sociological science

Is it Okay to Use Generative AI in Sociological Research?

Oscar Stuhler,[a] Dustin S. Stoltz,[b] AJ Alvero,[c] Marshall A. Taylor[d]

a) Northwestern University; b) Lehigh University; c) Cornell University; d) CREST, ENSAE, Institut Polytechnique de Paris

Abstract: This article presents the results of a survey of what uses of generative AI (GenAI) sociologists consider acceptable across different aspects of the research process, including literature review and ideation, data collection and analysis, writing, and peer review. Despite some areas of agreement, sociologists diverge considerably in their views on the acceptability of GenAI use in research. Many of the 19 practices we asked about were deemed generally unacceptable by a significant share of respondents even when disclosed in resulting publications, while an often similarly sized share deemed them acceptable even without disclosure. For example, 40 percent consider it unacceptable to generate an outline for an article with GenAI while for 34 percent doing so is acceptable and need not be disclosed. Few sociologists categorically reject the use of GenAI in research, with only 5 percent finding all uses unacceptable. GenAI use is more acceptable to quantitative researchers, non-native English speakers, and men. Although 82.2 percent of sociologists support journals' adoption of mandatory disclosure statements on GenAI use, only 22 percent trust such statements. We reflect on these findings and discuss their broader implications.

Keywords: generative AI; methodology; computational sociology; survey; large language models

Reproducibility Package: A replication repository for this article can be found at https://osf.io/fpjzn. The data for this article are hosted on the Harvard Dataverse at https://doi.org/10.7910/DVN/HNKOSO.

Citation: Stuhler, Oscar, Dustin S. Stoltz, AJ Alvero, and Marshall A. Taylor. 2026. "Is it Okay to Use Generative AI in Sociological Research?" Sociological Science 13: 1382-1403.

Received: August 21, 2026

Accepted: September 29, 2026

Published: November 17, 2026

Editor(s): Stephen Vaisey, Jeremy Freese

DOI: 10.15195/v13.a52

Copyright: © 2026 The Author(s). This open-access article has been published under a Creative Commons Attribution License, which allows unrestricted use, distribution and reproduction, in any form, as long as the original author and source have been credited.

RECENT surveys suggest that a large portion of sociologists regularly use or have at least tried using generative artificial intelligence (GenAI) for their research (Alvero et al. 2026; Lyttelton, Massenkoff, and Wilmers 2026). Over the past years, a quickly evolving conversation has emerged on how GenAI might be used to expand the sociological toolkit, spanning quantitative (e.g., Boelaert et al. 2025; Broska, Howes, and Loon 2025; Karell, Sachs, and Barrett 2025; Law and Roberto 2025; Stuhler, Ton, and Ollion 2025; Waight et al. 2025) and qualitative (e.g., Dunivin 2025; Than et al. 2025; Abramson et al. 2026; Ibrahim and Voyer 2026) modes of inquiry. Although discussed less publicly, many sociologists also use GenAI for everyday aspects of research, such as writing, literature review, brainstorming ideas, coding assistance, or analyzing data (Alvero et al. 2026). An analysis of sociology articles' abstracts estimates that in late 2025, as many as 25 percent of the sentences are large language model (LLM)-modified (Cao 2026). And, some sociologists have begun to explore the capacity of GenAI "agents" to assist or replace humans in most parts of the research process (Engzell and Wilmers 2026; Kozlowski 2026; Lee 2026; Zhang 2026).

Alongside the incorporation of GenAI into sociological research, there is also concern about its impact on the scientific community. An open letter published in

December 2025 and signed by more than 400 researchers from the social sciences and humanities argues that "qualitative research should remain a distinctly *human* practice" in which GenAI must have no role (Jowsey et al. 2025:1; for the same argument in management, see Nguyen and Welch 2026). A joint statement by editors of several leading social science journals (including this one) warns that the use of GenAI in the production of academic articles threatens an already strained peer review system by accelerating submission rates "beyond the breaking point" (Munger et al. 2026:1). Confronted with "a large volume of AI-generated slop articles," a major preprint server for social science research issued a 90-day moratorium for articles on AI topics (Cohen 2025). Subsequently, it adopted a policy that forbids, among other practices, LLM-based simulation of human subjects (Cohen 2026). Alvero et al. (2026) find that only a minority of the sociologists they surveyed in 2025 believed that GenAI will have a positive effect on the discipline, while many worry about bias it could introduce into research. Furthermore, concern about GenAI's negative social and environmental impacts is widely shared among sociologists (Jowsey et al. 2025; Alvero et al. 2026; Timmermans and Tavory 2026).

