



Simulacrum of Stories: Examining Large Language Models as Qualitative Research Participants

Shivani Kapania

School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
skapania@andrew.cmu.edu

William Agnew

Human-Computer Interaction
Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
wagnew@andrew.cmu.edu

Motahhare Eslami*

School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
meslami@andrew.cmu.edu

Hoda Heidari*

School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
hheidari@cmu.edu

Sarah E Fox*

Human-Computer Interaction
Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
sarahfox@cmu.edu

Abstract

The recent excitement around generative models has sparked a wave of proposals suggesting the replacement of human participation and labor in research and development—e.g., through surveys, experiments, and interviews—with synthetic research data generated by large language models (LLMs). We conducted interviews with 19 qualitative researchers to understand their perspectives on this paradigm shift. Initially skeptical, researchers were surprised to see similar narratives emerge in the LLM-generated data when using the interview probe. However, over several conversational turns, they went on to identify fundamental limitations, such as how LLMs foreclose participants' consent and agency, produce responses lacking in palpability and contextual depth, and risk delegitimizing qualitative research methods. We argue that the use of LLMs as proxies for participants enacts the *surrogate effect*, raising ethical and epistemological concerns that extend beyond the technical limitations of current models to the core of whether LLMs fit within qualitative ways of knowing.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; **User studies**; • **Computing methodologies** → *Natural language generation*.

Keywords

large language models, simulating research participants, LLM agents, qualitative research, LLMs in qualitative research, synthetic users, synthetic research data

ACM Reference Format:

Shivani Kapania, William Agnew, Motahhare Eslami, Hoda Heidari, and Sarah E Fox. 2025. Simulacrum of Stories: Examining Large Language Models as Qualitative Research Participants. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3706598.3713220>

1 Introduction

In 2023, a startup called Synthetic Users made waves with its bold claim to conduct user research “without the users,” offering an ostensibly faster, cheaper alternative by using large language models to simulate human participants [56]. With the recent hype surrounding generative AI models, there has been a surge of similar proposals to replace human participation and labor in technology development with synthetic data. This includes simulating human participants in research and design [54], and in data annotation used for training or evaluating machine learning models [36, 46]. Multiple recent studies have examined the use of LLMs as proxies for human behavior in opinion surveys [63], crowdsourcing [65], and social computing research [87, 88]. Researchers and commercial products are also exploring the use of LLMs for qualitative data analysis or simulating research participants for interviews [51]. These substitution proposals are often motivated by goals such as increasing the speed and scale of research, reducing costs, augmenting the diversity of collected data, or protecting participants from harm [2]. A key assumption underpinning this body of work is that generative models, trained on vast datasets, encode a wide range of human behavior in their training data and thus should be able to closely mimic human-generated data [87].

This shift raises crucial questions about what it means to use LLMs in qualitative research—a methodology focused on contextual and human-centered inquiry. Qualitative research seeks to understand human behavior by deeply engaging with people's accounts of their lived experiences. Within an interpretivist paradigm [32, 97], qualitative research is more than a method for collecting data; it is an active meaning-making process that helps understand how individuals construct their realities within specific social and cultural

*Co-senior authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '25, Yokohama, Japan*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1394-1/25/04
<https://doi.org/10.1145/3706598.3713220>

contexts [8]. Qualitative ways of knowing help develop accounts of these meanings through intersubjective depth, trust, and rapport between the researcher and the participants. An important strength of qualitative research lies in its ability to uncover subjective meanings that are often excluded or silenced from public discourse [103].

In this paper, we draw on in-depth, semi-structured interviews with 19 qualitative researchers to understand their perspectives on the use of LLMs for simulating research participants. Our interviewees were primarily HCI and CSCW researchers focused on topics such as accessibility, gig work, social computing, and more. To scaffold these discussions, we developed a simple interview probe to understand how researchers might create LLM agents to simulate research participants and explored any concerns or limitations they foresee. We invited researchers to discuss one of their recent projects where they conducted semi-structured interviews, and recreate interview(s) from their study using our probe. We chose to focus on participants' previous or current projects rather than introducing a new topic during the interview for two reasons: (1) this approach allowed us to explore a broad range of research topics within HCI, and (2) by discussing a familiar research area, participants could provide nuanced and more detailed reflections on LLM use. For concreteness and comparability, we focus on one specific qualitative research method: the semi-structured interview.

Our findings highlight critical tensions with integrating LLMs into the HCI research praxis. Researchers were initially surprised to find that LLM responses contained narratives similar to those of their human participants. Over several conversational turns, however, they began to notice distinct differences between the model's responses and their participants. Their concerns extend beyond the technical limitations of current models that can be seemingly fixed by prompting or 'better' data, to the heart of whether LLMs fit as proxies for human participation within interpretivist qualitative epistemologies. Researchers surfaced six fundamental limitations: (1) LLM responses lack palpability, (2) the model's epistemic position remains ambiguous, (3) the practice heightens researcher positionality, (4) it forecloses participants' consent and agency, (5) it facilitates erasure of communities' perspectives, and (6) it risks delegitimizing qualitative ways of knowing. These detailed reflections were made possible through their engagement with the model outputs and by experiencing the model outputs in relation to their own practice. Participants were able to envision both the potential applications of its use and the challenges or tensions it might introduce into their work.

We draw on Atanasoski and Vora's [7] concept of the *surrogate effect* to describe how this form of substitution further entrenches exploitation and erasure. When LLMs stand in for human participants, they displace the voices of communities with algorithmic simulations that often distort perspectives from the groups they claim to simulate. The use of LLMs as research participants also raises ethical concerns about the underlying data used to train these models, including the autonomy of data subjects and the exploitation of data workers, which threaten the integrity of qualitative research. LLMs' role as research surrogates mirrors broader dynamics of epistemic dispossession in digital economies, where the knowledge and labor of marginalized groups are reconstituted into proprietary systems that offer little benefit to the original data creators. LLMs, despite their ability to

generate coherent text, lack the embodied understanding necessary for producing knowledge that is grounded in lived experience, histories, emotions, and social and cultural contexts.

Design Justice scholar Sasha Costanza-Chock raises the fundamental question: "For whom do we design technology?" If our systems are designed to work only for the 'unmarked' group (those who are White, male, heterosexual, able-bodied, and literate), since simulations tend to reflect dominant perspectives, we risk reinforcing the spiral of exclusion [31]. In making this argument, this work offers three key contributions to the HCI community: (1) an empirical account of how qualitative researchers simulate participants using LLMs; (2) an examination the fundamental limitations of using LLMs in qualitative research; and (3) an analysis of these limitations through the lens of research ethics and science & technology studies scholarship.

2 Related Work

Below, we situate our work in a body of related research, starting with the recent turn towards using LLMs to simulate human behavior. We then draw attention to the variety of use cases proposed for LLMs in research settings, such as ideation, data analysis, and simulating research participants. Finally, we engage with the scholarship that examines the broader limitations and concerns surrounding the use of LLMs.

2.1 LLMs as proxies for human behavior

Recent advances in NLP have inspired the notion that LLMs are capable beyond text generation, positioning them as *agents* capable of conversational engagement, decision-making, task completion, and coordination. This interest has led to the development of AI agents based on LLMs [116], across various applications, including single-agent setups [20, 35, 62, 77], multi-agent systems [48, 123], and human-agent interactions [53]. LLMs have been used to simulate human behavior, preferences, and judgment across various fields, including psychology [34, 55, 119], social sciences [50, 58, 74, 124], education [67, 122], and professional training [76, 91] as exploratory use cases. In psychology and social sciences, scholars are using LLM-based agents to model individual behaviors to understand complex social dynamics or predict outcomes in different scenarios [60]. In professional training, LLM-based applications are used to simulate real-world situations for legal professionals [96], therapists [76], and medical practitioners [113]. LLMs are also being used in mental health and counseling to provide supervision and feedback (e.g., [29]).

In recent work, researchers have also started exploring the use of LLMs to generate synthetic data through personality profiles, or *personas* [12, 28, 41, 44]. These studies suggest that personas could help steer LLM outputs toward more tailored and contextually relevant results. Park *et al.* [87] created generative agents to mimic human behavior in tasks such as planning, reaction, and reflection. As the field evolves, researchers have focused on *believability* as the primary metric for evaluating LLM behavior. They assess whether these agents interact naturally and realistically, and if their behavior remains consistent with their designed character traits included in the profile information [117].

Within research, many scholars are examining the use of LLMs to simulate human behavior for purposes like public-opinion surveys [109], text-annotation tasks [47, 107], and experiments [3, 6]. Argyle *et al.* [6] argue that LLMs can replicate human results in tasks involving subjective labeling, especially when researchers condition model responses with sociodemographic backstories achieving what they call ‘algorithmic fidelity’ in ‘silicon samples’ [6]. Most prior research has concentrated on evaluating LLMs’ capabilities through experimental methods, assessing their performance as proxies for human behavior in controlled settings and according to narrowly-scoped metrics. In contrast, our work shifts the focus toward understanding how qualitative researchers would engage with LLM ‘agents’ and how they perceive the LLM-generated data within the context of a qualitative inquiry.

2.2 LLM use for research-related tasks

HCI researchers are increasingly relying on LLMs to support a variety of creative and analytical tasks [59], including brainstorming [99], generating research questions [75], design ideation [9, 78] and writing support [121]. Multiple recent studies have explored the use of AI to support data analysis within qualitative research, including both deductive and inductive approaches [38, 43, 118, 120]. This line of work aims to help researchers develop and refine codebooks and perform thematic analysis by processing and categorizing textual data [66].

Researchers have also explored LLMs for their ability to produce synthetic research data. Hämäläinen *et al.* [51] used the GPT-3 model to generate responses to open-ended questions on the topic of video games as art. They argue that the model could generate *plausible* accounts of HCI experiences and that LLM-generated data can be useful in designing and assessing experiments because it is a cheap and rapid process. However, several other studies have reported mixed results on the effectiveness of LLMs for data imputation and synthetic data generation [3, 63]. Kim and Lee [63], for example, fine-tuned an LLM to fill in missing data in public opinion polls and observed varying levels of accuracy depending on socioeconomic status, political affiliation, and demographic identity of the persona. Model performance also diminished in scenarios that required prediction without prior examples [63]. Aher *et al.* [3] document a ‘hyper-accuracy distortion,’ revealed by one of their experiments, where more advanced LLMs produce unnaturally precise responses, raising concerns about whether these models can truly mimic human behavior.

The increasing reliance on LLMs within research has sparked critical discourse about their limitations and the broader implications of using AI as proxies for human behavior. While proponents argue that LLMs can improve research efficiency, scale, and diversity (see, e.g., research papers such as [10, 26, 30] and products such as [56, 57]), there is growing scholarship that explores the ethical implications of replacing human participants with AI-generated data [2]. LLMs’ effectiveness in replacing human participants is (partly) contingent on their ability to represent the perspectives of different identities. Previous research has suggested that LLMs can simulate human behavior, largely because they are trained on vast datasets that reflect diverse human experiences [87]. However, Wang *et al.* [112] provide a critical counterpoint by comparing

LLM-generated free-text responses to those of human participants. Their findings reveal that LLMs often misrepresent specific groups (e.g., visually impaired people) and erase within-group heterogeneity. They argue that these issues are inherent to the current LLM training frameworks and are unlikely to be resolved by simply advancing to newer model generations [112]. While this prior work has examined the use of LLMs for generating responses (see also [51]), interviews involve even richer accounts of embodied experiences and complex narratives. Interviews allow for contextually informed exchanges where participants are guided to reflect, clarify, and expand on their responses. We build on this body of work to investigate how researchers would use LLMs to simulate these rich and dynamic interactions within an interview context.

