



Shaping the Emerging Norms of Using Large Language Models in Social Computing Research

Hong Shen
hongs@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, PA, United States

Tianshi Li
tianshil@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, PA, United States

Toby Jia-Jun Li
toby.j.li@nd.edu
University of Notre Dame
Notre Dame, IN, United States

Joon Sung Park
joonspk@stanford.edu
Stanford University
Stanford, CA, United States

Di Yi Yang
diyi@cs.stanford.edu
Stanford University
Stanford, CA, United States

ABSTRACT

The emergence of Large Language Models (LLMs) has brought both excitement and concerns to social computing research. On the one hand, LLMs offer unprecedented capabilities in analyzing vast amounts of textual data and generating human-like responses, enabling researchers to delve into complex social phenomena. On the other hand, concerns are emerging regarding the validity, privacy, and ethics of the research when LLMs are involved. This SIG aims at offering an open space for social computing researchers who are interested in understanding the impacts of LLMs to discuss their current practices, perspectives, challenges when engaging with LLMs in their everyday work and collectively shaping the emerging norms of using LLMs in social computing research.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing theory, concepts and paradigms.**

KEYWORDS

Large Language Models, Research Methods, Validity, Privacy, Ethics, Social Computing

ACM Reference Format:

Hong Shen, Tianshi Li, Toby Jia-Jun Li, Joon Sung Park, and Di Yi Yang. 2023. Shaping the Emerging Norms of Using Large Language Models in Social Computing Research. In *Computer Supported Cooperative Work and Social Computing (CSCW '23 Companion)*, October 14–18, 2023, Minneapolis, MN, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3584931.3606955>

1 BACKGROUND

The development of large language models (LLMs) such as ChatGPT has brought both excitement and concerns to the field of social computing. On the one hand, LLMs present important opportunities that leverage the vast amount of human behavioral data

captured in the model [13] to analyze and augment interactions in social computing systems. For instance, LLM-powered tools have been shown to assist researchers to more efficiently analyze textual data [17, 18], replicate social science studies [3], and enable new ways to prototype emergent social dynamics in social computing systems where it is infeasible or dangerous to conduct in the wild studies [14].

On the other hand, concerns about applying LLMs to social computing research [2] that parallel previous discussions on privacy and consent in computational social science research [8, 15] have also arisen. In particular, we focus on three key themes: validity, privacy and ethics. Can we ensure the validity of findings generated by a non-deterministic black-box model whose output may depend on the nuances captured in the input prompts? Can we protect the privacy of the subjects whose data maybe captured in the models' training data? And can we put in place proper guardrails that will encourage ethical application of LLMs in social computing research in the face of their risks, including but not limited to their potential misuse of LLMs for manipulation or deception?

The primary goal of this Special Interest Group (SIG) is to provide an inclusive platform for social computing researchers who wish to explore the implications of LLMs in their day-to-day research activities. In particular, we aim to explore the following questions:

- How to better design online studies to effectively prevent LLM-based spams? For example, how to prevent, recognize and filter out spammers who use LLMs to fill out online surveys?
- How to utilize LLMs to analyze human-generated data and effectively evaluate their performance? For example, how to construct ground truth when using LLMs in analyzing qualitative data (e.g., open-ended survey responses)? What are the evaluation metrics should we use (e.g., shall we calculate inter-coder reliability)?
- How to accurately document the utilization of LLMs in research while simultaneously acknowledging its inherent limitations and biases in an effective manner?
- How to preserve the privacy of the study data when we use LLMs in data analysis? For example, how to effectively remove Personal Identifiable Information (PII) from survey responses, interview transcripts, and data scraped from the web?



This work is licensed under a Creative Commons Attribution International 4.0 License.

CSCW '23 Companion, October 14–18, 2023, Minneapolis, MN, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0129-0/23/10.
<https://doi.org/10.1145/3584931.3606955>

- How to address ethical concerns related with using LLMs in social computing? For example, how to craft informed consent to inform study participants the potential of using LLMs in the study design? How to ethically use dataset containing real-world LLMs usage data in research?
- How to mitigate the equity concerns associated with the substantial cost, computational resources, and technical expertise required for employing LLMs, considering the unequal access to these resources among different research teams?

