

Not Everything That Counts Can Be Counted: A Case for Safe Qualitative AI

Stine L. Beltoft^[0009–0006–5412–0050]
and Lukas Galke^[0000–0001–6124–1092]

University of Southern Denmark
{stinelb,galke}@imada.sdu.dk

Abstract. Artificial intelligence (AI) and large language models (LLM) are reshaping science, with most recent advances culminating in fully-automated scientific discovery pipelines. But qualitative research has been left behind. Researchers in qualitative methods are hesitant about AI adoption. Yet when they are willing to use AI at all, they have little choice but to rely on general-purpose tools like ChatGPT to assist with interview interpretation, data annotation, and topic modeling – while simultaneously acknowledging these system’s well-known limitations of being biased, opaque, irreproducible, and privacy-compromising. This creates a critical gap: while AI has substantially advanced quantitative methods, the qualitative dimensions essential for meaning-making and comprehensive scientific understanding remain poorly integrated. We argue for developing dedicated qualitative AI systems built from the ground up for interpretive research. Such systems must be transparent, reproducible, and privacy-friendly. We review recent literature to show how existing automated discovery pipelines could be enhanced by robust qualitative capabilities, and identify key opportunities where safe qualitative AI could advance multidisciplinary and mixed-methods research.

Keywords: AI for Science · Mixed methods · Large language models

1 Introduction

In the past five years, artificial intelligence has made an undeniable impact on the practice of scientific research. The discourse surrounding AI in research has grown rapidly with students using ChatGPT to draft essays, scholars debating bias in large language models, and new guidelines on responsible AI for writing papers. Much of this development, however, has focused on quantitative methods. AI is routinely used to analyze medical images, classify tumors, run complex statistical models, wrangle large numerical datasets, and gather literature for surveys – everything from suggesting statistically sound question formats to simulating response patterns. Most recently, these efforts have culminated in fully-automated scientific discovery pipelines [36,13,22]. But even these pipelines remain fundamentally incomplete, as they lack a qualitative component.

Unlike quantitative research, where AI tools have become increasingly sophisticated and specialized, qualitative research currently relies on general-purpose

AI tools like ChatGPT which were not built for the demands of interpretive, contextual inquiry and come with severe limitations regarding reproducibility (models are not deterministic and may become unavailable), transparency (bias-prone black box decisions), and privacy (cloud-based processing of sensitive data). This mismatch is one of the reasons why AI use in qualitative research often raises skepticism. Moreover, available tools do not support the epistemological depth, reflexivity, or human-centered reasoning at the heart of qualitative inquiry.

This current state, however, reflects design limitations rather than inherent AI constraints. AI has far greater potential for qualitative research, if developed with safety (transparency, reproducibility, privacy) and interpretive goals in mind. Ideally, such a tool would identify recurring themes across dozens of interviews and compare qualitative narratives to quantitative trends or even assist in the ethical and context-sensitive collection of qualitative data. Such tools could fundamentally expand, not replace, the capacities of qualitative researchers.

Here, we argue that *if AI is to truly support scientific inquiry in its full breadth, it must do more than accelerate data processing. It must also engage with meaning, nuance, and complexity.* That requires understanding qualitative inquiry as requiring fundamentally different computational approaches and designing tools that serve its purposes from the ground up. After establishing the importance of qualitative research in Section 2, we structure our argument around the following key observations:

- Current practice: Researchers already use AI tools despite their limitations (Section 3)
- Contradictions in practice: AI seems to be used in qualitative research out of necessity, but simultaneously labeled as a limitation (Section 4)
- Structural inequalities: The most powerful AI systems are biased towards high-resource languages such as English and aligned with Western norms and values (Section 5)
- Dangers of replacement: Synthetic participants undermine scientific validity (Section 6)
- Quantification bias: AI for science and recent advances in automated scientific discovery pipelines exclude qualitative dimensions (Section 7)

After these considerations, we give concrete examples, for how AI could in fact support and, thereby, advance the qualitative research process (Section 8), before we conclude.

2 Why qualitative research matters

Qualitative research plays a vital role in scientific inquiry by offering tools and methods to explore meaning, context, and complexity, elements that often elude numerical representation. At its core, qualitative research is concerned with thick description, contextual meaning, lived experience, and interpretive reasoning[11,8]. Rather than isolating variables or seeking predictive patterns,

qualitative research attends to how people understand and construct their realities, often through narrative, contradiction, and ambiguity.