This raises challenging questions for researchers and editors. Is my use of GenAI acceptable to my readers and colleagues? What obligations do I have to document and acknowledge my use? Is my refusal to use GenAI shared by my peers and is it putting me at a disadvantage? What uses of GenAI, if any, should journals allow? What advice should be given to graduate students? Establishing norms for GenAI use and documentation is crucial for the discipline. Although this conversation is underway, it has so far lacked systematic evidence on what applications of GenAI sociologists consider acceptable and under what circumstances.

Here, we present the results of a survey that asked authors about the acceptability of GenAI use across an array of sociological research practices, including literature review and ideation, data collection and analysis, writing, and peer review. For this, we contacted a sample of 3,425 authors who published in 50 sociology journals in the past five years. Among those contacted, 300 authors responded, 259 of them sociologists.

We find that despite some areas of agreement, sociologists are divided over the use of GenAI. Many of the research practices we asked about were deemed generally unacceptable by a large share of respondents even when disclosed in resulting publications, while an often similarly sized share deemed them acceptable even without disclosure. For example, 40 percent consider it unacceptable to use GenAI to generate a first outline for a article while 34 percent said that doing so is acceptable and does not need to be disclosed. Similarly, using GenAI to summarize peer reviews and "discuss" revision strategies is unacceptable to 37 percent, while for 40 percent it is acceptable and need not be acknowledged. Across scenarios, our analyses suggest that GenAI use is perceived as more acceptable by quantitative researchers, non-native English speakers, and men. The vast majority of respondents want journals to adopt mandatory AI disclosure statements (82.2 percent). However, only 22 percent trust such statements. We reflect on the implications of our findings for the discipline in our concluding section.

Background

Adoption of GenAI in Sociological Research

A methodological discussion has emerged on how–and whether–to use LLMs for sociological research. Although initial contributions focused on text classification (Bonikowski, Luo, and Stuhler 2022; Do, Ollion, and Shen 2022; Wankmüller 2024), *generative* LLMs—models that produce text, generally referred to as *generative AI*[1]—have opened the technology to a variety of other applications (Bail 2024; Davidson 2024; Davidson and Karell 2025). Although we cannot offer a systematic overview, this includes simulating survey responses (Boelaert et al. 2025; Broska, Howes, and Loon 2025; Kozlowski and Evans 2025; Kim and Lee 2026) and social interactions (Karell, Sachs, and Barrett 2025), encoding image or multimodal content (Law and Roberto 2025; Davidson 2026), extracting granular information from text (Stoltz 2025; Stuhler, Ton, and Ollion 2025), generating adaptive stimuli for online experiments (Argyle et al. 2023), qualitative content coding (Dunivin 2025; Than et al. 2025; Abramson et al. 2026; Ibrahim and Voyer 2026), and formalizing narrative structure (Waight et al. 2025).

There is also growing evidence that some sociologists use GenAI for procedural, "every-day" aspects of research. Recent surveys suggest that GenAI is used for brainstorming research ideas, literature review, programmatic coding, and writing. For instance, Alvero and colleagues (2026) find that in early 2025, about a third of sociologists have used GenAI for writing while roughly a fifth indicate having used it for planning and data analysis. Other studies suggest that adoption may be even higher: a survey conducted in March 2026 by three sociologists on behalf of Anthropic estimates that 78 percent of "quantitative" sociologists have used GenAI for their research (Lyttelton, Massenkoff, and Wilmers 2026)—though this number is likely to be inflated given that, as the authors caution, participation in the survey was incentivized with access to GenAI products. A text analysis suggested that in the second half of 2025, almost a quarter of the sentences in sociology articles' abstracts appeared LLM-modified (Cao 2026), though the article reports the rate to be considerably lower for mainstream journals.

