing [27]. Evaluating a large number of paraphrased prompts instead would allow designers to effectively assess the impact of changing the underlying structure of a prompt [14], [27], making the prototyping process of an LLM-based SE assistant more systematic and robust. Applying such techniques, however, requires automatic quantitative metrics. Finally, quantitative benchmarks are useful in demonstrating the effectiveness of new assistants [10] and in comparing them with existing solutions.

### R4. Provide qualitative insights

While quantitative metrics do allow a designer to quickly assess the impact of a change, they do not offer any information on usability issues or potential improvements. Conversational agents are often tested with users to gather qualitative feedback throughout the design process [28]. LLM-based SE assistants are no different, and feedback is essential to ensure these assistants are developed in line with users’ needs [5]. Automatic generation of qualitative feedback on SE assistants would allow designers to anticipate and fix usability issues [18], complementing user studies.

## III. CURRENT SITUATION

Here, we discuss user simulations and LLM-as-a-Judge approaches and how these methods address the identified requirements. A comparison of the methods can be found in Table I.

TABLE I  
COMPARISON BETWEEN SIMULATED USERS AND LLM-AS-A-JUDGE

Requirement	Simulated Users	LLM-as-a-Judge
R1. Explore realistic conversations	✓	✗
R2. Support diversity	✓	✓
R3. Provide quantitative metrics	?	✓
R4. Provide qualitative insights	✓	?

### A. Simulated Users

Recently, researchers have explored leveraging LLMs to simulate humans and generate synthetic research data. Hämäläinen et al. [29] found that GPT-3 was able to produce realistic qualitative data, substantially overlapping with data generated by humans and even producing additional insights. LLM-simulated simulated user interactions were found to be realistic, and useful for assessing the effects of design choices during prototyping of social media platforms [30]. Xiang et al. [16] used an LLM to simulate a user interacting with an interface, and found it was useful for identifying edge-cases, generating a substantial amount of usability feedback not identified by human users. These studies demonstrate simulated users can be leveraged for qualitative insights (**R4**), both by allowing designers to inspect interactions and by directly generating feedback.

An application of user simulations with especially much potential is in the evaluation of conversational agents. Simulated users do not rely on static benchmark datasets but can be interactively prompted, and can produce realistic, believable dialogue [17], [31] (**R1**). De Wit [17] found that researchers were able to identify several usability issues in a conversational agent by leveraging simulated user interactions.

When instructed to impersonate a persona, or description of a user, simulated users can align their manner of interaction with the persona’s characteristics [16], [31]. When employing simulated users with a broad range of personas, designers can support diversity by identifying inclusivity bugs [30] (**R2**).

There has been little research on producing quantitative insights through user simulations. De Wit [17] proposes having simulated users generate synthetic survey data to measure user experience of a conversational agent. ChatGPT has been able to replicate outcomes of some quantitative SE surveys, but this has not been attempted for user studies of SE tools [32]. As of yet, it is unclear whether user simulations can provide reliable quantitative metrics, and if so, how (**R3**).

### B. LLM-as-a-Judge

Zheng et al. [19] coined the term “LLM-as-a-Judge”, referring to approaches where LLMs are used as a “judge” to evaluate LLM-generated texts. These approaches come in several forms including grading a single LLM output, or selecting the best of a pair of outputs [19]. These judgments can be made based on a set of criteria [13], meaning LLM-as-a-Judge can be employed to provide a wide range of quantitative metrics (**R3**). The evaluation provided by LLM-as-a-Judge approaches often align well with human judgments [19], including for various SE tasks [20]. In addition to their high scalability and low cost compared to human evaluation, this has led to LLM-as-a-Judge being increasingly employed in practice for evaluating LLM-based tools [4].

Furthermore, judge LLMs can be instructed to provide an explanation for their assessment [19]. However, there is a lack of research on the usefulness of these explanations, and whether they can be used to pinpoint usability issues, provide feedback, and thus generate qualitative insights (**R4**).