Not everything of value in research can be counted or measured. Many of the most pressing questions in contemporary science like those related to mental health, social change, identity, and political opinion, require approaches that embrace nuance and plurality. In these domains, meaning often emerges not from aggregate data, but from situated knowledge and subjective accounts [14,11].

Qualitative research is considered both constructivist and constructionist [23]: A constructivist approach focuses on how individuals create meaning based on their lived experiences [6], while a constructionist approach emphasizes how these meanings are shaped by broader cultural and social contexts [20]. These epistemological commitments stand in contrast to positivist assumptions: qualitative research does not aim to discover a single, objective truth, but instead seeks to produce transferable insights that are situated, interpretive, and often provisional [24].

The qualitative research process typically involves stages such as data collection (e.g., interviews, observations), transcription, and interpretive analysis [7]. Each of these stages requires paying close attention to language, context, and power, and often involves iterative reflection by the researcher. This complexity and reflexivity are not flaws, they are essential features of a methodology designed to engage with human meaning in all its depth.

Despite its importance, qualitative research is often under-supported in computational settings. As this paper will argue, AI tools developed for scientific research, which are currently culminating in fully-automated scientific discovery pipelines [13,22,36], must begin to take seriously the methodological and epistemic demands of qualitative inquiry, rather than treating it as an afterthought or reducing it to mere text processing.

3 AI in the qualitative research pipeline today

A small-but steadily growing set of systems is being built specifically for the interpretive tasks that define qualitative research (described in more detail below). Although most remain at prototype stage or are add-ons to established software, they illustrate what an explicitly qualitative AI ecosystem could look like.

Interviewbot [3] shows how a large language model can be tuned for semi-structured interviewing. The system generates follow-up questions in real time, aiming to preserve conversational flow while leaving the researcher in control of topic depth and ethical boundaries. Cody [29] tackles one of the most labour-intensive steps in qualitative analysis; coding. It combines rule-based logic with supervised machine learning to propose initial codes, surface disagreements, and log provenance so that researchers can audit every suggestion before acceptance.

Moving into observations, The Ethnobot [32] embeds a chatbot in Internet-of-Things devices to capture contextual anecdotes from users as they interact with smart-mobility services. By prompting short, situated reflections, it augments traditional field-notes with time-stamped micro-ethnographies. Journey-

Bot [16] demonstrates AI-assisted sense-making. The tool lets design-research teams feed interview excerpts into an interactive journey-map visualizer; a transformer model highlights convergent pain points and alternative storylines without flattening contradictory accounts.

Two mainstream qualitative analysis suites are beginning to natively integrate AI. MAXQDA’s [34] experimental “AI Assist” module uses large language models to suggest code–text matches and memo prompts, while NVivo AI [17] offers auto-summaries of interviews in multiple languages and proposes sentiment categories that analysts can accept, revise, or delete. Because both platforms store data locally (or on researcher-controlled servers), they can be configured for GDPR-compliant projects, an essential step for sensitive fieldwork.

Taken together, these examples signal a shift from re-purposing general NLP tools toward building software that recognizes qualitative epistemologies: context-sensitive questioning, researcher reflexivity, and transparent human-AI collaboration. Yet they also reveal gaps. Most current systems, operate on English-language corpora and struggle with multilingual or oral data, excel at first-pass coding or prompting but rarely support deeper theoretical integration, and remain locked inside proprietary ecosystems, limiting community scrutiny and method sharing. Bridging these gaps will require dedicated efforts, open-source benchmarks grounded in qualitative theory, and closer partnership between computer scientists and interpretivist researchers. Only then, AI can become a reliable partner in the nuanced work of qualitative inquiry.