The growing accessibility of AI "agents" adds another dimension to this discussion (Zhang 2026). Unlike chatbot interfaces, agents can execute operations on behalf of users such as writing and reading files, running computer scripts, searching the web, and invoking user-written skills or heuristics. Engzell and Wilmers introduced "an automated workflow that generates a quantitative social science article from an initial prompt" and a dataset using agents (2026:2).[2] Others suggest that researchers co-develop a research plan with an agent, before any tasks are executed (Lee 2026; also see Kozlowski 2026; Zhang 2026; outside of sociology, see Lu et al. 2026). The common principle among these frameworks is that researchers write reusable procedural heuristics that serve as explicit instructions for the agents—say, for how to ask questions when developing a research plan, or how to compute marginal effects for a regression model.

Concerns about GenAI Use in Sociology

Sociologists' concern about GenAI's integration into research is less citable than the methodological contributions discussed above, but it is certainly part of the larger discourse at conferences, faculty meetings, and discussion threads on social media. Still, several statements express worry about its impact on scholarship and opposition to certain practices. Most notably, *Qualitative Inquiry* published an open letter signed by 419 researchers from the social sciences and humanities arguing that "use of GenAI is inappropriate in all phases of reflexive qualitative analysis" (Jowsey et al. 2025:2). GenAI, the letter argues, "is incapable of meaning-making," (p. 1) and therefore cannot be reflexive. More specifically, GenAI is predisposed to reinforce dominant paradigms, biases, and language patterns, leading to the "quieting of marginal voices and practices, including those of critical scholars." (p. 2; for commentary on the letter, see Ibrahim, Nelson, and Voyer 2026; Timmermans and Tavory 2026).

Meanwhile, a joint statement by editors of leading social science journals highlights the strain that the technology has put on the peer review system (Munger et al. 2026:1). This includes an increase in low-quality article submissions (also noted by others, see Cohen 2025), many of which superficially appear to be high-quality articles. The statement makes the case that while AI can advance research, "peer review is an essential and irreplaceably human component of social science." (p. 2) Assessing future scenarios, the authors argue that social scientists should increase their involvement with the evaluation of research.

More broadly, sociologists are contributing to the critical study of AI with research on its political economy, environmental impacts, social harms, or policy implications (see, e.g., Joyce et al. 2021; Zajko 2022; Joyce and Cruz 2024; Law and McCall 2024; Bender and Hanna 2025; Xiao and Verdegem 2025; Lupton 2026; Vertesi et al. 2026). This critical scholarship is bound up with whether and how we ought to use GenAI. Prior survey evidence suggests that sociologists are highly concerned about the technologies' broader impact: four in five sociologists report being "concerned" or "very concerned" about the technologies' social and environmental costs (Alvero et al. 2026). For some, this is a key reason for rejecting the technology outright (Jowsey et al. 2025).

Regulation of GenAI in Sociological Research

As of late July 2026, the *American Sociological Association* has, to our knowledge, not put forward a statement or policy on AI use in sociological research. To get a systematic overview of the field, we manually examined the websites of 50 sociology journals. Most journals do not explicitly mention AI use on their websites (42 percent). About a third (36 percent) have statements about the permissibility of AI use or on how to disclose it. We should note, however, that few of these journals had *detailed* rules. Meanwhile, 22 percent do not state any principles on AI use but refer to their publisher's general policy (for details and limitations, see Appendix J of the online supplement).

We did not contact journals for this study, and discussions around GenAI regulation appear to be very much underway among editors (see discussion in Munger

et al. 2026). For instance, in their annual report to the ASA, the editors of the *American Sociological Review* note that they spent considerable time reflecting on the issues posed by GenAI and are eagerly awaiting guidance from the ASA (Cort et al. 2026). Nonetheless, it appears that GenAI use in sociology is currently relatively unregulated or subject to rules set by publishers. We take this as a point of departure for our empirical analysis: what uses of GenAI do *sociologists* consider acceptable, and what could be consensus conventions for the discipline?