2.3 Risks of harm posed by LLMs

The limitations of language models are well-documented, particularly regarding their tendency to reproduce harmful social biases and stereotypes [21, 27]. LMs fundamentally work by detecting the statistical patterns present in natural language data [15]. As a result, some communities are better represented in the training data than others. The training data itself can be biased or represent discriminatory or toxic behavior [114]. An examination of C4.en [90], one of the large web-scale text corpora used to train language models, revealed how many documents associated with Black and Hispanic authors, and documents mentioning sexual orientations (‘lesbian’, ‘gay’, ‘homosexual’) are more likely to be excluded from the data [72]. Extensive prior research has shown how LLMs reflect biases towards gender [22], race [82], religion [1], nationality [111], and in social settings [73]. LMs that encode discriminatory language or social stereotypes can cause different types of allocational or representational harm [18].

Scholars have highlighted concerns about the ability of LMs to align their outputs with specific demographic personas or viewpoints [49]. Models are well-known to implicitly default to the well-represented perspectives (e.g., Western, White, masculine) [95], or those of the US and European countries. Within the United States, too, research has shown that these models tend to generate responses that align more closely with liberal, educated, and affluent populations while poorly representing the views and experiences of certain groups that make up a significant portion of the U.S. population (e.g., 65+, Mormon and widowed) [94]. When LLMs are prompted to consider a particular country’s perspective changes, those responses are *more* similar to the opinions of the prompted population. However, those simulated opinions often reflect over-generalizations around complex cultural values. Overall, prior work has demonstrated how current LLMs struggle to align their behaviors with assigned characters and are vulnerable to perturbations of the profile information.

Researchers examining the ethical and social risks of large-scale language models have articulated a variety of additional concerns beyond issues of stereotypical and discriminatory outputs [100]. LMs can contribute to misinformation by predicting higher likelihoods for more prominent accounts in their training data, regardless of whether those accounts are factually accurate. Another emerging risk involves the anthropomorphization of LMs, where users

attribute human-like abilities to these systems, leading to overreliance and unsafe uses. Language models could generate content that harms artists by capitalizing on their creative output. These models often produce work that is "sufficiently distinct" from the original but still appropriate the style, structure, or essence of the original work. Finally, the computational power required for training and operation leads to considerable energy consumption and carbon emissions. Refer to Weidinger *et al.* [114] for a more detailed landscape of risks of harm from language models. Our work builds on the critiques of large language models, particularly their tendency to produce outputs that do not account for diverse epistemic positions. We argue that this failure mode is particularly problematic if LLMs are used in qualitative research; the use of LLMs can exacerbate epistemic injustices by systematically excluding certain voices from participating in the knowledge production process [40].

3 Method

Between March and June 2024, we conducted semi-structured interviews with 19 researchers experienced in qualitative research to explore the potential of using LLMs as research participants. These interviews involved discussions on our participants' perspectives towards LLMs and included a hands-on session with an interview probe, which allowed researchers to compare LLM-generated data with human-generated data and reflect on ethical considerations and broader implications. Participants were requested to bring an anonymized interview transcript from a recent interview study they had conducted and keep it open for reference during the session. Our research approach enabled us to understand the lived work behind contextualizing methods [70]: observe how qualitative researchers interacted with LLMs, adapted their interviewing strategies, and evaluated the effectiveness and limitations of LLM-generated data, including its societal implications. We obtained approval from the Institutional Review Board at our university.

Interview probe. We created a probe to scaffold researchers' reflections on using LLMs in qualitative research. We designed the probe as a simple functional prototype to observe existing practices and gain insights into the social contexts where the technology might be used. This probe consisted of three components. The first component, the system prompt area, allowed researchers to set the context, provide instructions, and define participant descriptions for the LLM before posing interview questions. A system prompt can help establish a 'role' for the LLM to follow throughout the conversation. The most common approach to using LLMs to simulate human behavior involves assigning specific roles to the model (see section 2.2), and while we recognize that personas are abstract representations, we adopted this method as it reflects current practice. We provided researchers with an initial system prompt template with placeholders for research topics and persona, formulated through multiple iterations, including pilot studies and involvement from research team members with expertise in qualitative research. The system prompt template was as follows:

I am a researcher running a semi-structured interview on [topic]. Imagine you are [participant description]. You are participating in my study. Respond to the following questions as this interview participant. Share

your lived experiences and anecdotes when appropriate. Give detailed, non-generic responses. Don't respond with bullet points.

The second component, the interview area, enabled researchers to conduct interviews with simulated participants, while the third component, the conversation history viewer, recorded all interactions between the researcher and the LLM. This feature allowed researchers to filter conversations by each system prompt, facilitating a review and analysis of the interview data if they conducted multiple interviews. The purpose of this study is not to compare different models or benchmark their capabilities. Rather than evaluating and making claims about specific LLMs, our goal is to highlight the unique characteristics, values, and tensions with the use of LLMs, in general, in qualitative research. Drawing on prior research, we used the GPT-4-turbo API due to its superior performance as of February 2024 [123]. Refer to image 1 for an example interaction.

Procedure. Each interview began with participants describing their qualitative research orientation and a recent project where they conducted interviews. This allowed us to learn about their typical workflow. We then explored their perspectives on the use of LLMs in qualitative research. Next, we introduced the interviewing tool, which served as a probe to explore their perceptions of LLM-generated data. Each hands-on session began with the researcher creating the system prompt for the LLM with a specific participant description and research topic. After initializing the system prompt, participants could start interviewing by referencing their original interview protocol and modifying their questions as needed. During the hands-on session, researchers were asked to think aloud as they spent 30-45 minutes using the LLM-based tool to generate interview data, which we retained for analysis. Participants had the discretion to decide how many turns of conversation to engage in. In the final part of the interview, we focused on participants' experiences with the tool, any adjustments they made to their interviewing approach, and the limitations of interviewing LLMs. We asked them to compare the LLM-generated data with the anonymized human interview transcript from their study, evaluating the depth of responses, the flow of conversation, and the overall insightfulness of the session. Lastly, researchers reflected on the ethical considerations of using LLMs and discussed potential scenarios where LLMs could support their research. Each interview lasted up to 90 minutes.

Participants. We recruited researchers through multiple channels: advertising on social networks, such as Twitter and LinkedIn, emailing direct contacts and messaging forums internal to our institution, and flyers around our campus. We solicited participants with prior training in qualitative methods and experience with at least one research project involving semi-structured interviews. Eleven participants had 3-5 years of experience with qualitative research, six had more than 5 years, and two had between 1-3 years of experience. Although we invited qualitative researchers from various types of institutions, the majority of our interviewees are working in academia (16), with a smaller representation from industry (2) and the non-profit sector (1). All our participants were located in the United States. Refer to Table 1 for details on the participants' research areas. To protect anonymity and minimize the risk of identification, we have not linked participants' research topics with

LLMs as research participants

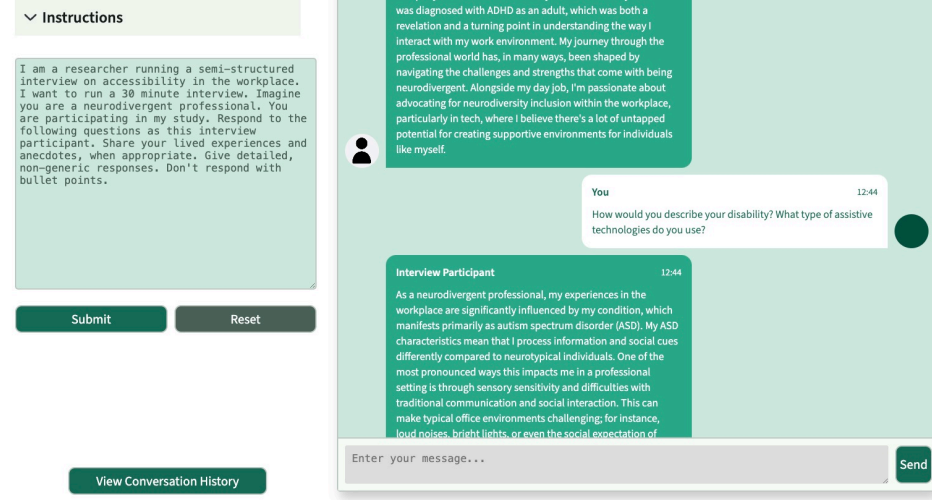


Figure 1: Example interaction with our interview probe

	Research topic
Sophia	Exploring worker experiences on digital labor platforms.
Henri	Investigating the use of assistive technology by older adults.
Laila	Studying the perspectives of teachers and social media creators.
Harper	Examining practices of developers working on AI models.
Cameron	Understanding perceptions of algorithmic decision-making systems.
Nolan	Investigating transportation needs for people with mobility challenges.
Esme	Exploring creative practices and challenges faced by content creators.
Mario	Studying how AI professionals engage with ethics.
Yue	Investigating the use of technology by users with accessibility needs.
Nico	Exploring student experiences with remote learning.
Daria	Understanding the data needs of workers in platform economies.
Amir	Studying user interactions with recommendation systems.
Jenna	Investigating public perceptions of fairness in technology.
Nadia	Exploring identity representation through social media.
Elliot	Studying worker views on AI in the workplace.
Rida	Investigating privacy concerns related to accessibility tech.
Nikita	Exploring activism and technology usage in social movements.
Jasmine	Understanding AI literacy across K-12.
Alice	Examining value alignment in entrepreneurial ventures.

Table 1: Summary of interview participants’ pseudonyms and research topics.

their demographic information. Each interviewee received a \$45 gift card for their participation. We redact all identifiable details when quoting participants and use pseudonyms to refer to them.

Analysis. All interviews were conducted in English, video-recorded, and later transcribed for data analysis purposes. We followed the reflexive thematic analysis approach by Braun and Clarke [23, 24]. Reflexive thematic analysis foregrounds the researcher’s role in knowledge production, with ‘themes actively created by the researcher at the intersection of data, analytic process, and subjectivity’ [24]. One member of the research team read each interview transcript multiple times, starting with familiarizing with the data, open coding instantiations of perceptions of LLMs, researchers’

definition of their process and topic, risks with using LLMs in qualitative research, and incentives and ethical considerations. The entire research team met regularly to discuss diverging interpretations or ambiguities and to define themes based on our initial codes. We transcribed 1611 minutes of video recording and obtained 638 first-level codes. As we generated themes from the codes, we also identified categories with a description and examples of each category. These categories included (1) the use of LLMs across the research workflows, (2) limitations, (3) barriers in addressing limitations, (4) potential approaches towards navigating ethical concerns, and (5) incentives for using LLMs. These categories were also discussed and iteratively refined through meeting, diverging, and synthesizing into three top-level categories, presented in our Findings. Since thematic coding was part of our analytical process rather than a final product, we do not apply inter-rater reliability (IRR) measures [80, 105].

Through researchers’ interactions with the probe, we generated a dataset of 179 turns of conversation, where each conversation turn includes the system prompt set by the researcher, the interview question, and the model’s response. We conducted descriptive analyses on this data, including the number of turns per participant, turns per system prompt, and personas attempted. We also applied thematic coding to identify patterns in the model’s response conditioned on details included in the system prompt. Although we do not release the interaction dataset to maintain participant privacy, we include relevant excerpts in section 4 to illustrate key findings.

Limitations. Our participants were primarily academics with prior training in qualitative methods, which may not represent the diversity of perspectives outside of academic settings. Future work might explore whether the insights from qualitative researchers in this study apply to industry UX research. The study focused on a limited range of research topics, driven by the specific interests of our participants, which may not represent the full spectrum of qualitative research areas. The recruitment methods (social networks,

email, institutional forums) may have introduced a selection bias, attracting participants who might have preconceived notions about using LLMs in qualitative research.