4 Contradictions in practice: The ethics and use of AI in qualitative research

Despite growing interest in AI across disciplines, AI for qualitative research is deeply contested. Researchers routinely raise concerns about its use with fears of epistemic reductionism, replication of bias, and the erosion of interpretive integrity. These critiques appear across disciplines, often centered on the limitations of opaque large language models like ChatGPT [9,25]. Besides criticism regarding bias of the models, discussed in detail in Section 5, the use of such opaque models comes with concerns around the following three main aspects:

1. **Transparency.** The models are not openly available, their training data is unknown, and it is not clear to what extent web interfaces employ internal routing (e.g., relaying supposedly easy queries to smaller and weaker models).
2. **Reproducibility.** As these AI systems can simply cease being available (e.g., models being replaced by a newer version), it is almost impossible to ensure reproducibility when using commercial general-purpose AI systems.
3. **Privacy.** For the most powerful AI systems, requests need to be sent to central compute clusters, and could leak into the training data of the models. While agreements can be made, sensitive data is still most often processed in areas where different privacy regulations apply.

Indeed, papers on the ethical implications of generative AI in qualitative research have proliferated, with many authors arguing that such tools should be used cautiously, if at all [23].

Yet these ethical warnings have not prevented the widespread and growing use of these tools in practice. *Many qualitative researchers now incorporate generative AI into their workflows, using it for tasks such as interview summarization, preliminary coding, literature synthesis, and manuscript drafting [21,27,33,35].* Paradoxically, these uses are often acknowledged in limitations sections only to be disavowed moments later, framed as experimental, provisional, or morally suspect. The result is a kind of institutionalized ambivalence: AI is used out of necessity or convenience, while its legitimacy is publicly questioned. This tension points to a deeper structural contradiction in current research culture.

The gap between public disapproval and private adoption reveals a failure not of ethics, but of infrastructure. Researchers are turning to AI because the tools that would truly support qualitative inquiry do not yet exist. The only available systems, such as ChatGPT, were not designed for qualitative science, and their use reflects a lack of better alternatives rather than a wholehearted endorsement.

Meanwhile, industry sectors that could substantially benefit from qualitative AI, such as healthcare, UX design, education, and policy, remain hesitant to adopt these tools at scale. This hesitancy is not due to a lack of need or potential, but rather due to economic, reputational, and legal concerns. *There is uncertainty around data privacy, fear of backlash for using AI in domains tied to empathy or ethics, and a general lack of confidence in current tools' validity.* As a result, adoption remains cautious and fragmented, even in spaces where AI could augment the work of human-centered analysis.

The situation is further compounded by the lack of investment in building qualitative AI tools within computer science itself. The stigma around qualitative reasoning as subjective, unscalable, or unscientific has discouraged many technical researchers from engaging seriously with the epistemic goals of interpretive inquiry. While quantitative AI enjoys decades of methodological refinement and dedicated tooling, qualitative AI is often dismissed as mere “text generation” and left to general-purpose models that are ill-equipped for the nuances of fieldwork, narrative, and theory.

This has left the field in a cycle of critique without construction. Scholars highlight the risks of qualitative AI while quietly relying on it; developers avoid building qualitative tools due to perceived reputational risk or unclear metrics of success. The result is a landscape where qualitative research is both underserved and overexposed, used widely, defended rarely, and rarely built for from the ground up. To move forward, we must break this cycle by shifting the conversation from whether qualitative researchers should use AI to how we can design AI systems that actually reflect the interpretive, reflexive, and situated nature of qualitative research.

5 Structural inequalities and epistemic loss

The increasing use of AI in qualitative research raises not only methodological concerns, but also critical questions about whose knowledge is being represented, and whose is being erased [2,10] (*inter alia*). When lived experiences are replaced by model-generated outputs, there is a real risk that the perspectives of marginalized, underrepresented, or politically contested groups are lost in translation. The apparent neutrality and fluency of AI-generated text can mask a deeper structural problem: the reproduction of dominant narratives at the expense of dissenting or non-normative voices.

Large language models like ChatGPT are trained on enormous corpora of text from the internet, academic literature, news, and social media [28]. While this suggests some kind of democratic inclusivity, in practice it mirrors and amplifies the imbalances already present in those sources [4]. Groups with less access to publication channels, fewer mentions in data-rich environments, or voices that exist primarily in oral, localized, or situated contexts are less likely to be represented, or are represented inaccurately. This leads to an epistemic loss that privileges what is already visible, already legible, and already powerful.