Data and Sampling

We aim to survey the population of those contributing to recent sociological scholarship. To do so, we created a sampling frame that includes all authors who published in 50 sociology journals (see Appendix A of the online supplement) in the past five years drawn from the Web of Science. We emailed a random sample of authors in April to June 2026, inviting them to take a survey on GenAI use in sociological research. In total, we invited 3,425 authors. We gave respondents the option to take the survey either via a personalized or an anonymous link. In total, 300 completed or nearly completed the survey, yielding a response rate of 8.7 percent.[3] Of these, 259 respondents (86.3 percent) answered "yes" when asked whether they consider themselves sociologists. Additional details on our sampling strategy and recruitment materials are provided in Appendix A of the online supplement. The survey data and codebook are available on the Harvard Dataverse (Stuhler et al. 2026).

Our survey is based on voluntary participation, so there is likely some self-selection bias in our sample. Descriptive statistics suggest that our survey was taken by a heterogeneous group of scholars with regard to methodological orientation, country of residence, gender, career status, GenAI use frequency, and subfield, among other factors (see Table A.1 in Appendix A of the online supplement; also see Panel B of Fig. 3). However, those with more developed stances toward GenAI were likely more inclined to take part in our survey. Such bias is hard to eliminate in voluntary, not institutionally mandated surveys. Although incentivizing participation may have increased our response rate, it would likely have introduced different forms of bias. Still, our results must be read with this important caveat in mind, and we discuss its potential implications for our findings along with other limitations further below.

Analysis

To contextualize our findings, we briefly note the state of GenAI adoption in sociological research. Twelve percent of those who identify as sociologists are *daily* users of GenAI, and another 21 percent use it at least *weekly* for research purposes. Meanwhile, about a quarter (23 percent) of sociologists say they *never* use GenAI for research (see Table A.1 in the online supplement). These numbers are quite similar to those found by Alvero et al. (2026) a year ago, suggesting that levels of GenAI use for sociological research have remained relatively constant and may be lower than some recent publications suggest (Cao 2026; Lyttelton, Massenkoff, and

Wilmers 2026). We offer a systematic time-trend analysis in Appendix I of the online supplement, which confirms that GenAI adoption has remained largely unchanged, with a stable group who avoid the technology entirely.

Uses of GenAI and Their Acceptability

What uses of GenAI, if any, do sociologists consider acceptable? We presented 19 scenarios for the use of GenAI in sociological research and asked respondents whether they consider these *generally not acceptable*, *acceptable if disclosed in resulting publications or presentations*, or *acceptable without disclosure*. Our survey stated that by "GenAI" we refer to models such as those of the GPT family (see Appendix B of the online supplement for details). Answers to these questions should therefore be interpreted as reflecting opinions on the use of proprietary, nonlocally executable language models. We recognize that some respondents may prefer other models/modalities—a point we return to—however, this is what is currently the most common mode of using GenAI and one that is likely familiar to all respondents. Scenarios were grouped into four research stages: *Literature review and ideation*, *collecting and analyzing data*, *writing manuscripts*, as well as writing and responding to *peer reviews*. Item wordings are presented in Table 1.

Figure 1 shows the distribution of the responses for each scenario, only for respondents who identify as sociologists. In the top panel, we see that 66 percent of sociologists think that using GenAI for finding literature is acceptable and need not be acknowledged. Meanwhile, 21 percent consider this practice acceptable, but only if disclosed, and 12 percent find this generally unacceptable. These percentages move toward lower acceptance rates as uses become more substantial. A quarter of sociologists (25 percent) consider it unacceptable to "discuss" research with a GenAI tool to identify gaps in the existing literature, and almost 42 percent consider it unacceptable to use GenAI tools to summarize articles and extract relevant information from them.

We also asked about GenAI use for data collection and analysis. Most sociologists (82 percent, cumulatively) consider it acceptable to use GenAI for generating code for data visualizations. This share drops considerably for more immersive uses of GenAI in data analysis: a workflow in which a researcher provides survey data to a GenAI tool, specifies variables to analyze, and has the tool run the statistical analysis is considered acceptable by only 52 percent (cumulatively). A thin majority of sociologists are open to using chatbots to conduct interviews (55 percent, cumulatively) so long as subjects are informed. What stands out in comparison with other stages of the research process is that "acceptable if disclosed" was chosen more often for the data collection and analyses practices we asked about. This suggests that sociologists value transparency here more than in other aspects of research.