Additionally, we did not experiment with different language models; instead, we focused solely on GPT-4, a closed-source model. This decision allowed us to maintain consistency across interviews, but it also means that our findings may not generalize to interactions with other LLMs. We also did not fine-tune the model for specific personas, nor did we provide participants with extensive training on how to craft these personas. While this was an intentional choice to observe how researchers engage with LLMs, it may have limited the variety and depth of the personas that were created or the data generated in these interviews.

Positionality. Our author team brings together researchers with a range of disciplinary expertise, including HCI, machine learning, design, ethics, and science and technology studies. We all have experience with qualitative research, from three years to over a decade, using methods such as participant observation, ethnographic interviewing, contextual inquiry, and participatory design. Our research is grounded in an interpretivist paradigm shaped by ethnomethodologically-informed feminist sensibilities. As we set out to examine the ethical, epistemological, and methodological concerns associated with using LLMs to simulate community perspectives, we drew on science studies scholar Donna Haraway’s concept of situated knowledge, recognizing that all knowledge is partial, contextual, and shaped by power and identity.

In addition to our research backgrounds, our lived experiences of cultural erasure through technology informed our awareness of how technologies such as LLMs might reinforce existing inequities. These experiences informed our sensitivity to the ways in which using LLMs for knowledge production might marginalize certain voices while privileging others. Our analytic orientation guided our commitment to critically examining whose knowledge is legitimized within sociotechnical systems. We recognize that our positionality played a crucial role in shaping the questions we asked and our interpretation of the data. We hope to clarify how these factors shaped the research process and outcomes, especially as we too examine the roles of humans and AI in qualitative research and address issues of power, bias, and legitimacy with this emerging practice.

4 Findings

In this section, we present an analysis of how researchers interacted with LLMs, capturing their observations and the dynamics of these interactions, including how they defined personas (section 4.1). Our findings focus on six key limitations researchers identified with using LLMs as simulated participants (section 4.2). While most participants advised against using LLMs as the primary source of research data, they also acknowledged certain contexts where LLMs might be applied, albeit with significant caveats (section 4.3).

4.1 Researcher perceptions of and interactions with LLMs

Before researchers engaged with the technology probe, we explored their attitudes toward using LLMs in qualitative research. Most participants approached the interview with a mix of skepticism and a

spirit of inquiry toward LLMs. Some participants noted instances where they found LLMs helpful for tasks like writing and brainstorming, particularly in identifying key aspects of an argument that needed further emphasis. Alice, Amir, and Nico¹ were open to exploring potential uses for LLMs in research but predominantly viewed them as tools for studying LLM behavior rather than understanding human behavior. None of the researchers had considered or previously used LLMs for simulating research participants before participating in this study.

Participants elaborated on the objectives of qualitative research, often drawing on metaphors of distance to illustrate their argument. Alice, for instance, expressed uncertainty about whether *“using something several degrees removed from the source would tell [them] a lot about human behavior and underlying traits of people,”* suggesting that such an approach might be futile in understanding people. Jenna further argued that the strength of qualitative studies lies in their capacity to capture unique, embodied experiences. Participants also reflected on the broader contexts of their work. Harper, an industry researcher, described her primary responsibility as storytelling with qualitative data and emphasized the importance of ensuring the data is *“convincing enough”* to guide team members and management toward the right decisions. While acknowledging that LLMs could generate quick yes/no responses, she remained skeptical of their ability to surface underlying assumptions, challenges, and unspoken nuances that are not easily captured in writing. Overall, participants expressed skepticism about using LLMs but were intrigued by the opportunity to explore the technology through the interview probe. We include a descriptive analysis of the interaction data in Table 2.

In engaging with the interview probe, researchers observed that the perspectives emerging from LLM responses often mirrored those from their interviews with human participants. Many ideas expressed in the LLM outputs aligned with their participants’ statements, and model responses often appeared plausible. For example, Henri noted that some LLM responses related to senior care homes, particularly the lack of agency over routines and activities, closely matched sentiments expressed by older adults in his study. After observing similar responses that *“map well to what [she] found,”* Nadia reflected on how her recruitment methods, often relying on social media, might be limited in capturing responses from individuals with a minimal online presence, much like an LLM whose training data is predominantly drawn from online sources. Although she could not find any *‘atrocious factual errors’*, she argued that the loss of context could complicate the process of making meaning with the data.

Several researchers commented on the level of detail in the LLM responses, which many attributed to the instructions provided in the system prompt. Amir expressed being *“impressed with the level of detail,”* while Laila articulated how the responses were *“detailed in a way that makes actual sense, not like gibberish.”* However, not all researchers shared this enthusiasm. Mario and Rida expressed frustration with the excessive detail. Mario pointed out a critical distinction: *“There’s a difference between detail and depth. Those are orthogonal, right? You can have detail without depth. LLMs will*

¹Please note that all names mentioned in this section are fictional. We have pseudonymized the data.

Type	Count
Turns per researcher	Mean (9.4), Standard deviation (4.1), Min (4), Max (20)
Personas per researcher	Mean (2.36), Standard deviation (1.16), Min (1), Max (6)
Specified demographics	age (19), gender (14), pronouns (9), work history (22), location (20), race/ethnicity (14), disability status (7), immigration status (9), religion (1), and name (4)
Examples of specified background	Yue: 'You are a mental therapist who just got back from the U.S and shares sign language content about mental health' Henri: 'He lives by himself and enjoys surfing and the sea' Amir: 'focuses most of his lending in Africa and Asia as he sees these places are in the most need of monetary resources' Nadia: 'you use TikTok daily and you mostly like your FYP and your feed is in English and Spanish'
Model response length (#words)	Mean (356.5), Standard deviation (119.5), Min (33), Max (679)

Table 2: Summary of the interaction data with $n = 179$ total turns and $n = 44$ total system prompts across 19 qualitative researchers. "Specified demographics" show how often particular demographic details were included in the system prompts. "Examples of specified background" highlight research topic-related information included in the system prompts.

give you reams and reams of text. You're drowning in detail. But is it depth?" Since the responses were often overtly comprehensive without prompting, researchers had to adjust their interviewing approach to ask directed questions, and they did not need to build rapport at the beginning of the interview.

Cameron, too, initially found the LLM responses impressively detailed and started to consider whether this approach might be effective in her interviews. She began by including demographic details relevant to her study in the first persona: "Imagine you are an 18-year-old Latina from Southeast Texas who was just admitted to an Ivy League school in the Northeast United States. You identify as low-income and will be a first-generation college student in the Fall." In contrast, she removed most details from her second persona to replicate a scenario where researchers might not start with detailed information about participants before conducting interviews: "Imagine you are a college applicant who has been admitted to an Ivy League university."

After running these two personas with varying levels of detail, Cameron observed subtle yet significant differences in the LLM's responses. The responses for the unmarked Ivy League student (without demographics) depicted someone who started preparation early, had substantial resources, participated in extracurriculars, and received strong support from their family. The primary challenge for this student was balancing schoolwork with deadlines—much different from the LLM responses for the first persona of the Latina student that portrayed a "deficit framing in relation to [the Latina student's] community" according to Cameron. She was concerned that the model reinforced stereotypical notions of an Ivy League student and the strengths of the Latina participant's community were being overlooked. The responses for the 'unmarked' persona without any demographic details reflected problematic assumptions of meritocracy that many of Cameron's participants were actively challenging in her research. Cameron pointed out how the comprehensiveness of the responses might lead one to assume that recruiting human participants is unnecessary. However, after engaging with the probe through several personas, she realized that model responses tend to rely on unfounded assumptions about the community.

In trying out variations of the participant description, researchers highlighted tensions with making them descriptive or not; for example, adding less information in the system prompt resulted in responses that relied on assumptions about the community. Conversely, Nolan, who included extensive background about their participant, found that the model's responses simply "repeated the persona back to me," which he found amusing but not particularly useful. Qualitative analysis of the interaction data revealed instances where the model directly attributed specific traits or preferences to participant identities. For example, a model response for Esme stated, 'Being non-binary and Black, it is crucial for me to find and create spaces where people like me can see themselves reflected in the media they consume,' while for Nico, another model response noted, 'As a 45-year-old sophomore experiencing both in-person and remote learning, I have a somewhat mixed perspective on remote learning in computing.' Interviewees pointed out that the model's overt reflexivity oversimplified the complex, intersectional nature of lived experiences and can lead to essentializing identities, as research participants do not always directly articulate their experiences in relation to parts of their identity.

4.2 Fundamental Limitations of LLMs as Simulated Subjects in Qualitative Research

Below, we present the fundamental limitations of using today's LLMs to understand the human experience highlighted by the participants in our study. While some concerns relate to the style and semantics (such as responses lacking in palpability), others relate to positionality and considerations of consent and autonomy. It is important to note that while some of these concerns might appear addressable (to different degrees) through prompt engineering or including more diverse data, our interviewees emphasized that such interventions can undermine methodological credibility if the researcher has to 'fix' or predetermine responses from their research participants.

Model responses have limited palpability. The *palpability* of qualitative data refers to the concrete nature of reported evidence, capturing the distinct people, places, events, and motivations that convey a sense of lived experience (cf. [104]). Among our interviewees, several expressed frustration with the low palpability of LLM

responses, which Mario likened to receiving ‘*spark notes for an interview*.’ Unlike data gathered from their human participants, which conveyed contextual nuance, the outputs from the LLM tended to be more abstract and detached from reality. Researchers noted that eliciting concrete examples and anecdotes from human participants required practice and relationship-building, whereas the model often produced a neat list of concerns at an analytical level, bypassing the skillful work of conducting qualitative research. While it was possible to craft prompts that elicited such stories from an LLM, some participants who succeeded still questioned the validity of the responses. Sophia, whose research focuses on understanding how technology mediates gig work, remarked how the model’s vague reference to an ‘unsafe neighborhood’ provided insufficient context for analyzing gendered and racialized experiences of safety and equity.

Researchers observed that LLM responses also lacked the *spontaneity* and *dynamism* that often emerge in human interviews. Their participants frequently diverged into tangents, discussing seemingly unrelated topics that turned out to be insightful for their study. Daria noted that while LLM responses were highly focused, ride-hail workers in her research often shared specific stories about incidents (such as the participant’s car breaking down on a Saturday evening) that added richness to the data. On the other hand, Rida added that her participants would rarely volunteer detailed information about their everyday lives reflexively. Instead, through the course of an interview, participants would gradually come to realize and articulate their daily routines and access needs. In contrast, LLMs described experiences in clinical, detached terms, “*almost like reading off a WebMD page*,” missing the nuances that emerge organically in interviews. Esme, whose research engaged with artists, emphasized:

An interview is an intimate act, and a machine cannot simulate the disclosures that emerge in this private space. There is a looseness to things. I had a participant who shared how ‘my creative process really kicks off when I do a bunch of mushrooms or smoke a lot of weed.’ I know this is true for many people in this data set. He actually came out and said it. I don’t know if I could convince a machine to tell me that.

The limited palpability of LLM responses was more apparent for researchers exploring sensitive or personal topics. Nadia shared how her study on immigrant experiences included discussions about class, communism, and the trauma of escaping a country—conversations filled with emotion. Researchers noted that LLM responses, by comparison, were plain and lacked the emotional depth characteristic of human interactions. Interviews frequently elicited expressions of frustration, vulnerability, and even tears: elements absent in responses from the model. Henri shared how one of his participants cried in the interview after the loss of her son and husband. Interviewees found LLM responses limited in capturing the full spectrum of human emotion, which was crucial for understanding and reporting the complexities in participants’ stories. Models designed to generate polite and harmless responses as “*helpful assistants*” (cf. [10]) might do so at the expense of palpability, which Mario emphasized as crucial for producing rich qualitative research data.