AI systems that are not grounded in qualitative reasoning frameworks risk stripping away exactly the kind of positionality, contradiction, and contextual awareness that make qualitative research valuable in the first place. Without tools that account for situated knowledge and interpretive depth, we risk creating outputs that appear comprehensive but are, in fact, hollow, flattened representations of reality that obscure power dynamics rather than exposing them.

Trained on a vast body of human knowledge, technologies like ChatGPT are not a self that contains multitudes, but multitudes absent of a self [12]. By design, these tools cannot hold a single, describable point-of-view, and thus cannot offer the positionality necessary for trust, accountability, or critical engagement. In qualitative research, perspective is not a bug, it is a method. It is precisely through the recognition of voice, standpoints, and relational knowledge that qualitative inquiry makes power visible and subject to analysis.

When AI is used uncritically in place of human narratives, there is not only a risk of misrepresentation but also of epistemic loss: the gradual erasure of lived realities, local knowledge, and subaltern perspectives. Without intentional intervention, through the design of tools that preserve rather than overwrite human complexity, qualitative AI threatens to become a machine for amplifying existing inequalities, all while appearing neutral.

In this light, the development of qualitative AI is not merely a technical challenge, it is a political and ethical imperative. To support just and inclusive research, we must design AI systems that can recognize, preserve, and even foreground the contested, incomplete, and situated nature of knowledge itself.

6 The dangers of replacing qualitative inquiry

As the capabilities of generative AI systems continue to grow, so does their use in domains once considered the exclusive domain of human interpretation, notably

qualitative research. *An emerging trend is the use of models like ChatGPT to simulate participant responses, particularly in exploratory studies where human data is costly, time-consuming, or ethically sensitive to collect*, as discussed in [1,30], for instance. In such cases, researchers have begun experimenting with AI as a proxy for human subjects, treating model-generated text as replacements for real-world opinions and perspectives.

This practice raises serious concerns: First, language models are not participants. They are not agents with lived experience, emotional investment, or socio-cultural embeddedness. They generate plausible-sounding language by predicting patterns based on prior text data, but they cannot hold beliefs, be shaped by life events, or reflect on the world as social beings. This means that any insights they produce are simulations without grounding, textual echoes of discourse rather than situated narratives. Take, for instance, a model trained on data up to 2020 asked to represent Danish public opinion toward the United States. It may claim that Danes “generally have a favorable view of the U.S.,” reflecting sentiments prevalent at that time. However, this output ignores any subsequent political events, cultural shifts, or changes in public sentiment. The result is a temporally lagging, context-agnostic snapshot presented with the confidence of current truth, but based on outdated data and probabilistic assumptions. This kind of simulation carries several dangers. First, there is the risk of misrepresentation: the views of real communities are overwritten by generic, statistically averaged language. Second, there is false confidence: because AI-generated responses are fluent and coherent, they may appear more authoritative than they are. Third, and most concerning, is the erasure of lived realities: When researchers substitute actual participant voices for model outputs, they lose access to the richness, contradiction, and affect that define qualitative data.

The use of generative models as proxy for the human perspective fundamentally misunderstands what qualitative research is meant to do. Qualitative inquiry is not simply the collection of words: it is the practice of engaging with meaning in context, conversation, and relationship. It is attentive to temporality, voice, and positionality. These are dimensions that no current AI model can simulate with fidelity.

If we begin to treat models as equivalent to people, or model outputs as equivalent to data, we risk reducing qualitative inquiry to a kind of textual ventriloquism where synthetic language is mistaken for authentic experience. This is not only a methodological mistake but an ethical one, particularly in research contexts involving marginalized or vulnerable populations whose realities cannot and should not be approximated by statistical generalization.

Rather than replacing qualitative research, AI should be used to support it. The dangers of treating generative models as stand-ins for participants make clear that the future of AI in qualitative science depends not on mimicry, but on collaborative, grounded, and interpretive design. Only then can AI contribute meaningfully to the complexity that qualitative inquiry seeks to illuminate.