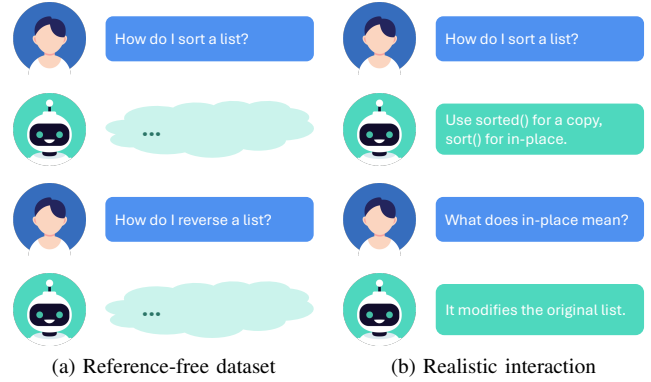


Fig. 2. Reference-free datasets miss conversations where responses refer to earlier messages.

Although LLM-as-a-Judge approaches can be used to evaluate outputs without a reference, datasets of inputs are still needed to serve as test cases. When evaluating conversational agents using multi-turn datasets, the lack of reference outputs makes it impossible to anticipate all potential interactions, as demonstrated in Figure 2. These reference-free datasets are likely not representative of real users’ interactions (**R1**), who might ask follow-up questions an SE assistants’ earlier response. Besides, constructing domain- and tool-specific input datasets for LLM-as-a-Judge may cost substantial effort. When designing an LLM-based SE assistant, changing or adding features would involve updating this dataset to accommodate for new interactions, slowing down the prototyping process.

Through fine-tuning on human feedback, LLMs have been aligned with the preferences of an “average user” [33]. While input datasets can be constructed to reflect the interactions of a diverse group of users, LLMs employed as judges will only evaluate responses from the point of view of this average user. However, LLMs can predict individuals’ preferences regarding natural language when provided personal context in the form of personas [34]. Providing personas is effective in LLM-as-a-Judge settings too, allowing judge LLMs to achieve performance comparable to human annotators in predicting individuals’ assessments [35]. Therefore, LLM-as-a-Judge approaches can support diversity when leveraging representative interaction datasets and user personas (**R2**).

## IV. TOWARDS A SOLUTION

### A. Proposed method

We highlight that the two discussed methods complement each other well. Bridging the gaps between HCI and AI research by combining the two methods can fulfill the discussed requirements and provide automatic human-centered evaluation of LLM-based conversational assistants for SE. We are the first to propose combining simulated users and LLM-as-a-Judge. Although providing LLM-as-a-Judge with a persona could arguably be seen as analogous to having a simulated user give quantitative assessments, the novelty in the proposed method lies in its framing that enables the transfer of practical insights between the fields of HCI and AI.

Employing the proposed combined method, one should start by creating a representative set of personas and relevant evaluation criteria, both quantitative and qualitative. As designers using LLM-as-a-Judge approaches have been found to adapt and add evaluation criteria in response to new insights [13], we imagine the set of personas may evolve similarly throughout the design process. A round of evaluation using the proposed method would require the following steps, for each persona:

- 1) Generate interactions between the assistant and a simulated user leveraging the persona.
- 2) Generate qualitative feedback on the interaction with a simulated user leveraging the persona.
- 3) Employ LLM-as-a-Judge for a quantitative assessment based on the persona and evaluation criteria.

Note that we do not intend for this method to replace expert analysis or user studies. Rather, we expect it will complement the design process by enabling quick evaluation during prototyping and increasing the scope of evaluation with real users. Repeated comparison of the proposed method with pilot user studies is necessary to ensure that the personas, simulated interactions, and evaluation align with real users.

### B. Challenges

We identify several challenges that must be addressed before successfully applying the proposed evaluation method to conversational SE assistants.

*C1) Create representative personas:* Simulated users will need to be representative of the intended end users. Personas described in the literature may serve as a starting point for these simulations. The GenderMag method [21] describes personas based on several facets of using problem-solving tools, and a study at Microsoft [36] characterized software engineers' personas based on differences in tasks, collaboration styles, and perspectives on autonomy. More research into individuals' intents, interaction styles, and preferences regarding conversational SE assistants is needed to extend and adapt these personas to the context of conversational SE assistants, increasing the realism of simulated users' interactions and evaluations through LLM-as-a-Judge.