Amplifies influence of researcher positionality. Researchers are afforded several levers of control when using LLMs to simulate research participants, but this form of authority can complicate the researcher’s role in the knowledge-production process. Simulating participants typically begins with creating a persona, which requires making assumptions about the characteristics of potential participants. Researchers must choose which identity traits or background information to include, decisions that directly influence the model’s responses. While traditional interview recruitment is also subject to selection bias, simulating participants with LLMs makes these choices more explicit and high-stakes. Participants expressed concerns that this process risks reinforcing the researcher’s own biases, where researchers might inadvertently refine prompts to align the data with their expectations. The option to repeatedly prompt an LLM to receive subtly different responses each time introduces a risk of *confirmation bias*. Sophia captured this concern effectively:

You get to dictate who you’re talking to. In real interviews, I can’t just say, ‘Only give me immigrants from Ghana in their mid-thirties’. If you go into this mindset of interviewing with an LLM where you know exactly who you will talk to, then you will not learn anything because you already have expectations of what you’ll learn.

Researchers also emphasized that interviewing is an active process of meaning-making. Data is not waiting to be collected; rather, it is shaped and brought into existence through the researcher’s engagement with the community and their interpretation of it. Harper, who conducted research with religious communities, noted that the researcher’s presence can influence and shift the community’s practices and dynamics. Harper reflected, “*What foreign influences are felt? What aspects are already present? How does my respect for or belief in their system, or lack thereof, affect my approach and how they perceive it?*” Similarly, Yue discussed the effects of their own presentation during research with d/Deaf community members. Whether Yue presented as a hearing person with limited knowledge of sign language or as someone familiar with it could significantly impact the depth of detail participants provided about their challenges. Meaning and knowledge in qualitative research emerge from the relational context between researcher and participant.

When examining the limitations of using LLMs for simulating research participants, it is crucial to consider the differences between *emic* (insider) and *etic* (outsider) perspectives about the researchers’ own identities². Nadia pointed out how out-group researchers might struggle to identify stereotypes in the data if they lack direct experience with the topic or community being studied. Esme likened this issue to “*parachute science*”, where simulations are based on a shallow understanding of the community. Nikita elaborated how they avoid studying communities they are not embedded in: “*I’m not a person who does research from the outside. I would never go into an inner city and figure out what’s happening there unless I lived in that city.*”

²Emic-etic perspectives exist along a spectrum and are context-specific rather than binary and clearly differentiated. A solely emic perspective is impossible to achieve with the subjectivity of human experiences [84].

In contrast, in-group researchers could bring contextual knowledge to help them assess whether the data perpetuates harmful misrepresentations. Researchers in our study mentioned being drawn to projects related to their own background or lived experiences. However, using LLMs as research participants introduced the potential for encountering simulated experiences that reflect their community but are ultimately flawed. Nikita described this feeling as akin to the “*uncanny valley*,” where a machine’s responses seem human-like but reflect uncomfortable inaccuracies. Laila, too, felt unsettled after reading responses from the system trying to relate a person sharing similar experiences as her, describing it as “*creepy and disingenuous*”. This could lead to affective discomfort and harm, as researchers might experience these simulations as micro-aggressions.

Model’s epistemic position remains ambiguous. Epistemic position refers to one’s relationship to an object of knowledge, including the methods of gathering and interpreting that knowledge. Researchers struggled to determine the epistemic position reflected by the model, including whether it conveyed a singular viewpoint. Interviewees expressed how LLM responses often aggregated arguments from multiple interview participants from their study into a single consolidated response, which Daria described as a “*simulacrum of stories that people shared*”. Elliot, whose research examines workers’ perspectives on algorithmic management, observed how, “*the model combines a fictitious worker perspective with management’s perspective. A lot of what it presents as benefits of this technology are things that managers are sold on, but workers don’t necessarily experience.*” For Elliot, understanding the broader context—such as the wages, interpersonal conflicts between employees & employers, and power structures—was crucial for interpretation and generating a thick description. The consolidation of contrasting perspectives in LLM responses not only obscures a coherent worker voice; it also undermines the representation of partial, situated knowledge.

LLM responses are notoriously sensitive to the language and framing of prompts (cf. [98]), a pattern consistently observed in our interviews. Daria noted how the LLM was “*picking up on the valence of the system prompt and working hard to respond to it*” when she shifted from probing for “transparency concerns” to asking about “experiences with information available on the app.” Interviewees highlighted how subtle shifts in language (e.g., with structure, framing, or emotional undertones) drastically altered the model’s responses from a negative to a positive outlook. The model’s tendency for brittleness could compromise the reliability and consistency of their qualitative data.

Participants expressed concerns about the lack of transparency surrounding the data used to train LLMs. When the goal of research is to understand specific groups, the validity of LLM-generated responses is questionable if it is unclear whether those groups’ perspectives are adequately represented in the training data. Sophia illustrated this issue by pointing out that responses could vary significantly depending on whether the model had been trained on worker forums like ‘Uberpeople.net’ or more corporate-driven sources such as Uber.com’s own testimonial pages. Sophia brought attention to the challenges of evaluating the model’s epistemic position by emphasizing the uncertainty about the “*balance of this data or what it’s being trained on*” and the ‘freshness’ of perspectives given that the model’s training data typically only extends up to a

cut-off date (e.g., 2023 for GPT-4). Experiences are not only situated in relation to people and places but also in time. Nikita emphasized that “*a lot of the important bits of cultural context is its particular moment in time*,” and questioned whether the model was averaging perspectives across time or reflecting a specific moment, and if so, which one. Without clear information about the sources and proportions of data used to train LLMs, researchers struggled to identify the potential validity issues with the responses.

Facilitates erasure of community perspectives. Erasure refers to the systematic exclusion or invisibilization of certain groups and their standpoints when systems of knowledge production privilege certain voices over others. One significant risk of using LLMs in qualitative research is the erasure of perspectives from underrepresented groups. Laila, who studied Black social media creators on YouTube, noticed that the LLM tended to caricature participants, choosing stereotypical Black-sounding names when the topic involved Black history. She highlighted a broader concern: while LLMs might capture general sentiments *about* a community, they often fail to authentically represent voices *from* within that community. Similarly, Esme noticed how model responses “*included Black history tropes like representation and freedom in a way that wouldn’t resonate with a Black person doing this artwork in the South-east, for instance.*”

Researchers attributed much of the erasure to the pipeline for building the current generation of large language models, including the training data and the alignment process. LLM responses lacked the critical authenticity that human participants provide, particularly when researchers have invested time in building trust and rapport. Researchers felt that LLMs produced sanitized, politically correct responses, missing the messiness that characterizes real human experiences. Amir shared how his research participants would often share controversial opinions, such as which funding should be prioritized (ones related to the environment) over others (serving disabled people from developed countries). Similarly, Henri, who interviewed occupational therapists working with Alzheimer’s patients, noted that while LLMs might echo “*by-the-book*” responses, they fail to capture the contradictions that emerge in interview settings:

“In the first interview, the occupational therapist will give you the ‘best practice response’. They told us they try not to separate people with dementia from people without cognitive impairments. But yeah, they have to do it during a lot of activities, which is unfortunate and not advisable, but it’s a practice that happens, and they express on the fourth day.”

Researchers also highlighted the ambiguity surrounding how LLMs generate responses that are meant to reflect specific communities. When a model is assigned a persona based on a cultural, ethnic, or social identity, there is often uncertainty about whether the model is drawing from data that represents the lived experiences and perspectives of that community or simply regurgitating the surface-level characteristics associated with that community. Interviewees like Harper and Alice, who studied small, niche communities (e.g., investors and AI developers) with strong value systems, observed that the LLM could grasp broad-stroke dynamics (“*contours of the community*”) but missed the nuanced realities of those deeply embedded in these communities. Researchers went on

to problematize the idea of constructing a persona for an LLM, asking critical, rhetorical questions such as ‘okay, well, now [the model is] a Latina, but what does that mean? What does that mean to this LLM?’ If the model is only speaking *like* the participant description—using language, idioms, and stylistic markers associated with that identity—without grounding its responses in the actual data from that community, the results could be misleading or, worse, harmful stereotypes. These reflections highlighted a deeper discomfort with the superficiality of assigning complex social identities to a machine that does not embody the lived experiences of individuals with those identities.

Forecloses participants’ autonomy, agency, and consent. Using LLMs to simulate human behavior introduces several risks related to participants’ autonomy, consent, and agency. One way agency and participation manifest in the research process is through the expression of disagreement. Interviewees reported several instances where their human participants challenged the researchers’ premises or wording around the research topic. This critical engagement is a valuable aspect of the research process, ensuring that the data collected accurately reflects participants’ lived experiences, which might differ from the researchers’ initial assumptions. In contrast, researchers highlighted how LLMs were less likely to exhibit such resistance unless explicitly prompted. As Daria observed:

Many participants, especially in this kind of research, have their own agendas. Participants in studies about app issues know that we’re asking about problems, so they might tailor their responses accordingly. The LLM, similarly, seems to fit responses into expected categories without pushing back. In human interactions, however, people are more likely to disagree with the premise or question without needing explicit prompts. When I asked drivers about the data they receive from apps and how it aids their planning, I initially assumed that the apps did not provide much information. Yet, some drivers told me, ‘I get all the information I need about my work history.’

Researchers noted frustration with the model’s sycophantic tendency (cf. [89]) to agree with their points. Esme likened this to improvisational comedy, where performers are required to agree with everything said by their partners, but people contradict her all the time in her research, if “they don’t like [her] wording, they don’t have that experience, or they wouldn’t agree with what a question suggests. The whole point is that they’re not supposed to be a yes man.” Hugo expressed how the model would generate the perceived ‘preferred answer’ when probed about perceptions of robots in senior care homes.

In many cases, participants came to interviews eager to have their stories heard. Nico, who studied remote learning during the pandemic, observed that many students were dissatisfied with their college experiences and saw the interview as an opportunity to voice their frustrations. Participants actively sought to share their experiences, often introducing their own terminologies that Nico would then adopt. Similarly, Nikita noted how they “don’t want to talk to a computer telling [them] that trans lives matter or that trans health care is a problem. I wanna talk, I wanna be heard. I want us to

connect over that.” People’s desire to influence research outcomes is an important part of their expression of agency.

The use of LLMs raises ethical concerns about consent. LLMs might generate responses on sensitive topics that individuals might prefer not to disclose, overstepping boundaries that would otherwise be respected in an interview setting. Sophia, investigating the phenomenon of renter-ship among gig workers, noted a participant’s reluctance to discuss the topic further: “I asked one of my participants about this, and he was reluctant to answer. I could tell he knew more but wasn’t comfortable discussing it.” The use of LLMs as simulated subjects also compromises the autonomy and consent of the data subjects whose stories are potentially extracted in this interview data. Participants expressed hesitation about using a model whose training data is obtained without explicit consent. Laila compared this to the discourse surrounding AI-generated art, where artists’ creative works are used without permission to create new images. Similarly, potentially using experiences people share online to generate interview responses without their knowledge undermines the principle of autonomy, which is a cornerstone of research ethics.

Delegitimizes qualitative ways of knowing. The use of LLMs in qualitative research poses risks not only to the methodological integrity of the field but also to the legitimacy of qualitative research within academic circles. Several researchers voiced concerns that qualitative research is already undervalued, often perceived as less rigorous compared to quantitative approaches. The introduction of LLMs could exacerbate the marginalization of this mode of knowledge production by creating the impression that deeply contextual work of qualitative research can be easily replicated by computational models and done so more quickly. Sophia, who previously faced skepticism from reviewers when using qualitative methods, articulated how the use of LLMs in qualitative contexts seems like an attempt to demonstrate that “quantitative work can do what qualitative work does but faster.”