For writing, about half of sociologists (49 percent) believe it is acceptable to use GenAI to "polish a paper," specifically "adopting suggestions to refine word choice, restructure sentences, adjust tone, and correct grammatical errors." Additionally, 32 percent consider this acceptable if it is acknowledged. However, almost one in five respondents (17 percent) rejects such practices. Sentiments shift toward less acceptance as we increase the level of GenAI involvement in the writing process.

Table 1: GenAI Use Scenarios.

Item	Scenario text
	Literature Review
Discuss research ideas	A researcher "discusses" research with a GenAI tool to identify gaps in the existing literature that could motivate a new study.
Search for literature	A researcher uses a GenAI tool to search for relevant articles during their literature review.
Summarize articles	A researcher uses a GenAI tool to summarize articles and extract relevant information from them.
	Data Collection and Analysis
Conduct interview (undisclosed)	A researcher uses a GenAI tool to conduct interviews. Respondents are not told they are responding to a chatbot.
Conduct interview (disclosed)	A researcher uses a GenAI tool to conduct interviews. Respondents are told they are responding to a chatbot.
Organize fieldnotes	A researcher uses a GenAI tool to organize daily ethnographic fieldnotes and create short summaries for each day.
Code qualitative interviews	A researcher creates a codebook based on a subset of interviews. They then input the remaining interviews along with the codebook into a GenAI tool and request that it apply codes to each transcript.
Write code for visualization	A researcher describes the type of visualization they would like to use to a GenAI tool and requests that it return code to create the visualization.
Run survey analysis	A researcher provides survey data to a GenAI tool, specifies variables to analyze, and has the tool run the statistical analysis and produce output tables. After reviewing output, the researcher includes the table in their article.
	Writing
Polish paper	A researcher uses a GenAI tool to polish an article draft. This includes adopting suggestions to refine word choice, restructure sentences, adjust tone, and correct grammatical errors.
Generate outline for article	A researcher uses a GenAI tool to generate an outline, which they then use to write a full article.
Produce first draft	A researcher prepares an outline and notes, then uses a GenAI tool to produce a first draft from these materials, which they subsequently revise.
Solve writer's block	A researcher encounters a momentary bout of writer's block. They use a GenAI tool to get suggestions for the next paragraph.
Give feedback on article	A researcher uses a GenAI tool to get substantive feedback on a draft. They rewrite their article based on the feedback.

Table 1: *(Continued)*

Peer Review	
Summarize review	A researcher pastes peer reviews they received into a GenAI tool to summarize key points.
Summarize review, strategize on revision	A researcher pastes peer reviews they received along with their original article into a GenAI tool. In "discussion" with the model, they develop a revision strategy.
Discuss other's article's content	A reviewer "discusses" an article's content and methods with a GenAI tool when drafting their review.
Paste other's article and discuss	A reviewer pastes a article into a GenAI tool to "discuss" its strengths and weaknesses. Based on this conversation, they write their review.
Paste other's article and write review	A reviewer pastes a article into a GenAI tool and asks the model to produce an initial review. They subsequently edit the model's review before submitting it.



Figure 1: Perceived acceptability of GenAI use scenarios. *Note:* The percentages represent the share of respondents in each category. Items offered a "don't know/can't tell" response, which is reflected in the percentages but was rarely chosen (between 0 percent and 6.6 percent, see Appendix C of the online supplement).

About a third (32 percent) of sociologists consider it unacceptable to use GenAI to overcome writer's block by asking GenAI for suggestions for the next paragraph. And 61 percent reject using GenAI to produce a first draft based on notes. A noteworthy observation here is that rewriting an article based on GenAI feedback is more acceptable to sociologists than using GenAI for more substantial writing.

Finally, we asked about employing GenAI in the review process—both when addressing reviews as an author and when reviewing other's work. Sociologists are divided on whether it is acceptable to use GenAI to address reviews, though a majority thinks that this is legitimate when acknowledged. Meanwhile, sociologists

are highly opposed to reviewers employing GenAI, especially when this involves sharing others' work directly with a model. A majority (56 percent) considers it unacceptable for reviewers to discuss the content of an article with GenAI, and 76 percent are opposed to pasting an article into a GenAI model. Only 10 percent are open to having GenAI create a first draft of a review.