*C2) Identify and mitigate biases:* LLMs have been found to produce biased and stereotyped outputs when simulating humans [37], [38] or when used in LLM-as-a-Judge approaches [19]. There are techniques to mitigate such biases [19], [37], [39], and biased user simulations may still prove useful in increasing the scale and coverage of evaluation as long as they do not replace real user studies [39]. However, the proposed method requires thorough testing to mitigate potential biases and ensure that its effects are not harmful.

*C3) Handle context:* There is much potential in conversational SE assistants providing tailored support by leveraging local context such as code repositories, documentation, and information about the current task [5], [40]. Techniques such as Retrieval-Augmented Generation (RAG) have extensively been researched to provide relevant context to LLMs while handling limitations on input length [41]. Although there is a wide range of metrics to evaluate RAG [41], more research is

needed on how these can be integrated with user simulations and LLM-as-a-Judge approaches.

### V. CONCLUDING DISCUSSION

Besides the discussed challenges, more research is needed regarding implementation choices such as prompt design, temperature settings, and the number and length of conversations. However, we expect existing research on simulated users [16], [31] and LLM-as-a-Judge [19] to provide useful guidance on these implementation choices.

Although employing an LLM for evaluation of another LLM may seem counterintuitive, existing literature provides evidence for the effectiveness of this approach [19]. However, the performance of LLM-as-a-Judge approaches varies based on the SE task being evaluated [20], and simulated users occasionally suffer from hallucination [16] and may not fully capture the cognitive processes of behavior and cognition [17]. To assess the viability of the combined method proposed in this paper and explore the potential of LLM-based evaluation in SE, we intend to employ it during the design of an LLM-based conversational SE assistant, and compare the synthesized results to those of a user study.

It should be noted that the goal is not to replace the human aspect by relying only on automatic methods for evaluation. Instead, the proposed approach is intended to complement manual human-centered methods by extending the scale, coverage, and frequency of evaluation. By allowing for comparison and alignment with human judgments, we consider this method to be an automatic human-centered approach. Still, it must be extensively tested and employed carefully to uphold ethical considerations and human-centered priorities.

To conclude, we have identified several requirements for automatic human-centered evaluation of LLM-based conversational assistants for SE and discuss whether simulated users and LLM-as-a-Judge address these requirements. We proposed a method combining both approaches and identified challenges down the road. In this way, we contribute to aligning future SE assistants with developers' needs.

### REFERENCES

- [1] L. Erlenhov, F. Gomes de Oliveira Neto, R. Scandariato, and P. Leitner, "Current and Future Bots in Software Development," in *2019 IEEE/ACM 1st International Workshop on Bots in Software Engineering (BotSE)*, May 2019, pp. 7–11.
- [2] 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," *ACM Transactions on Software Engineering and Methodology*, p. 3695988, Sep. 2024.
- [3] M. Wessel, A. Abdellatif, I. Wiese, T. Conte, E. Shihab, M. A. Gerosa, and I. Steinmacher, "Bots for pull requests: the good, the bad, and the promising," in *Proceedings of the 44th International Conference on Software Engineering*. Pittsburgh Pennsylvania: ACM, May 2022, pp. 274–286.
- [4] H. Li, C.-P. Bezemer, and A. E. Hassan, "Software Engineering and Foundation Models: Insights from Industry Blogs Using a Jury of Foundation Models," Oct. 2024, arXiv:2410.09012.
- [5] 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.