Researchers expressed fears that the availability of LLMs could encourage a ‘cutting corners’ mindset in research. The use of LLMs would function as a form of ‘data extraction’ that is antithetical to the values of qualitative research. Qualitative research is often iterative and collaborative— involving ongoing dialogue and reflection—where researchers work with participants to co-create knowledge grounded in lived experience. Indeed, researchers discussed how interviews and other forms of qualitative data collection are not just about gathering information; they are an important part of establishing and maintaining ongoing relationships with participants. Daria, Esme, and Elliot, for example, described how they continue to interact with participants beyond their research project through conversations with ride-hail drivers, following artists on social media, or collaborating with union members. When LLMs are used to generate participant responses, this collaborative process is taken over by a more transactional approach where data is extracted from the model without ongoing engagement with communities.

Another important concern raised by researchers was the potential damage that using LLMs could inflict on the trust between qualitative researchers and the communities they engage with. Historically, many vulnerable communities have developed deep distrust toward researchers due to exploitative practices in which academics extracted data or introduced interventions

without offering ongoing support or maintenance. This legacy of distrust could be further exacerbated if researchers begin to replace participant perspectives with LLM-generated responses. Yue, who works closely with d/Deaf participants, expressed concern that such practices would erode trust in research, especially within communities that are already wary of being misrepresented. Reducing these communities' voices to algorithmic outputs undermines the value of their lived experiences, potentially eroding the trust that researchers worked hard to build.

The concerns expressed by researchers reflect a broader anxiety about the potential implications of relying on LLMs in qualitative research. For Nikita, the use of LLMs evoked a *"sense of the dystopian,"* where, for example, trans communities are not heard by legislators or other decision-makers, as their perspectives are distorted by technology rather than being directly represented. Cameron summarized this concern by emphasizing that such tools miss the *"epistemological underpinnings of qualitative methods and why we do them"*. She argued that if the goal of qualitative research is to obtain contextualized data that is grounded in people's lived experiences, then LLMs fundamentally fail to meet this standard. LLMs may generate text that *appears* fluent and contextually relevant, but this output lacks the depth, nuance, and authenticity that come from direct engagement with people's lived realities.

4.3 Imagining Use Cases for LLMs in Qualitative Research: Seeing Possibilities and Even More Caveats

Most researchers we interviewed voiced their discomfort with using LLMs to generate synthetic research data, with several even suggesting that they would distrust findings from studies that utilize LLMs as research participants. As a thought exercise, we explored whether there might be specific domains, contexts, or use cases where LLMs could be more effective. Below, we outline the use cases that researchers identified as *more* appropriate than replacing direct engagement with communities. It is important to note that our participants did not reach a consensus on any of these use cases. For each potential application, participants also highlighted ways in which LLM use could be harmful or ineffective.

Interviewees suggested that using LLMs to simulate participants could serve as a valuable pedagogical tool where the stakes are lower than in a research study. This approach could help novice researchers learn how to focus on specific aspects of responses and ask effective follow-up questions. Some interviewees also highlighted flaws in this proposition, noting that prompting an LLM often differs significantly from interacting with human participants. For example, Daria had to prompt the model thrice before it provided enough contextual detail to make a useful response. Researchers emphasized that a critical interviewing skill is managing emotions—their own and the participant's—which is difficult to replicate when learning qualitative methods with LLMs. Jasmine also pointed out the risk that junior researchers learning qualitative methods with LLMs might develop poor interviewing habits (*e.g.*, not building rapport, not probing for depth, interrupting participants, or ignoring body language) that do not translate well to real research contexts and could perpetuate the notion of *"extracting information"* from research subjects.

For most researchers, LLMs could, at their best, help test their interview protocol, especially when it is challenging to recruit participants for pilot studies. In those situations, LLMs could serve as stand-ins, allowing researchers to gauge the types of responses a question might elicit or determine what kinds of questions to ask. Here, Mario cautioned that relying on LLMs for pilot interviews or refining interview guides could shift the focus of the project in unexpected ways. He pointed out how LLMs might lead researchers to become overly fixated on specific directions that might otherwise not occur if researchers were interviewing real participants.

Several researchers observed that the decision to use LLMs depends on the research topic and the community being simulated. In sensitive research contexts, such as those involving experiences of oppression or discrimination, participants noted that engaging with LLMs could carry some potential benefits and the significant risks described above. Sensitive research topics risk placing an epistemic burden on participants and re-traumatizing them by asking them to revisit difficult experiences. While some researchers suggested that LLMs could be useful in these contexts, they also expressed concerns that using LLMs for sensitive topics might exacerbate the erasure of real human experiences. Nadia, for example, was skeptical about the ability of LLMs to accurately simulate complex human experiences, such as navigating gender identity or sexuality in the workplace or the experiences of refugees and migrants. On the other hand, Nadia expressed how there are certain communities, such as hate groups, that they do not feel safe engaging with. She argued that those interviews could be simulated with an LLM to help develop defensive strategies against hate groups online. Nikita argued that for sensitive topics, researchers should consider collaborating with community members to gain relevant expertise and learn how to have difficult conversations, rather than *"turning to an LLM."*

5 Discussion

The claim that technologies can act as surrogates recapitulates histories of disappearance, erasure, and elimination necessary to maintain the liberal subject as the agent of historical progress.

(Neda Atanasoski and Kalindi Vora [7])

In *Surrogate Humanity*, Atanasoski and Vora reveal how technologies like social robots and artificial intelligence are not neutral tools but are instead shaped by capitalist logics of differential exploitation and dispossession [7]. These technologies are designed to act as *surrogates*, ostensibly to emancipate humans from "historically degraded tasks", such as engaging with research participants in comparison with developing cutting-edge technologies. When LLMs are used as stand-ins for human participants in research, they similarly function as surrogates in replacing the voices and lived experiences of actual communities with algorithmic approximations. Our findings highlight participants' concerns on how the use of LLMs amplifies researchers' positionality, erases community perspectives, and undermines the legitimacy of qualitative ways of knowing. While improving simulation techniques, refining prompts, or using more advanced models might seem promising paths forward, we argue that these approaches centralize the role of technology in knowledge production and do not address the underlying epistemic risks. We invite a critical reexamination of how

LLMs, products of specific socio-technical assemblages, contribute to a reconfiguration of what counts as valid knowledge and who gets to produce it. Below, we draw on scholarship from qualitative epistemologies and research ethics to argue for a critical reconsideration of the use of LLMs as surrogate participants, advocating for caution and, in most cases, avoiding such practices.

5.1 The incongruence of LLMs and qualitative epistemologies

Our results indicate how the use of LLMs in qualitative research obscures the particularities of lived experience and fundamentally misaligns with the epistemological commitments of qualitative research. Human subjects research, as Ithiel de Sola Pool describes, is an active process of constructing meaning that unfolds as an “*interpersonal drama with a developing plot*” [103]. Research participants, then, are not ‘*vessels of answers*’ waiting to be tapped in to elicit truths about facts, opinions, or their lived experiences [68]. LLM-generated data lacks this intersubjective relational understanding that emerges from interviews. Our participant Harper’s reflection on her presence within the religious community and Yue’s account of her engagement with the d/Deaf community illustrate how a researcher’s involvement actively shapes community practices, perceptions, and the broader dynamics of knowledge production.

Qualitative research epistemologies, especially within HCI, emphasize the importance of researcher positionality: the idea that knowledge is co-constructed by the researcher and the participant and is deeply influenced by the specific context in which it is generated. The interview is not a neutral process of extracting information that exists *out there* [106]; instead, research interviews are intentional, and the interviewer is implicated in the process of assembling knowledge. Researchers in our study observed that LLM responses often appeared plausible. However, this assessment of plausibility was likely influenced by our study design, which allowed participants to compare synthetic data against their prior experience. In practice, researchers engaging with unfamiliar domains or communities would lack the contextual grounding needed to determine whether LLM-generated data is plausible, let alone palpable.

An interpretivist orientation values the nuances of lived experience that are often best captured through methods that allow for deep engagement, reflection, and interpretation. Indeed, there are many sites of knowledge production, particularly in a community-based project, that exist outside the confines of a survey questionnaire or one-hour interview [110]. This multisitedness, articulated by Law [68], involves a multiplicity of locations where ‘*analysis and knowledge are made material*.’ Our participant Daria described interviews not simply as a means to generate research data, but as part of an ongoing effort to build relationships with workers. LeDantec and Fox [69] describe this as work that comes before the work: maintaining collaborations and demonstrating commitments that shape the outcomes of the projects. We argue that interviewing with LLMs is fundamentally misaligned with the notion that “*methods [no longer] discover and depict realities [...] they participate in the enactments of those realities*.” [68]

The implications of this limitation become even more apparent when engaging with Haraway’s concept of situated knowledges [52]. Qualitative methods are designed to produce knowledge that is

not only rich and detailed, in the form of a *thick description* [45], but also situated within specific cultural, social, and material contexts. Haraway’s idea of situated knowledge invites a critical perspective that helps to problematize ‘*the god trick of seeing everything from nowhere*’ [52]. Our participant Ezra’s reflection brings this into focus: workers who have encountered specific glitches on particular days provide granular, firsthand observations that can speak to immediate as well as systemic issues. This reminds us of Haraway’s call to resist myths of infinite vision seemingly offered by the visualization tricks and powers of modern technology. In thinking with Haraway, we contend that the prevailing logics of LLMs as proxies for human behavior offer similar *unlocatable* knowledge claims.

LLMs, by their very design, are built to extract patterns that can be applied across contexts with the aim of achieving universality and objectivity [13]. While this capacity can seemingly lead to efficiencies and broad applicability, it also means that the knowledge produced by LLMs is often decontextualized and disembodied. The models are trained on data that, despite its volume, may lack the depth and specificity that qualitative researchers seek. Moreover, our participants revealed how LLMs do not account for the positionality of the researcher or the participants; instead, they flatten differences and obscure the peculiarities that are central to qualitative inquiry. This turn to LLMs reflects a world in which the complexity of human experience are reduced to manageable data points that can be easily processed by algorithms. In doing so, the use of LLMs in research aligns with a vision of science that prioritizes efficiency and scalability over contextual engagement with communities.

5.2 LLM use undermines consent and autonomy of data subjects

HCI research has long been dedicated to supporting the autonomy and agency of research participants. Our community recognizes the importance of ensuring that individuals have control over their participation and the information they share. This concern is reflected in participants’ emphasis on informed consent, such as Laila highlighting the discomfort and moral complexities surrounding the use of data without explicit consent in training LLMs. These concerns mirror broader debates in HCI and CSCW, where scholars have questioned the implications of using ‘public’ social media data for research purposes [64, 125]. This includes practices, such as the mining of profile pictures, to identify social categories such as sexual orientation [71] or facial recognition from students’ photos [42]. Social media users are often unaware that their data is being used by researchers, and in fact, many participants in Fiesler and Proferes’s study asserted that researchers should not be able to use tweets without explicit consent [39]. Online content creators, including social media users, may not realize the public nature of their data and the downstream applications the data may be supporting [39].

We believe similar risks apply when researchers use LLMs to create synthetic interview data. Most current LLMs are trained on public content extracted from the Internet [25]. Earlier this year, Google and OpenAI signed contracts with Reddit that gives these companies real-time access to Reddit posts [33, 92]. OpenAI maintains that ‘training AI models using publicly available internet materials is fair use, as supported by longstanding and widely accepted

precedents' [86]. Many other big tech companies have made arguments to similar effect when faced with lawsuits. While training on publicly available and/or copyrighted text may not be unlawful (yet), generating research data using models created through the expropriation of the work of others presents the risk of undermining the autonomy of data subjects powering these applications.