Overall, few sociologists categorically reject using GenAI in research, with only 5 percent finding all uses unacceptable. Sociologists are relatively open to GenAI being employed for literature search, coding, and low-level writing tasks. Meanwhile, the vast majority of sociologists reject GenAI use for reviewing others' work, expansive writing tasks, and manipulative modes of data collection. Still, there were many practices for which respondents were divided, such that large shares accept them without disclosure while many others find them generally unacceptable.

What is the dimensionality of these opinions? To test this, we examine the correlations between all 19 items, where responses were encoded as continuous (assigning 0 for *generally unacceptable*, 1 for *acceptable if disclosed*, and 2 for *acceptable without disclosure*). Associations between items are strong, with an average pairwise correlation of 0.39. In a principal component analysis based on complete cases ($N = 216$), the first dimension accounts for 45.1 percent of the variation, while subsequent dimensions account for only 7.3 percent and 5.7 percent, respectively (see Appendix C of the online supplement for details). Because of our small sample size, the variance accounted for by the first dimension is likely inflated, and the population-level component structure is likely to be somewhat less clearly unidimensional. Still, this suggests that while there may be factions that reject or accept certain sets of practices, heterogeneity in opinions on the acceptability of GenAI use is primarily structured along a single division line. If this is the case, what arguments and views stand behind these positions?

Open-Ended Responses: Themes and Arguments

About half of the respondents elaborated on their responses via an open text entry, giving us the chance to characterize positions and identify several recurring themes. In Figure 2, we present selected quotes, together with the respective respondent's position on a scale measuring acceptance of GenAI. The scale was constructed by averaging responses across the 19 scenarios discussed above. A value of 0 means that someone considered all 19 practices as unacceptable while a value of 2 would indicate that all practices were found to be acceptable without disclosure, though the observed range is between 0 and 1.94 ($mean = 0.89$; $sd = 0.49$). We opted for this basic approach to scale construction because it simplifies interpretation (factor analysis yields substantively identical results, see Appendix C of the online supplement). Quotes in Figure 2 were chosen to represent ideas expressed by multiple respondents.

At the low end of the scale, we find scholars who more or less categorically reject using GenAI. This view does not distinguish between use cases and was often justified with GenAI's negative environmental and social impacts (see Q2, Q5, and Q4 in Fig. 2). Views on this end of the spectrum were often expressed