7 Quantitative bias in AI for science

Artificial intelligence has become deeply integrated into the infrastructures of contemporary scientific discovery, but its development has been shaped by a strong, and often unacknowledged, quantitative bias. The most prominent AI tools used in science today are built to serve data-centric workflows. *They perform tasks such as literature mining, hypothesis generation, statistical modeling, and simulation. From predicting protein structures [19] to automating survey design [18] and analyzing high-resolution images [15], these tools are optimized for speed, scalability, and measurable outcomes. These systems are typically designed to support quantitative inference. Their strength lies in identifying patterns, classifying data, and making predictions from large, structured datasets. They excel at streamlining standardized processes: optimizing experimental design, calculating statistical correlations, and drawing generalizable conclusions from massive data volumes. This emphasis reflects long-standing assumptions within both AI and scientific research, that knowledge is best when it is objective, replicable, and quantifiable.*

While this focus has undoubtedly advanced many fields, it has also created a significant asymmetry in how different forms of knowledge are supported. The vast majority of AI tools are designed for environments where variables are controlled, and data is clean. As a result, there is a striking lack of infrastructure for tasks fundamental to qualitative research: open-ended interpretation, coding of narrative data, theory-building, and contextual analysis. Qualitative research, which thrives on ambiguity, contradiction, and reflexivity, is often sidelined by systems that treat complexity as noise rather than insight.

This imbalance has broader epistemological consequences. When AI development in science disproportionately serves quantitative paradigms, it subtly reinforces the idea that interpretive or meaning-based inquiry is less scientific, less rigorous, or less important. Yet some of the most critical issues we face, how people experience injustice, how institutions shape behavior, how beliefs shift over time, are not well captured through statistical models alone. These are questions that require interpretive depth and contextual nuance. Moreover, this automation of quantitative reasoning often comes at the expense of interpretation. When large-scale pattern recognition replaces close reading, and when prediction supersedes understanding, we risk stripping research of its human-centered focus.

Interpretation is not a delay in the pipeline or an inefficiency to be optimized away, it is the space where meaning is constructed, where ethical questions are confronted, and where knowledge is made accountable to lived reality. Without AI tools designed to engage these interpretive processes, qualitative researchers are often forced to retrofit general-purpose systems, like generative language models, into workflows they were not built to support. These tools can be helpful in limited ways, transcribing interviews or summarizing texts, but they do not attend to the core epistemic practices of qualitative inquiry. The result is a patchwork approach that may satisfy surface-level efficiency but leaves the deeper work of meaning-making unsupported.

Recent work by recognizes similar limitations in current AI trajectories, proposing ‘Scientist AI’ systems that explain observations rather than merely act upon them [5] – an approach that aligns more closely with the interpretive goals of qualitative inquiry.

To move forward, the development of AI in science must recognize this imbalance and broaden its scope. A truly inclusive research pipeline would support not only the extraction of data, but also the construction of meaning. This requires an intentional investment in tools that respect and enable the complex, reflexive, and socially embedded nature of qualitative research.

8 Designing AI to support qualitative research

If AI is to meaningfully contribute to the future of qualitative research, it must be designed not to replace interpretive processes, but to support, scaffold, and extend them. This requires a fundamental reorientation toward systems that are responsive to the specific epistemologies, ethics, and workflows of qualitative inquiry. Rather than treating text as interchangeable data or analysis as reducible to classification, qualitative AI should engage with how meaning is produced, who it is produced by, and in what context. It must operate with the understanding that language is not only a medium of information, but also a medium of identity, ideology, and experience. There is growing potential to build AI tools that are not just assistants, but thoughtful collaborators in the research process. For example interview coding assistants could help researchers surface recurring themes across transcripts while remaining transparent about ambiguity and disagreement in codes. Or theory-building collaborators that might suggest relationships between concepts or offer alternative framings drawn from specified theoretical traditions, rather than from general web-trained embeddings. Even reflective tools that could prompt researchers to articulate their positionality, assumptions, or blind spots, making reflexivity a supported part of the analytic workflow rather than an afterthought.

Underlying these possibilities are a set of design principles that differ markedly from the assumptions built into most existing AI tools:

1. *Context-sensitivity* is key: systems should not treat language as decontextualized input but instead be aware of who is speaking, when, and in what setting.
2. *Temporal awareness* matters, especially in domains like social scientific or policy research, where attitudes and conditions change over time.
3. All AI used in qualitative workflows should be *human-in-the-loop* [26], ensuring that interpretation remains a human act, shaped by ethical judgment and domain expertise.
4. Rather than present as another tool to master, AI should *seamlessly extend* the self of the qualitative researcher [31].
5. Qualitative AI must embrace *non-reductive reasoning*: ambiguity, contradiction, and complexity are not bugs to be optimized away, but essential features of the data.