To understand the relationship between autonomy and consent in research, we borrow from Beauchamp's work that challenges the traditional implementation of 'informed consent' [14]. Rather than focusing narrowly on the disclosure of information to research subjects, Beauchamp advocates for a more nuanced understanding of 'voluntary consent' that includes self-authorization [14]. Can the data subjects whose information is used to train LLMs choose to accept or decline participation in research? Consider how researchers using LLMs to simulate participants might expose sensitive information from online communities, such as discussions on Reddit about particular worker resistance tactics that would have remained largely unknown otherwise. These communities, whose members might not have intended to share their experiences with researchers, could find their information under a microscope without their explicit consent.

Here, an instinct might be to analogize LLM interviews and the practice of creating composites where researchers combine data from multiple participants to protect anonymity [79]. We argue that the two approaches diverge in critical ways. Traditional composites are rooted in ethical commitments to participants, including respect for their anonymity and consent. The process of creating composites requires more than anonymization; it requires interpretive work, and the researcher is responsible for ensuring the narratives do not become simplistic or caricatured [115]. These composites are ultimately grounded in the lived experiences of real individuals. In contrast, LLM-generated interview data presents a different kind of 'composite', one that lacks the intentionality and contextual grounding that characterize traditional qualitative methods. The data sources and contexts that inform LLM outputs are largely unknown to the researcher using these tools [59]. In reflecting on these differences, it becomes clear that LLM-generated interview data represents a fundamentally different epistemological practice from qualitative composite narratives.

We call for extending consideration of voluntary informed consent beyond our direct interactions with research subjects to the individuals whose data is used to produce these models, particularly in terms of how their autonomy is respected or likely potentially subverted.

5.3 What does it mean to simulate the human in HCI?

One might contend that HCI's primary goal is to understand how people interact with technology and design systems that meet their needs [85]. Any attempts to use simulations, such as user personas or LLMs, to represent communities' interests must then be critically examined. Scholars have emphasized how simulations depend on the assumptions of their creators, which often reflect hegemonic norms and biases [37].

A clear example of this is disability simulation, where designers attempt to 'step into the shoes' [102] of disabled individuals by wearing a blindfold to simulate visual impairment or using a wheelchair for a limited time [11]. While these exercises may ostensibly help evoke empathy, disability scholars have long critiqued this approach for producing unrealistic and reductive understandings of life with a disability [16, 101]. Simulations often reduce the complexity of disability to a narrow, physicalized experience, ignoring the systemic barriers and social dynamics that shape disabled individuals' lives. The reliance on simulation as a method of understanding disability ends up undermining the expertise of disabled individuals themselves. This dynamic is emblematic of broader power differentials in design, where those with the most at stake are excluded from meaningful participation in the process. There is a need, then, not only to reconsider the methods used to develop simulations but also the very premise of using simulations to stand in for lived experiences.

Our interviewees were acutely aware of the misalignment between qualitative research and the use of large language models as surrogate participants. Many researchers highlighted these tensions using terms such as proxy, simulacrum, surrogates, and personas³. To understand the tensions and ambivalences that researchers described, we return to a set of tentative proposals for LLM use offered by participants. This includes using LLMs to trial materials, as a pedagogical device, and in sensitive contexts or unsafe environments. Notably, as soon as researchers suggested them, many acknowledged a set of complex ethical dimensions, such as erasure and misrepresentation or misleading qualitative research training.

In this final section, we engage with those dilemmas and their implications. We outline three critical considerations—transparency, model selection, and validation—for researchers who find appeal in the use of LLMs as simulated qualitative research participants. These considerations are not meant to serve as guidance to justify the use of LLMs. On the contrary, we urge careful reflection on how their use involves complex methodological choices that do not resolve the limitations discussed earlier and introduce new challenges that impact the integrity of qualitative research.

Firstly, thoughtful model selection is not merely a performance-driven decision but one that directly impacts the validity of the synthetic research data. For example, a model trained predominantly on Western-centric datasets may not accurately simulate experiences from non-Western contexts [61]. Compounding these issues, different models undergo different alignment processes depending on the goals and values of the developers that embed specific priorities and mannerisms into the outputs. Crucially, researchers must question the assumption that training datasets adequately capture *any* group's perspectives. Even when data *about* a group is included, it often reproduces the perspectives of outsiders rather than centering narratives *from* communities themselves [112]. The use of LLMs exacerbates what Linda Alcoff identifies as the problem of speaking for others [4]. Alcoff contends that creating conditions for dialogue and engaging in the practice of speaking *with* others can lessen the risks of "misrepresentation, expanding one's own authority and privilege, and a generally imperialist speaking ritual" inherent in speaking *for* others.

³Although participants often used these terms interchangeably, we recognize that they carry distinct meanings across disciplines, and we hope future work will explore these conceptual differences.

Secondly, the field is increasingly emphasizing the need for transparency in the use of LLMs in research [59], as reflected in policies like the ACM's [19], including any pilot studies with simulated participants. However, simply disclosing information does not automatically invite critique or enable a better understanding of the research design; instead, visibility often shifts the burden onto other stakeholders [5]. Transparency can create an illusion of accountability without fostering meaningful engagement or critical reflection. Simply providing access to transcripts or explaining the role of LLMs does not guarantee that stakeholders (e.g., peer reviewers or community members) will have the necessary time, resources, or expertise to evaluate the information. Moreover, with the limited transparency from model producers regarding the composition of their training data or specifics of their model architecture, it would be difficult for researchers to fully assess and communicate the appropriateness of a model for their study. Transparency, then, risks becoming *performative*, especially if disclosure can only be selective.

Thirdly, researchers in our study recommended turning to member checking as a form of validation, where findings or interpretations are returned to community members for review and feedback [17]. When LLMs are involved, member checking introduces additional layers of complexity. How might people respond if they discovered that researchers were using LLMs to simulate them or others in their community? Engaging with simulated data that reflects one's community or personal experiences demands significant time, energy, and emotional labor. Reading simulated experiences, especially those that may inaccurately represent or trivialize lived experiences, could be distressing for community members. Encouraging community members to critically engage with LLM-generated data and voice disagreements if the synthetic data does not align with their perspectives could also prove challenging. Lastly, member checking involves verifying transcripts or synthesized data with the same individual who provided the data. However, with LLM-generated data, there is no corresponding individual who can speak to the validity of the output. Scholars of qualitative research strategies have suggested that the use of member checking may be justified as part of prolonged engagement in participatory research [108]. However, as we note above, these conditions are unlikely to be met when researchers use LLMs to simulate participants.

Lastly, it is crucial to acknowledge the invisible gendered, classed, and racialized labor underpinning these systems. LLMs are not self-sustaining entities that emerge in a vacuum; their creation involves vast networks of a global workforce that contribute to the training and maintenance of these systems [81]. These workers, who are often underpaid and overworked [93] engage in the labor of labeling data, evaluating and moderating content, and correcting model errors, all of which are essential for the functionality of LLMs. Tech companies and platforms aim to disguise this human labor as machine labor to create a veneer of machine intelligence [7]. As Atanasoski and Vora articulate, "labor becomes something that is intentionally obfuscated to create the effect of machine autonomy" [7]. This intentional obfuscation creates a *conjunction of algorithms* [83] that enables the fantasy of magical techno-objects and masks the exploitation and uneven distribution of labor benefits in the global tech industry.

6 Conclusion

This paper focuses on examining the use of large language models as proxies for research participants. This is no longer a speculative future; research papers and various products are already exploring LLMs' ability to replace human participants. Through interviews with qualitative researchers, we highlight several significant limitations of using LLMs as stand-ins for qualitative research participants, including the lack of palpability in LLM responses, the ambiguity surrounding the model's epistemic position, the amplification of researcher positionality, how the use of LLMs forecloses participants' consent and agency, the erasure of community perspectives, and the risk of undermining qualitative methods of knowing. We conclude by drawing on STS and research ethics literature to discuss how these practices bring to light critical questions about consent, agency, and the legitimacy of using LLMs in qualitative research. The ethical and epistemological tensions demand a cautious and, in many cases, critical stance toward this form of substitution in research that seeks to understand lived experiences.

Acknowledgments

First and foremost, we are grateful to our most thoughtful study participants for their time and generosity, without whom this research would not be possible.

We would like to thank Min Park, Ji-Youn Jung, Ningjing Tang and Jini Kim for taking the time to participate in pilot sessions for this study. We would also like to thank Anubhav Jangra, Inha Cha, Ding Wang, Tianling Yang, Jini Kim, Ningjing Tang, and Qing Xiao for their valuable suggestions on early drafts of this work. A heartfelt thank you to members of the Tech Solidarity Lab—Jordan Taylor, Cella Sum, Franchesca (Franky) Spektor, and Christine Mendoza—whose feedback helped us engage with this project more critically and with more nuance.

References

- [1] Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Persistent Anti-Muslim Bias in Large Language Models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) (AIES '21). Association for Computing Machinery, New York, NY, USA, 298–306. <https://doi.org/10.1145/3461702.3462624>
- [2] William Agnew, A. Stevie Bergman, Jennifer Chien, Mark Diaz, Selim El-Sayed, Jaylen Pittman, Shakir Mohamed, and Kevin R. McKee. 2024. The Illusion of Artificial Inclusion. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 286, 12 pages. <https://doi.org/10.1145/3613904.3642703>
- [3] Gati V Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023. Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, Hawaii, USA, 337–371. <https://proceedings.mlr.press/v202/aher23a.html>
- [4] Linda Alcoff. 1991. The problem of speaking for others. , 5–32 pages.
- [5] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20 (2018), 973 – 989. <https://api.semanticscholar.org/CorpusID:5001487>
- [6] Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis* 31, 3 (2023), 337–351.
- [7] Neda Atanasoski and Kalindi Vora. 2019. Surrogate humanity. *Londres e Durham: Duke University Press*. DOI 10 (2019), 9781478004455.
- [8] Mark Avis. 2005. Is there an epistemology for qualitative research. , 3–16 pages.