In an attempt to ground the above questions in a more concrete manner, below we discuss the impacts of LLMs in social computing research on the following four aspects: Data collection, data generation, data analysis as well as system deployment and evaluation.

2 THE IMPACT OF LLMS ON DATA COLLECTION

Social computing researchers are already experiencing the impacts of LLMs in their work. One particular area is during the data collection stage. On the one hand, the rapidly advancing capabilities of LLMs to mimic human behaviors have presented numerous promising opportunities. For instance, researchers can leverage LLMs to generate hypothetical scenarios (e.g., vignettes) to collect data from human participants. Additionally, LLM-based agents can also be introduced into multiplayer games, opening up new avenues for studying human behavior in interactive settings [2]. However, these capabilities also give rise to a variety of concerns. For example, researchers have shown that chatbots powered by LLMs can effectively mimic survey respondents with diverse backgrounds [1]. This poses a challenge for researchers engaged in online studies, as it becomes increasingly difficult to differentiate between AI-based spammers/bots from genuine human participants. Traditional methods used to identify and filter out spammers may no longer be effective in this context. Moreover, the introduction of LLMs in different parts of the research design, including using LLMs to analyze human-generated data and/or using LLMs to simulate human behavior, will also likely need an update on the consent process. What is the best practices to craft informed consent with human participants when LLMs is involved?

3 THE USE OF LLMS IN GENERATING SYNTHETIC DATA

LLMs capture the human behaviors that are represented in their training data [3, 14], and as such, these models can replicate these behaviors when prompted. Recent studies have shown that the human behavior generated by these models is qualitatively believable [13] and, at times, accurate enough to replicate some social science studies [18] and surveys [1]. This capacity for the model to generate human behavior offers opportunities to enable new ways of studying and augmenting social computing systems. For instance, these models can allow the designers of a social system to prototype the social dynamics that only emerge at scale, to iterate without exposing the users to potentially flawed system design. They can also bring about new ways of conducting computational social science by replicating results that were only achievable via crowd participants or empirical studies. However, the applications of LLMs inherit the imperfections and biases of the underlying model. Their

output might depend on the subtle nuances of a prompt, while their biases might misrepresent certain populations. We posit that our community will need to continuously validate and benchmark the use of LLMs in social computing while emphasizing the importance of directly connecting with human stakeholders.

4 THE USE OF LLMS IN ANALYZING DATA

At the data analysis stage of social computing research, we are witnessing emerging use of LLMs (and recently, *large* pre-trained ones) in both qualitative and quantitative methods.

In qualitative research, academic tools such as PaTAT [6] and CollabCoder [5] as well as commercial products such as the AI Coding feature in ATLAS.ti¹ utilize LLMs to aid in the qualitative coding process for textual data. Generally, these tools employ LLMs to analyze data, propose new codes for inductive coding procedures, understand the semantics of codes as users create and assign them, and suggest codes for data items. They also assist with the sensemaking process by summarizing, synthesizing, aggregating, or visualizing data.

For quantitative methods, AI and LLM-enabled assistants and “pair programmers” have been developed to recommend analyses and statistical procedures based on data characteristics, conduct data queries, and implement data analysis code in response to user prompts in exploratory data science scenarios [10, 12, 16]. Commercial products such as Tableau AI² suggest metrics based on the data domain, generate insights, and allow users to ask questions about the underlying data using natural language.