These design principles must be underpinned by the technical foundations that current commercial systems lack: transparency, reproducibility, and privacy protection. Unlike the black-box ChatGPT queries that researchers currently resort to, qualitative AI systems should provide clear trails of their reasoning processes (i.e., explainability built-in), allowing researchers to understand and critique how interpretations were generated. They must further support reproducible workflows where the same data processed with the same parameters yields consistent results, which is impossible with constantly-updating, non-deterministic models. Moreover, these systems must be designed to run on local infrastructure, ensuring that sensitive interview data, personal narratives, and confidential fieldnotes never leave the researcher’s control. This represents a fundamental departure from the cloud-based, privacy-compromising tools that researchers currently have no choice but to use when they turn to AI at all.

Going further, when qualitative AI systems are transparent and interoperable, they could finally enable the kind of interdisciplinary collaboration and mixed-methods integration that researchers have long wanted but struggled to achieve. Instead of forcing researchers to choose between quantitative precision and qualitative depth, properly designed systems could let them work seamlessly across both approaches. Imagine being able to trace how statistical patterns connect to narrative themes, with clear documentation of how each insight was generated and how they inform each other. This is a far cry from the current reality where quantitative and qualitative analyses typically happen in separate silos, often using completely different tools that have no interaction.

Rather than asking whether AI should be used in qualitative research, the more pressing question is: what kinds of AI are we willing to build? The future of qualitative AI depends on shifting from a culture of workaround to one of design, where tools are shaped in conversation with qualitative epistemologies, not in spite of them.

This vision aligns with recent calls for AI systems that prioritize understanding over action. Bengio et al.’s proposal for “Scientist AI” systems that explain the world from observations rather than simply generate text or take actions [5] offers a compelling framework and anchor for designing safe qualitative AI. Such systems, designed to build explanatory models with explicit uncertainty quantification, naturally complement qualitative research’s interpretive goals. Rather than replacing human judgment with automated decisions, these approaches could help researchers explore multiple theoretical framings, identify patterns across complex datasets, and articulate the limits of their interpretations, while maintaining the human researcher’s central role in meaning-making.

9 Conclusion

Qualitative research remains an essential pillar of scientific understanding. It allows us to ask not only what is happening, but why, how, and for whom. It reveals the meaning behind behavior, the complexity behind consensus, and the voices often excluded from dominant narratives. As artificial intelligence

continues to reshape research workflows, we must ensure that these technologies do not flatten science into what is merely countable, predictable, or extractable.

The current AI landscape reflects a lopsided investment in quantitative methods, an approach that, while powerful, is incomplete. To build a more inclusive and epistemically diverse research infrastructure, researchers, developers, and funders must turn their attention to qualitative AI: systems that support context, interpretation, reflexivity, and lived experience while maintaining the transparency, reproducibility, and privacy protections that rigorous research demands. These tools are not simply technical add-ons but necessary complements to a scientific practice that aims to understand the world in its full human complexity.

Good science does not rely on a single way of knowing. It draws from numbers, narratives, models, and meanings. If AI is to serve science in its fullest sense, it must be built not only for efficiency, but also for empathy, ambiguity, and ethical depth. Supporting qualitative research through thoughtful AI design is not just a methodological improvement, it is a commitment to better, fairer, and more human-centered knowledge.

Acknowledgements This research has been in parts supported by the Danish Foundation Models project. Parts of the text (in particular, abstract and introduction) have been refined against critique from Claude Sonnet 4.