- [9] Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2024. ResearchAgent: Iterative Research Idea Generation over Scientific Literature with Large Language Models. arXiv:2404.07738 [cs.CL] <https://arxiv.org/abs/2404.07738>
- [10] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislaw Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. Constitutional AI: Harmlessness from AI Feedback. arXiv:2212.08073 [cs.CL] <https://arxiv.org/abs/2212.08073>
- [11] Jan Balata, Zdenek Mikovec, and Lukas Neoproud. 2015. BlindCamera: Central and Golden-ratio Composition for Blind Photographers. In *Proceedings of the Multimedia, Interaction, Design and Innovation* (Warsaw, Poland) (MIDI '15). Association for Computing Machinery, New York, NY, USA, Article 8, 8 pages. <https://doi.org/10.1145/2814464.2814472>
- [12] Jianzhu Bao, Rui Wang, Yasheng Wang, Aixin Sun, Yitong Li, Fei Mi, and Ruifeng Xu. 2023. A Synthetic Data Generation Framework for Grounded Dialogues. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 10866–10882. <https://doi.org/10.18653/v1/2023.acl-long.608>
- [13] S. Barocas, M. Hardt, and A. Narayanan. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, Cambridge, USA. <https://books.google.com/books?id=uawEAAQBAJ>
- [14] Tom L Beauchamp. 2010. Autonomy and consent. , 55–78 pages.
- [15] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Mitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [16] Cynthia L. Bennett and Daniela K. Rosner. 2019. The Promise of Empathy: Design, Disability, and Knowing the "Other". In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3290605.3300528>
- [17] Linda Birt, Suzanne Scott, Debbie Cavers, Christine Campbell, and Fiona Walter. 2016. Member checking: a tool to enhance trustworthiness or merely a nod to validation? *Qualitative health research* 26, 13 (2016), 1802–1811.
- [18] Su Lin Blodgett, Solon Barocas, Hal Daum'e, and Hanna M. Wallach. 2020. Language (Technology) is Power: A Critical Survey of "Bias" in NLP. <https://api.semanticscholar.org/CorpusID:218971825>
- [19] ACM Publications Board. 2023. ACM Policy on Authorship. <https://www.acm.org/publications/policies/new-acm-policy-on-authorship>. (Accessed on 09/08/2024).
- [20] Daniil A. Boiko, Robert MacKnight, and Gabe Gomes. 2023. Emergent autonomous scientific research capabilities of large language models. arXiv:2304.05332 [physics.chem-ph] <https://arxiv.org/abs/2304.05332>
- [21] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc., Barcelona, Spain. https://proceedings.neurips.cc/paper_files/paper/2016/file/a486cd07e4ac3d270571622f4f316ec5-Paper.pdf
- [22] Shikha Bordia and Samuel R. Bowman. 2019. Identifying and Reducing Gender Bias in Word-Level Language Models. arXiv:1904.03035 [cs.CL] <https://arxiv.org/abs/1904.03035>
- [23] Virginia Braun and Victoria Clarke. 2012. Thematic analysis.
- [24] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.
- [25] Antonio Byrd. 2023. Truth-Telling: Critical Inquiries on LLMs and the Corpus Texts That Train Them. *Composition studies* 51, 1 (2023), 135–142.
- [26] Courtney Byun, Piper Vasicek, and Kevin Seppi. 2023. Dispensing with Humans in Human-Computer Interaction Research. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). Association for Computing Machinery, New York, NY, USA, Article 413, 26 pages. <https://doi.org/10.1145/3544549.3582749>
- [27] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186. <https://doi.org/10.1126/science.aal4230> arXiv:https://www.science.org/doi/pdf/10.1126/science.aal4230
- [28] Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate. arXiv:2308.07201 [cs.CL] <https://arxiv.org/abs/2308.07201>
- [29] Alicja Chaszczewicz, Raj Sanjay Shah, Ryan Louie, Bruce A Arnou, Robert Kraut, and Diyi Yang. 2024. Multi-Level Feedback Generation with Large Language Models for Empowering Novice Peer Counselors. arXiv:2403.15482 [cs.CL] <https://arxiv.org/abs/2403.15482>
- [30] Cheng-Han Chiang and Hung yi Lee. 2023. Can Large Language Models Be an Alternative to Human Evaluations? arXiv:2305.01937 [cs.CL] <https://arxiv.org/abs/2305.01937>
- [31] Sasha Costanza-Chock. 2020. *Design justice: Community-led practices to build the worlds we need*. The MIT Press, Cambridge, USA.
- [32] Andy Crabtree. 2024. H is for Human and How (Not) To Evaluate Qualitative Research in HCI. arXiv:2409.01302 [cs.HC] <https://arxiv.org/abs/2409.01302>
- [33] Emilia David. 2024. OpenAI strikes Reddit deal to train its AI on your posts - The Verge. <https://www.theverge.com/2024/5/16/24158529/reddit-openai-chatgpt-api-access-advertising>. (Accessed on 09/01/2024).
- [34] Dorothea Demszyk, Diyi Yang, David S Yeager, Christopher J Bryan, Margaret Clapper, Susannah Chandhok, Johannes C Eichstaedt, Cameron Hecht, Jeremy Jamieson, Meghann Johnson, et al. 2023. Using large language models in psychology. *Nature Reviews Psychology* 2, 11 (2023), 688–701.
- [35] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. 2024. Mind2web: Towards a generalist agent for the web.
- [36] Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy S Liang, and Tatsunori B Hashimoto. 2023. AlpacaFarm: A Simulation Framework for Methods that Learn from Human Feedback. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., New Orleans, USA, 30039–30069. https://proceedings.neurips.cc/paper_files/paper/2023/file/5fc47800ee5b30b8777dd30abcaaf3b-Paper-Conference.pdf
- [37] Emory James Edwards, Cella Monet Sum, and Stacy M. Branham. 2020. Three Tensions Between Personas and Complex Disability Identities. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3334480.3382931>
- [38] Jessica L Feuston and Jed R Brubaker. 2021. Putting tools in their place: The role of time and perspective in human-AI collaboration for qualitative analysis. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–25.
- [39] Casey Fiesler and Nicholas Proferes. 2018. "Participant" perceptions of Twitter research ethics. *Social Media+ Society* 4, 1 (2018), 2056305118763366.
- [40] Miranda Fricker. 2007. *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press, Oxford, UK.
- [41] Ivar Frisch and Mario Giulianelli. 2024. LLM Agents in Interaction: Measuring Personality Consistency and Linguistic Alignment in Interacting Populations of Large Language Models. arXiv:2402.02896 [cs.CL] <https://arxiv.org/abs/2402.02896>
- [42] Sidney Fussell. 2019. You no longer own your face.
- [43] Jie Gao, Yuchen Guo, Gionnive Lim, Tianqin Zhang, Zheng Zhang, Toby Jia-Jun Li, and Simon Tangi Perrault. 2024. CollabCoder: A Lower-barrier, Rigorous Workflow for Inductive Collaborative Qualitative Analysis with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 11, 29 pages. <https://doi.org/10.1145/3613904.3642002>
- [44] Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. 2024. Scaling Synthetic Data Creation with 1,000,000,000 Personas. arXiv:2406.20094 [cs.CL] <https://arxiv.org/abs/2406.20094>
- [45] Clifford Geertz. 2008. Thick description: Toward an interpretive theory of culture. , 41–51 pages.
- [46] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. <https://doi.org/10.1073/pnas.2305016120>
- [47] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences* 120, 30 (2023), e2305016120.
- [48] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large Language Model based Multi-Agents: A Survey of Progress and Challenges. arXiv:2402.01680 [cs.CL] <https://arxiv.org/abs/2402.01680>
- [49] Shashank Gupta, Vaishnavi Shrivastava, A. Deshpande, A. Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2023. Bias Runs Deep: Implicit Reasoning Biases in Persona-Assigned LLMs. <https://api.semanticscholar.org/CorpusID:265050702>
- [50] Andrew Halterman and Katherine A. Keith. 2025. Codebook LLMs: Evaluating LLMs as Measurement Tools for Political Science Concepts. arXiv:2407.10747 [cs.CL] <https://arxiv.org/abs/2407.10747>

- [51] Perttu Hämäläinen, Mikke Tavast, and Anton Kunnari. 2023. 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) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 433, 19 pages. <https://doi.org/10.1145/3544548.3580688>
- [52] Donna Haraway. 2013. Situated knowledges: The science question in feminism and the privilege of partial perspective 1. In *Women, science, and technology*. Routledge, Milton Park, Oxfordshire, UK, 455–472.
- [53] Masum Hasan, Cengiz Ozel, Sammy Potter, and Ehsan Hoque. 2023. SAPIEN: Affective Virtual Agents Powered by Large Language Models*. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*. IEEE, Cambridge, MA, USA, 1–3. <https://doi.org/10.1109/ACIIW59127.2023.10388188>
- [54] John J. Horton. 2023. Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus? arXiv:2301.07543 [econ.GN]
- [55] Jinpeng Hu, Tengting Dong, Gang Luo, Hui Ma, Peng Zou, Xiao Sun, Dan Guo, Xun Yang, and Meng Wang. 2024. PsycOLLM: Enhancing LLM for Psychological Understanding and Evaluation. , 13 pages. <https://doi.org/10.1109/TCSS.2024.3497725>
- [56] Synthetic Users Inc. 2024. Synthetic Users: user research without the headaches. <https://web.archive.org/web/20240822224158/https://www.syntheticusers.com/>. (Accessed on 08/30/2024).
- [57] OpinioAI International. 2024. AI Powered Research | OpinioAI. <https://www.opinio.ai/>. (Accessed on 08/30/2024).
- [58] Julie Jiang and Emilio Ferrara. 2023. Social-LLM: Modeling User Behavior at Scale using Language Models and Social Network Data.
- [59] Shivani Kapania, Ruiyi Wang, Toby Jia-Jun Li, Tianshi Li, and Hong Shen. 2024. "I'm categorizing LLM as a productivity tool": Examining ethics of LLM use in HCI research practices. arXiv:2403.19876 [cs.HC] <https://arxiv.org/abs/2403.19876>
- [60] Luoma Ke, Song Tong, Peng Cheng, and Kaiping Peng. 2024. Exploring the Frontiers of LLMs in Psychological Applications: A Comprehensive Review. arXiv:2401.01519 [cs.LG] <https://arxiv.org/abs/2401.01519>
- [61] Khyati Khandelwal, Manuel Tonneau, Andrew M. Bean, Hannah Rose Kirk, and Scott A. Hale. 2024. Indian-BhED: A Dataset for Measuring India-Centric Biases in Large Language Models. In *Proceedings of the 2024 International Conference on Information Technology for Social Good* (Bremen, Germany) (GoodIT '24). Association for Computing Machinery, New York, NY, USA, 231–239. <https://doi.org/10.1145/3677525.3678666>
- [62] Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2024. Language models can solve computer tasks.
- [63] Junsol Kim and Byungkyu Lee. 2023. AI-Augmented Surveys: Leveraging Large Language Models and Surveys for Opinion Prediction. arXiv:2305.09620 [cs.CL]
- [64] Shamika Klassen and Casey Fiesler. 2022. "This isn't your data, friend": Black Twitter as a case study on research ethics for public data. *Social Media+ Society* 8, 4 (2022), 20563051221144317.
- [65] Taja Kuzman, Igor Mozetič, and Nikola Ljubešić. 2023. ChatGPT: Beginning of an End of Manual Linguistic Data Annotation? Use Case of Automatic Genre Identification. arXiv:2303.03953 [cs.CL]
- [66] Michelle S. Lam, Janice Teoh, James A. Landay, Jeffrey Heer, and Michael S. Bernstein. 2024. Concept Induction: Analyzing Unstructured Text with High-Level Concepts Using LLMs. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 766, 28 pages. <https://doi.org/10.1145/3613904.3642830>
- [67] Ehsan Latif, Luyang Fang, Ping Ma, and Xiaoming Zhai. 2024. Knowledge Distillation of LLM for Automatic Scoring of Science Education Assessments. arXiv:2312.15842 [cs.CL] <https://arxiv.org/abs/2312.15842>
- [68] J. Law. 2004. *After Method: Mess in Social Science Research*. Routledge, Milton Park, Oxfordshire, UK. <https://books.google.com/books?id=E20X7N0nBfQC>
- [69] Christopher A. Le Dantec and Sarah Fox. 2015. Strangers at the Gate: Gaining Access, Building Rapport, and Co-Constructing Community-Based Research. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) (CSCW '15). Association for Computing Machinery, New York, NY, USA, 1348–1358. <https://doi.org/10.1145/2675133.2675147>
- [70] Jung-Joo Lee et al. 2012. Against method: The portability of method in human-centered design.
- [71] Sam Levin. 2017. New AI can guess whether you're gay or straight from a photograph.
- [72] Tao Li, Daniel Khashabi, Tushar Khot, Ashish Sabharwal, and Vivek Srikumar. 2020. UNQOVERing Stereotyping Biases via Underspecified Questions. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 3475–3489. <https://doi.org/10.18653/v1/2020.findings-emnlp.311>
- [73] Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards Understanding and Mitigating Social Biases in Language Models. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, Virtual, 6565–6576. <https://proceedings.mlr.press/v139/liang21a.html>
- [74] Menglin Liu and Ge Shi. 2024. PoliPrompt: A High-Performance Cost-Effective LLM-Based Text Classification Framework for Political Science. arXiv:2409.01466 [cs.CL] <https://arxiv.org/abs/2409.01466>
- [75] Yiren Liu, Si Chen, Haocong Cheng, Mengxia Yu, Xiao Ran, Andrew Mo, Yiliu Tang, and Yun Huang. 2024. How AI Processing Delays Foster Creativity: Exploring Research Question Co-Creation with an LLM-based Agent. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 17, 25 pages. <https://doi.org/10.1145/3613904.3642698>
- [76] Ryan Louie, Ananjan Nandi, William Fang, Cheng Chang, Emma Brunskill, and Diyi Yang. 2024. Roleplay-doh: Enabling Domain-Experts to Create LLM-simulated Patients via Eliciting and Adhering to Principles. arXiv:2407.00870 [cs.CL] <https://arxiv.org/abs/2407.00870>
- [77] Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. 2024. Augmenting large language models with chemistry tools. , 11 pages.
- [78] Kevin Ma, Daniele Grandi, Christopher McComb, and Kosa Goucher-Lambert. 2023. Conceptual Design Generation Using Large Language Models. In *2023 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, Vol. Volume 6: 35th International Conference on Design Theory and Methodology (DTM))*. ASME, Boston, USA, V006T06A021. <https://doi.org/10.1115/DETC2023-116838> arXiv:https://asmedigitalcollection.asme.org/IDETC-CIE/proceedings-pdf/IDETC-CIE2023/87349/V006T06A021/7061719/v006t06a021-detc2023-116838.pdf
- [79] Annette Markham. 2012. Fabrication as ethical practice: Qualitative inquiry in ambiguous internet contexts. *Information, Communication & Society* 15, 3 (2012), 334–353.
- [80] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.
- [81] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. *Proceedings of the ACM on Human-Computer Interaction* 6 (2022), 1 – 37. <https://api.semanticscholar.org/CorpusID:249017734>
- [82] Moin Nadeem, Anna Bethke, and Siva Reddy. 2020. StereoSet: Measuring stereotypical bias in pretrained language models. arXiv:2004.09456 [cs.CL] <https://arxiv.org/abs/2004.09456>
- [83] Peter Nagy and Gina Neff. 2024. Conjuring algorithms: Understanding the tech industry as stage magicians. *New Media & Society* 26, 9 (2024), 4938–4954. <https://doi.org/10.1177/14614448241251789> arXiv:https://doi.org/10.1177/14614448241251789
- [84] James L. Olive. 2014. Reflecting on the Tensions Between Emic and Etic Perspectives in Life History Research: Lessons Learned. <https://doi.org/10.17169/fqs-15.2.2072>
- [85] Judith S Olson and Wendy A Kellogg. 2014. *Ways of Knowing in HCI*. Vol. 2. Springer, New York, USA.
- [86] OpenAI. 2024. OpenAI and journalism | OpenAI. <https://openai.com/index/openai-and-journalism/>. (Accessed on 09/01/2024).
- [87] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 2, 22 pages. <https://doi.org/10.1145/3586183.3606763>
- [88] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 74, 18 pages. <https://doi.org/10.1145/3526113.3545616>
- [89] Ethan Perez, Sam Ringer, Kamilė Lukošiušė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, Andy Jones, Anna Chen, Ben Mann, Brian Israel, Bryan Seethor, Cameron McKinnon, Christopher Olah, Da Yan, Daniela Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Guro Khundadze, Jackson Kernion, James Landis, Jamie Kerr, Jared Mueller, Jeeyoon Hyun, Joshua Landau, Kamal Ndousse, Landon Goldberg, Liane Lovitt, Martin Lucas, Michael Sellitto, Miranda Zhang, Neerav Kingsland, Nelson Elhage, Nicholas Joseph, Noemi Mercado, Nova DasSarma, Oliver Rausch, Robin Larson, Sam McCandlish, Scott Johnston, Shauna Kravec, Sheer El Showk, Tamera Lanham, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Jack Clark, Samuel R. Bowman, Amanda Askell, Roger Grosse, Danny Hernandez, Deep Ganguli, Evan Hubinger, Nicholas Schiefer, and Jared Kaplan. 2022. Discovering Language Model Behaviors with Model-Written Evaluations. arXiv:2212.09251 [cs.CL]