Despite their potential in improving the efficiency of the analysis process, facilitating additional insight discovery, and reducing learning curves, the use of LLMs also presents new challenges and raises concerns regarding their application in social computing research. For instance, LLMs have been found to display biases and stereotypes in their outputs [9, 11], which could influence the data analysis process, especially in domains where understanding the socio-cultural context is essential for data interpretation. Moreover, the collaboration between human researchers and LLM-enabled tools in analysis tasks poses challenges in ensuring user autonomy, preventing over-reliance, and promoting effective human learning about data and patterns, this challenge can be amplified by the lack of interpretability in LLMs [6]. The application of LLMs in data analysis introduces additional data privacy challenges since many LLMs lack transparency regarding their usage of user data. This raises questions about creating informed consent protocols to notify participants about the use of LLMs in analyzing their data. Furthermore, the community needs new guidelines and norms regarding evaluation metrics when LLMs are used alongside human coders in data analysis.

5 THE METHODS FOR DEPLOYING/STUDYING LLM-ENABLED SOCIO-TECHNICAL SYSTEMS

The success of LLMs is going to have a profound impact on socio-technical systems in various domains that humans use natural languages to interact with. Over the past few months, news articles

¹<https://atlasti.com/ai-coding>

²<https://www.tableau.com/solutions/ai-analytics>

have covered systems built with LLMs used for mental health support³, education⁴, legal services⁵, job searching advice⁶, and many other purposes. While they demonstrate exciting opportunities of advancing these fields, concerns and backlashes have been raised by the public. For example, in January 2023, a company called Koko that offers mental health services tested responses generated by GPT-3 on thousands of its users. A co-founder tweeted about their experiments and it soon sparked a heated discussion around the ethics of this research. People questioned their informed consent process, the legitimacy of testing an unproven technology on real users, and even the appropriateness of involving AI in such a process at all. Although the co-founder later clarified that Koko users knew the messages were co-written by a bot, it did not resolve all these concerns.

Researching systems built with LLMs is a delicate process. However, the lack of clear guidelines for conducting research in this field will affect both researchers who design, develop, and deploy socio-technical systems powered by LLMs, and researchers who conduct empirical studies to investigate how real-world users interact with these systems. In this SIG, we aim to take the first step towards the development of the guidelines. The questions that need in-depth discussions include but are not limited to the following. For researchers who aim to deploy a novel LLM-enabled system, how should they determine whether AI-based intervention is appropriate for the selected use case? How should they disclose the use of LLMs to their users? Deploying certain services (e.g., mental health support) may inevitably lead the users to expose sensitive information about themselves and other people (i.e., interdependent privacy [7]). How should researchers process traces from the study that may involve such sensitive information? Relatedly, the natural language interfaces give users a great amount of flexibility, which means the LLM-enabled services may not have definite use scenarios (e.g., ChatGPT) or the users may use them in unexpected ways. Hence, researchers who want to study the use of LLM-enabled systems are facing challenges in handling unexpected privacy harms (e.g., economical, reputational, psychological harms [4]), which may affect the choice of tools for analysis (e.g., local vs. cloud-based tools).

6 CONCLUSION

The rapid development of Large Language Models (LLMs) has already had a significant impact on various aspects of social computing research, encompassing areas such as data collection, data generation, data analysis, as well as system deployment and evaluation. However, alongside the excitement surrounding these advancements, concerns have emerged regarding issues of validity, privacy, and ethics. This Special Interest Group (SIG) aims to provide a much-needed space for researchers who are interested in comprehending the impacts of LLMs on their work. It offers an opportunity for the members of the community to openly discuss

their current practices, perspectives, and challenges when engaging with LLMs in their day-to-day activities and collectively shaping the emerging norms of LLMs-impacted social computing research.