- [6] 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.
- [7] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le, and C. Sutton, "Program Synthesis with Large Language Models," Aug. 2021, arXiv:2108.07732 [cs].
- [8] R. Choudhuri, D. Liu, I. Steinmacher, M. Gerosa, and A. Sarma, "How Far Are We? The Triumphs and Trials of Generative AI in Learning Software Engineering," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. Lisbon Portugal: ACM, Apr. 2024, pp. 1–13.
- [9] D. Lo, N. Nagappan, and T. Zimmermann, "How practitioners perceive the relevance of software engineering research," in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2015. New York, NY, USA: Association for Computing Machinery, Aug. 2015, pp. 415–425.
- [10] B. A. Myers, A. J. Ko, T. D. LaToza, and Y. Yoon, "Programmers Are Users Too: Human-Centered Methods for Improving Programming Tools," *Computer*, vol. 49, no. 7, pp. 44–52, Jul. 2016.
- [11] M. C. Davis, E. Aghayi, T. D. Latoza, X. Wang, B. A. Myers, and J. Sunshine, "What's (Not) Working in Programmer User Studies?" *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 5, pp. 120:1–120:32, Jul. 2023.
- [12] M. Allouch, A. Azaria, and R. Azoulay, "Conversational Agents: Goals, Technologies, Vision and Challenges," *Sensors*, vol. 21, no. 24, p. 8448, Jan. 2021.
- [13] Q. Pan, Z. Ashktorab, M. Desmond, M. Santillán Cooper, J. Johnson, R. Nair, E. Daly, and W. Geyer, "Human-Centered Design Recommendations for LLM-as-a-judge," in *Proceedings of the 1st Human-Centered Large Language Modeling Workshop*, N. Soni, L. Flek, A. Sharma, D. Yang, S. Hooker, and H. A. Schwartz, Eds. TBD: ACL, Aug. 2024, pp. 16–29.
- [14] E. Jiang, K. Olson, E. Toh, A. Molina, A. Donsbach, M. Terry, and C. J. Cai, "PromptMaker: Prompt-based Prototyping with Large Language Models," in *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. New Orleans LA USA: ACM, Apr. 2022, pp. 1–8.
- [15] C.-W. Liu, R. Lowe, I. Serban, M. Noseworthy, L. Charlin, and J. Pineau, "How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, J. Su, K. Duh, and X. Carreras, Eds. Austin, Texas: Association for Computational Linguistics, Nov. 2016, pp. 2122–2132.
- [16] W. Xiang, H. Zhu, S. Lou, X. Chen, Z. Pan, Y. Jin, S. Chen, and L. Sun, "SimUser: Generating Usability Feedback by Simulating Various Users Interacting with Mobile Applications," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Honolulu HI USA: ACM, May 2024, pp. 1–17.
- [17] J. de Wit, "Leveraging Large Language Models as Simulated Users for Initial, Low-Cost Evaluations of Designed Conversations," in *Chatbot Research and Design*, A. Følstad, T. Araujo, S. Papadopoulos, E. L.-C. Law, E. Luger, M. Goodwin, S. Hobert, and P. B. Brandtzaeg, Eds. Cham: Springer Nature Switzerland, 2024, pp. 77–93.
- [18] M. Gerosa, B. Trinkenreich, I. Steinmacher, and A. Sarma, "Can AI serve as a substitute for human subjects in software engineering research?" *Automated Software Engineering*, vol. 31, no. 1, p. 13, Jan. 2024.
- [19] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, and I. Stoica, "Judging LLM-as-a-judge with MT-bench and Chatbot Arena," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS '23. Red Hook, NY, USA: Curran Associates Inc., May 2024, pp. 46 595–46 623.
- [20] T. Ahmed, P. Devanbu, C. Treude, and M. Pradel, "Can LLMs Replace Manual Annotation of Software Engineering Artifacts?" Aug. 2024, arXiv:2408.05534.
- [21] M. Burnett, S. Stumpf, J. Macbeth, S. Makri, L. Beckwith, I. Kwan, A. Peters, and W. Jernigan, "GenderMag: A Method for Evaluating Software's Gender Inclusiveness," *Interacting with Computers*, vol. 28, no. 6, pp. 760–787, Nov. 2016.
- [22] M. Guizani, I. Steinmacher, J. Emard, A. Fallatah, M. Burnett, and A. Sarma, "How to debug inclusivity bugs?: a debugging process with information architecture," in *Proceedings of the 2022 ACM/IEEE 44th International Conference on Software Engineering: Software Engineering in Society*. Pittsburgh Pennsylvania: ACM, May 2022, pp. 90–101.