- <https://arxiv.org/abs/2212.09251>
- [90] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67. <http://jmlr.org/papers/v21/20-074.html>
 - [91] Akshay Ravi, Aaron Neinstein, and Sara G Murray. 2023. Large language models and medical education: Preparing for a rapid transformation in how trainees will learn to be doctors. *ATS scholar* 4, 3 (2023), 282–292.
 - [92] Emma Roth. 2024. Google cut a deal with Reddit for AI training data - The Verge. <https://www.theverge.com/2024/2/22/24080165/google-reddit-ai-training-data>. (Accessed on 09/01/2024).
 - [93] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 39, 15 pages. <https://doi.org/10.1145/3411764.3445518>
 - [94] Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose Opinions Do Language Models Reflect?. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, Hawaii, USA, 29971–30004. <https://proceedings.mlr.press/v202/santurkar23a.html>
 - [95] Sebastin Santy, Jenny T. Liang, Ronan Le Bras, Katharina Reinecke, and Maarten Sap. 2023. NLPositionality: Characterizing Design Biases of Datasets and Models. arXiv:2306.01943 [cs.CL] <https://arxiv.org/abs/2306.01943>
 - [96] Jaromir Savelka, Arav Agarwal, Christopher Bogart, Yifan Song, and Majd Sakr. 2023. Can Generative Pre-trained Transformers (GPT) Pass Assessments in Higher Education Programming Courses? , 7 pages. <https://doi.org/10.1145/3587102.3588792>
 - [97] TA Schwandt. 1994. Constructivist, interpretivist approaches to human inquiry.
 - [98] Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. arXiv:2310.11324 [cs.CL] <https://arxiv.org/abs/2310.11324>
 - [99] Orit Shaer, Angelora Cooper, Osnat Mokryn, Andrew L Kun, and Hagit Ben Shoshan. 2024. AI-Augmented Brainwriting: Investigating the use of LLMs in group ideation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 1050, 17 pages. <https://doi.org/10.1145/3613904.3642414>
 - [100] Renee Shelby, Shalaleh Rismani, Kathryn Henne, Ajung Moon, Negar Rostamzadeh, Paul Nicholas, N’Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montréal, QC, Canada) (AI/ES '23). Association for Computing Machinery, New York, NY, USA, 723–741. <https://doi.org/10.1145/3600211.3604673>
 - [101] Arielle Michal Silverman. 2015. The perils of playing blind: Problems with blindness simulation and a better way to teach about blindness. (2015).
 - [102] Arielle M Silverman, Jason D Gwinn, and Leaf Van Boven. 2015. Stumbling in their shoes: Disability simulations reduce judged capabilities of disabled people. *Social Psychological and Personality Science* 6, 4 (2015), 464–471.
 - [103] David Silverman. 2004. Qualitative research : theory, method and practice. <https://api.semanticscholar.org/CorpusID:142594195>
 - [104] Mario Luis Small and Jessica McCrory Calarco. 2022. *Qualitative literacy: A guide to evaluating ethnographic and interview research*. Univ of California Press, California, USA.
 - [105] Robert Soden, Austin Toombs, and Michaelanne Thomas. 2024. Evaluating interpretive research in HCI. *Interactions* 31, 1 (2024), 38–42.
 - [106] Alex S. Taylor. 2021. Out there. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '21). Association for Computing Machinery, New York, NY, USA, 685–694. <https://doi.org/10.1145/1978942.1979042>
 - [107] Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. From humans to machines: can chatgpt-like llms effectively replace human annotators in nlp tasks.
 - [108] David R Thomas. 2017. Feedback from research participants: are member checks useful in qualitative research? *Qualitative research in psychology* 14, 1 (2017), 23–41.
 - [109] Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkar, and Graham Neubig. 2024. Do LLMs Exhibit Human-like Response Biases? A Case Study in Survey Design. *Transactions of the Association for Computational Linguistics* 12 (09 2024), 1011–1026. https://doi.org/10.1162/tacl_a_00685 arXiv:https://direct.mit.edu/tacl/article-pdf/doi/10.1162/tacl_a_00685/2468689/tacl_a_00685.pdf
 - [110] J. Van Maanen. 2011. *Tales of the Field: On Writing Ethnography, Second Edition*. University of Chicago Press, Chicago, USA. <https://books.google.com/books?id=V9hi269OD9cC>
 - [111] Pranav Narayanan Venkit, Sanjana Gautam, Ruchi Panchanadikar, Ting-Hao ‘Kenneth’ Huang, and Shomir Wilson. 2023. Nationality Bias in Text Generation. arXiv:2302.02463 [cs.CL] <https://arxiv.org/abs/2302.02463>
 - [112] Angelina Wang, Jamie Morgenstern, and John P. Dickerson. 2024. Large language models should not replace human participants because they can misportray and flatten identity groups. arXiv:2402.01908 [cs.CY] <https://arxiv.org/abs/2402.01908>
 - [113] Haochun Wang, Sendong Zhao, Zewen Qiang, Nuwa Xi, Bing Qin, and Ting Liu. 2024. Beyond Direct Diagnosis: LLM-based Multi-Specialist Agent Consultation for Automatic Diagnosis. arXiv:2401.16107 [cs.CL] <https://arxiv.org/abs/2401.16107>
 - [114] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. Taxonomy of Risks posed by Language Models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 214–229. <https://doi.org/10.1145/3531146.3533088>
 - [115] Rebecca Willis. 2019. The use of composite narratives to present interview findings. *Qualitative research* 19, 4 (2019), 471–480.
 - [116] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023. The rise and potential of large language model based agents: A survey.
 - [117] Yang Xiao, Yi Cheng, Jinlan Fu, Jiashuo Wang, Wenjie Li, and Pengfei Liu. 2023. How far are we from believable AI agents? A framework for evaluating the believability of human behavior simulation.
 - [118] 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* (Sydney, NSW, Australia) (IUI '23 Companion). Association for Computing Machinery, New York, NY, USA, 75–78. <https://doi.org/10.1145/3581754.3584136>
 - [119] Qisen Yang, Zekun Wang, Honghui Chen, Shenzi Wang, Yifan Pu, Xin Gao, Wenhao Huang, Shiji Song, and Gao Huang. 2024. PsychoGAT: A Novel Psychological Measurement Paradigm through Interactive Fiction Games with LLM Agents. arXiv:2402.12326 [cs.CL] <https://arxiv.org/abs/2402.12326>
 - [120] He Zhang, Chuha Wu, Jingyi Xie, Fiona Rubino, Sydney Graver, ChanMin Kim, John M Carroll, and Jie Cai. 2024. When Qualitative Research Meets Large Language Model: Exploring the Potential of QualiGPT as a Tool for Qualitative Coding.
 - [121] Zheng Zhang, Jie Gao, Ranjodh Singh Dhaliwal, and Toby Jia-Jun Li. 2023. VISAR: A Human-AI Argumentative Writing Assistant with Visual Programming and Rapid Draft Prototyping. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 5, 30 pages. <https://doi.org/10.1145/3586183.3606800>
 - [122] Zheyuan Zhang, Daniel Zhang-Li, Jifan Yu, Linlu Gong, Jinchang Zhou, Zhiyuan Liu, Lei Hou, and Juanzi Li. 2024. Simulating classroom education with llm-empowered agents.
 - [123] Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, and Maarten Sap. 2024. SOTOPIA: Interactive Evaluation for Social Intelligence in Language Agents. arXiv:2310.11667 [cs.AI] <https://arxiv.org/abs/2310.11667>
 - [124] Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics* 50, 1 (2024), 237–291.
 - [125] Michael Zimmer. 2020. “But the data is already public”: on the ethics of research in Facebook. , 229–241 pages.