REFERENCES

- [1] Lisa P Argyle, Ethan Busby, Joshua Gubler, Chris Bail, Thomas Howe, Christopher Rytting, and David Wingate. 2023. AI Chat Assistants can Improve Conversations about Divisive Topics. *arXiv preprint arXiv:2302.07268* (2023).
- [2] Christopher A Bail. 2023. Can Generative AI Improve Social Science? (2023). <https://osf.io/preprints/socarxiv/rwtzs>
- [3] Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences* 120, 6 (feb 2023). <https://doi.org/10.1073/pnas.2218523120>
- [4] Danielle Keats Citron and Daniel J Solove. 2022. Privacy harms. *BUL Rev.* 102 (2022), 793.
- [5] Jie Gao, Yuchen Guo, Gionnieve Lim, Tianqin Zhan, Zheng Zhang, Toby Jia-Jun Li, and Simon Tangi Perrault. 2023. CollabCoder: A GPT-Powered Workflow for Collaborative Qualitative Analysis. *arXiv preprint arXiv:2304.07366* (2023).
- [6] Simret Araya Gebreegziabher, Zheng Zhang, Xiaohang Tang, Yihao Meng, Elena L. Glassman, and Toby Jia-Jun Li. 2023. PaTAT: Human-AI Collaborative Qualitative Coding with Explainable Interactive Rule Synthesis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 362, 19 pages. <https://doi.org/10.1145/3544548.3581352>
- [7] Mathias Humbert, Benjamin Trubert, and Kévin Huguenin. 2019. A survey on interdependent privacy. *ACM Computing Surveys (CSUR)* 52, 6 (2019), 1–40.
- [8] Kevin Lewis, Jason Kaufman, Marco Gonzalez, Andreas Wimmer, and Nicholas Christakis. 2008. Tastes, ties, and time: A new social network dataset using Facebook. *com. Social networks* 30, 4 (2008), 330–342.
- [9] Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards understanding and mitigating social biases in language models. In *International Conference on Machine Learning*. PMLR, 6565–6576.
- [10] Andrew M McNutt, Chenglong Wang, Robert A Deline, and Steven M. Drucker. 2023. On the Design of AI-Powered Code Assistants for Notebooks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 434, 16 pages. <https://doi.org/10.1145/3544548.3580940>
- [11] Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 5356–5371. <https://doi.org/10.18653/v1/2021.acl-long.416>
- [12] Zheng Ning, Zheng Zhang, Tianyi Sun, Yuan Tian, Tianyi Zhang, and Toby Jia-Jun Li. 2023. An Empirical Study of Model Errors and User Error Discovery and Repair Strategies in Natural Language Database Queries. In *Proceedings of the 28th International Conference on Intelligent User Interfaces* (Sydney, NSW, Australia) (IUI '23). Association for Computing Machinery, New York, NY, USA, 633–649. <https://doi.org/10.1145/3581641.3584067>
- [13] Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. *arXiv:2304.03442 [cs.HC]*
- [14] Joon Sung Park, Lindsay Popowski, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *In the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3526113.3545616>
- [15] Matthew J Salganik. 2019. *Bit by bit: Social research in the digital age*. Princeton University Press.
- [16] Dakuo Wang, Justin D. Weisz, Michael Muller, Parikshit Ram, Werner Geyer, Casey Dugan, Yla Tausczik, Horst Samulowitz, and Alexander Gray. 2019. Human-AI Collaboration in Data Science: Exploring Data Scientists' Perceptions of Automated AI. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 211 (nov 2019), 24 pages. <https://doi.org/10.1145/3359313>
- [17] Ziang Xiao, Xingdi Yuan, Q Vera Liao, Rania Abdelghani, and Pierre-Yves Oudeyer. 2023. Supporting Qualitative Analysis with Large Language Models: Combining Codebook with GPT-3 for Deductive Coding. In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces*. 75–78.
- [18] Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2023. Can Large Language Models Transform Computational Social Science? *arXiv preprint arXiv:2305.03514* (2023).

³<https://gizmodo.com/mental-health-therapy-app-ai-koko-chatgpt-rob-morris-1849965534>

⁴<https://fortune.com/2023/02/22/chatgpt-ai-openai-educatoin-tutor-teaching-school/>

⁵<https://gizmodo.com/donotpay-speeding-ticket-chatgpt-1849960272>

⁶<https://www.forbes.com/sites/jackkelly/2023/04/03/how-to-leverage-ai-and-use-chatgpt-in-your-job-search-according-to-rsum-writers-and-career-coaches/?sh=728117a5ac5a>