- [23] S. Amershi, D. Weld, M. Vorvoreanu, A. Fournay, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, 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.
- [24] 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.
- [25] J. Richards and M. Wessel, "What You Need is What You Get: Theory of Mind for an LLM-Based Code Understanding Assistant," in *2024 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2024, pp. 666–671.
- [26] Y. Dong, R. Mu, Y. Zhang, S. Sun, T. Zhang, C. Wu, G. Jin, Y. Qi, J. Hu, J. Meng, S. Bensalem, and X. Huang, "Safeguarding Large Language Models: A Survey," Jun. 2024, arXiv:2406.02622.
- [27] P. Thomas, S. Spielman, N. Craswell, and B. Mitra, "Large Language Models can Accurately Predict Searcher Preferences," in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Washington DC USA: ACM, Jul. 2024, pp. 1930–1940.
- [28] M. Sadek, R. A. Calvo, and C. Mougnot, "Trends, Challenges and Processes in Conversational Agent Design: Exploring Practitioners' Views through Semi-Structured Interviews," in *Proceedings of the 5th International Conference on Conversational User Interfaces*. Eindhoven Netherlands: ACM, Jul. 2023, pp. 1–10.
- [29] P. Hämäläinen, M. Tavast, and A. Kunnari, "Evaluating Large Language Models in Generating Synthetic HCI Research Data: a Case Study," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Hamburg Germany: ACM, Apr. 2023, pp. 1–19.
- [30] P. Törnberg, D. Valeeva, J. Uitermark, and C. Bail, "Simulating Social Media Using Large Language Models to Evaluate Alternative News Feed Algorithms," Oct. 2023, arXiv:2310.05984 [cs].
- [31] H. Jin, M. Yoo, J. Park, Y. Lee, X. Wang, and J. Kim, "TeachTune: Reviewing Pedagogical Agents Against Diverse Student Profiles with Simulated Students," Oct. 2024, arXiv:2410.04078.
- [32] I. Steinmacher, J. M. Penney, K. R. Felizardo, A. F. Garcia, and M. A. Gerosa, "Can ChatGPT emulate humans in software engineering surveys?" in *Proceedings of the 18th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ser. ESEM '24. New York, NY, USA: Association for Computing Machinery, Oct. 2024, pp. 414–419.
- [33] H. Wang, Y. Lin, W. Xiong, R. Yang, S. Diao, S. Qiu, H. Zhao, and T. Zhang, "Arithmetic Control of LLMs for Diverse User Preferences: Directional Preference Alignment with Multi-Objective Rewards," Mar. 2024, arXiv:2402.18571 [cs, stat].
- [34] T. Hu and N. Collier, "Quantifying the Persona Effect in LLM Simulations," Jun. 2024, arXiv:2402.10811 [cs].
- [35] Y. R. Dong, T. Hu, and N. Collier, "Can LLM be a Personalized Judge?" Jun. 2024, arXiv:2406.11657 [cs].
- [36] D. Ford, T. Zimmermann, C. Bird, and N. Nagappan, "Characterizing Software Engineering Work with Personas Based on Knowledge Worker Actions," in *2017 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, Nov. 2017, pp. 394–403.
- [37] L. P. Argyle, E. C. Busby, N. Fulda, J. R. Gubler, C. Rytting, and D. Wingate, "Out of One, Many: Using Language Models to Simulate Human Samples," *Political Analysis*, vol. 31, no. 3, pp. 337–351, Jul. 2023.
- [38] A. Taubenfeld, Y. Dover, R. Reichart, and A. Goldstein, "Systematic Biases in LLM Simulations of Debates," Feb. 2024, arXiv:2402.04049 [cs].
- [39] A. Wang, J. Morgenstern, and J. P. Dickerson, "Large language models cannot replace human participants because they cannot portray identity groups," Feb. 2024, arXiv:2402.01908 [cs].
- [40] G. Melo, L. F. Lins, P. Alencar, and D. Cowan, "Supporting Contextual Conversational Agent-Based Software Development," in *2023 IEEE/ACM 5th International Workshop on Bots in Software Engineering (BotSE)*, May 2023, pp. 9–13.
- [41] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, M. Wang, and H. Wang, "Retrieval-Augmented Generation for Large Language Models: A Survey," Mar. 2024, arXiv:2312.10997 [cs].