References

1. Abdurahman, S., Atari, M., Karimi-Malekabadi, F., Xue, M.J., Trager, J., Park, P.S., Golazizian, P., Omrani, A., Dehghani, M.: Perils and opportunities in using large language models in psychological research. *PNAS nexus* **3**(7), pgae245 (2024)
2. Barikeri, S., Lauscher, A., Vulić, I., Glavaš, G.: Redditbias: A real-world resource for bias evaluation and debiasing of conversational language models. In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. pp. 1941–1955 (2021)
3. Beltoft, S., Schneider-Kamp, P., Askegaard, S.: Interview Bot: Can Agentic LLM’s Perform Ethnographic Interviews?.. In: *Proceedings of the 17th International Conference on Agents and Artificial Intelligence*. pp. 702–709. SCITEPRESS - Science and Technology Publications, Porto, Portugal (2025). <https://doi.org/10.5220/0013387800003890>, <https://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0013387800003890>
4. Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S.: On the dangers of stochastic parrots: Can language models be too big? In: *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. pp. 610–623 (2021)
5. Bengio, Y., Cohen, M., Fornasiere, D., Ghosn, J., Greiner, P., MacDermott, M., Mindermann, S., Oberman, A., Richardson, J., Richardson, O., et al.: Superintelligent agents pose catastrophic risks: Can Scientist AI offer a safer path? *arXiv preprint arXiv:2502.15657* (2025)
6. Bochner, A.P.: Unfurling Rigor: On Continuity and Change in Qualitative Inquiry. *Qualitative Inquiry* **24**(6), 359–368 (Jul 2018). <https://doi.org/10.1177/1077800417727766>, <https://journals.sagepub.com/doi/10.1177/1077800417727766>

7. Brinkmann, S., Tanggaard, L.: *Kvalitative Metoder: En Grundbog*. Hans Reitzels Forlag, København, 1 edn. (2010)
8. Denzin, N.K., Lincoln, Y.S., Sage Publications (eds.): *The Sage handbook of qualitative research*. Sage Publ, Thousand Oaks, Calif., 3. ed., [nachdr.] edn. (2009)
9. Friedman, C., Owen, A., VanPuymbrouck, L.: Should chatgpt help with my research? a caution against artificial intelligence in qualitative analysis. *Qualitative Research* (2024)
10. Gallegos, I.O., Rossi, R.A., Barrow, J., Tanjim, M.M., Kim, S., Dernoncourt, F., Yu, T., Zhang, R., Ahmed, N.K.: Bias and fairness in large language models: A survey. *Computational Linguistics* **50**(3), 1097–1179 (2024)
11. Geertz, C.: *Thick Description: Towards an Interpretive Theory of Culture*. In: *The Interpretation of Cultures*. New York (1973)
12. Gillen, A.L.G.: Can we trust ai in qualitative research? [Blog post] (2024), <https://www.insidehighered.com/opinion/views/2024/10/09/can-we-trust-ai-qualitative-research-opinion>, Inside Higher Education
13. Gottweis, J., Weng, W.H., Daryin, A., Tu, T., Palepu, A., Sirkovic, P., Myaskovsky, A., Weissenberger, F., Rong, K., Tanno, R., et al.: Towards an AI co-scientist. arXiv preprint arXiv:2502.18864 (2025)
14. Haraway, D.: Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies* **14**(3), 575 (1988). <https://doi.org/10.2307/3178066>, <https://www.jstor.org/stable/3178066?origin=crossref>
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition (2015). <https://doi.org/10.48550/ARXIV.1512.03385>, <https://arxiv.org/abs/1512.03385>, version Number: 1
16. Hwang, S., Kim, D.: JourneyBot: Designing a chatbot-driven interactive visualization tool for design research (2023). <https://doi.org/10.57698/V17I3.06>, <https://www.ijdesign.org/index.php/IJDesign/article/view/4674>, publisher: International Journal of Design
17. International, Q.: NVivo, <https://www.qsrinternational.com/nvivo-qualitative-data-analysis-software/home>
18. Jacobsen, R.M., Cox, S.R., Griggio, C.F., Van Berkel, N.: Chatbots for Data Collection in Surveys: A Comparison of Four Theory-Based Interview Probes. In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. pp. 1–21. ACM, Yokohama Japan (Apr 2025). <https://doi.org/10.1145/3706598.3714128>, <https://dl.acm.org/doi/10.1145/3706598.3714128>
19. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Židek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S.A.A., Ballard, A.J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A.W., Kavukcuoglu, K., Kohli, P., Hassabis, D.: Highly accurate protein structure prediction with AlphaFold. *Nature* **596**(7873), 583–589 (Aug 2021). <https://doi.org/10.1038/s41586-021-03819-2>, <https://www.nature.com/articles/s41586-021-03819-2>
20. Lee, C.G.: Reconsidering Constructivism in Qualitative Research. *Educational Philosophy and Theory* **44**(4), 403–412 (Jan 2012). <https://doi.org/10.1111/j.1469-5812.2010.00720.x>, <https://www.tandfonline.com/doi/full/10.1111/j.1469-5812.2010.00720.x>

21. Lee, V.V., Van Der Lubbe, S.C.C., Goh, L.H., Valderas, J.M.: Harnessing ChatGPT for Thematic Analysis: Are We Ready? *Journal of Medical Internet Research* **26**, e54974 (May 2024). <https://doi.org/10.2196/54974>, <https://www.jmir.org/2024/1/e54974>
22. Lu, C., Lu, C., Lange, R.T., Foerster, J., Clune, J., Ha, D.: The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery (2024). <https://doi.org/10.48550/ARXIV.2408.06292>, <https://arxiv.org/abs/2408.06292>, version Number: 3
23. Marshall, D.T., Naff, D.B.: The Ethics of Using Artificial Intelligence in Qualitative Research. *Journal of Empirical Research on Human Research Ethics* **19**(3), 92–102 (Jul 2024). <https://doi.org/10.1177/15562646241262659>, <https://journals.sagepub.com/doi/10.1177/15562646241262659>
24. Maxwell, J.A.: *Qualitative research design: An interactive approach*. Sage, 3 edn. (2013)
25. Morgan, D.L.: Exploring the use of artificial intelligence for qualitative data analysis: The case of chatgpt. *International journal of qualitative methods* **22** (2023)
26. Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Bascarán, J., Fernández-Leal, Á.: Human-in-the-loop machine learning: A state of the art. *Artificial Intelligence Review* **56**(4), 3005–3054 (2023)
27. Naeem, M., Smith, T., Thomas, L.: Thematic Analysis and Artificial Intelligence: A Step-by-Step Process for Using ChatGPT in Thematic Analysis. *International Journal of Qualitative Methods* **24**, 16094069251333886 (Apr 2025). <https://doi.org/10.1177/16094069251333886>, <https://journals.sagepub.com/doi/10.1177/16094069251333886>
28. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research* **21**(140), 1–67 (2020)
29. Rietz, T., Maedche, A.: Cody: An AI-Based System to Semi-Automate Coding for Qualitative Research. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. pp. 1–14. ACM, Yokohama Japan (May 2021). <https://doi.org/10.1145/3411764.3445591>, <https://dl.acm.org/doi/10.1145/3411764.3445591>
30. Rossi, L., Harrison, K., Shklovski, I.: The problems of LLM-generated data in social science research. *Sociologica* **18**(2), 145–168 (2024)
31. Schneider-Kamp, A., Godono, A.: The AI-extended professional self: user-centric AI integration into professional practice with exemplars from healthcare. *AI & SOCIETY* (Mar 2025). <https://doi.org/10.1007/s00146-025-02319-5>
32. Tallyn, E., Fried, H., Gianni, R., Isard, A., Speed, C.: The Ethnobot: Gathering Ethnographies in the Age of IoT. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. pp. 1–13. ACM, Montreal QC Canada (Apr 2018). <https://doi.org/10.1145/3173574.3174178>, <https://dl.acm.org/doi/10.1145/3173574.3174178>
33. Turobov, A., Coyle, D., Harding, V.: Using ChatGPT for Thematic Analysis (2024). <https://doi.org/10.48550/ARXIV.2405.08828>, <https://arxiv.org/abs/2405.08828>, version Number: 1
34. VERBI: MAXQDA, <https://www.maxqda.com/>
35. Xiao, Z., Yuan, X., Liao, Q.V., Abdelghani, R., Oudeyer, P.Y.: Supporting Qualitative Analysis with Large Language Models: Combining Codebook with GPT-3 for Deductive Coding (2023). <https://doi.org/10.48550/ARXIV.2304.10548>, <https://arxiv.org/abs/2304.10548>, publisher: arXiv Version Number: 1

36. Yamada, Y., Lange, R.T., Lu, C., Hu, S., Lu, C., Foerster, J., Clune, J., Ha, D.: The AI Scientist-v2: Workshop-level automated scientific discovery via agentic tree search. arXiv preprint arXiv:2504.08066 (2025)