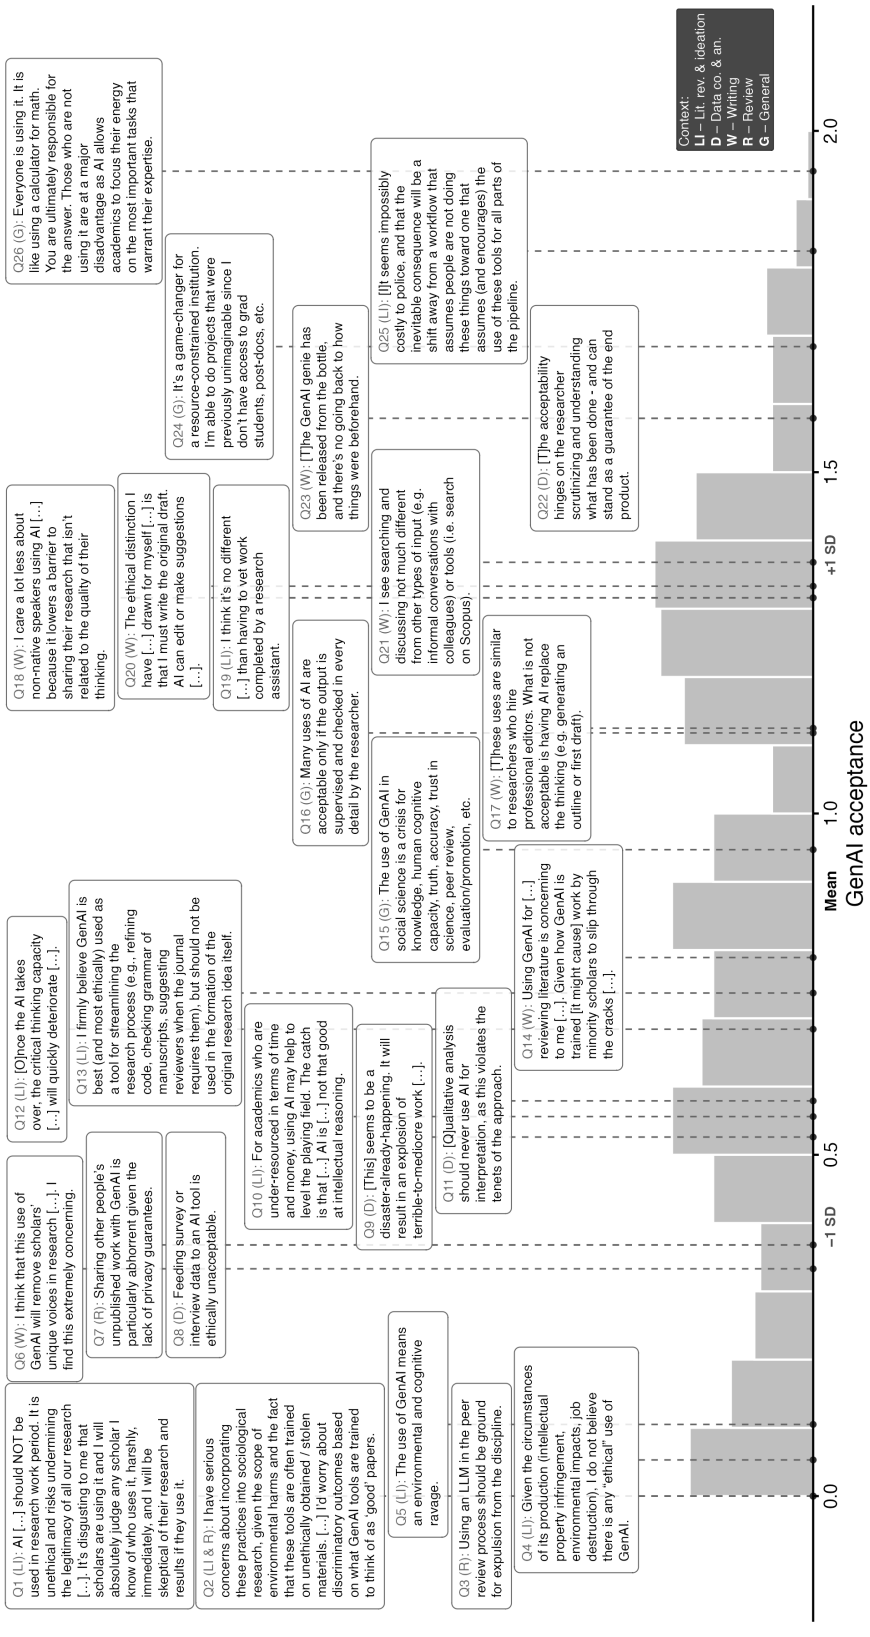


Figure 2: Quotes from open text entry. *Note:* The bottom of the figure shows a histogram of a scale measuring individuals' perceived acceptability of GenAI use in sociological research. It was constructed by averaging across 19 items shown in Figure 1. Nonresponses for single items were treated as missing and did not factor into the measure. The survey offered opportunities for open text entry after each set of scenarios as well as at the end of the survey (see the figure legend).

quite emphatically (Q1 and Q3), and those who reject using GenAI were generally somewhat more likely to elaborate on their views.

At the other end of the spectrum, scholars typically reason that GenAI is "a tool, no different from any other," as one person put it. Respondents stressed that the researcher maintains responsibility for the quality of the scholarship: cited sources must be verified, code must be reviewed, and results must be checked (Q26 and Q22). As long as the final research product meets scientific standards, however, using GenAI is legitimate and does not need any explicit acknowledgment in most cases. Several respondents highlighted that using GenAI enhances their productivity, was a "game changer for a resource-constrained institution" (Q24), and some of these scholars also noted that the increasing adoption of GenAI appears to be inevitable (Q25 and Q23).

Between these poles, we find that scholars search for principled ways to distinguish between what is acceptable and what is not. Some recurrent arguments were as follows: Most uses are acceptable as long as everything is disclosed. Many uses are acceptable as long as GenAI's work or suggestions are thoroughly checked (Q16). Use of GenAI in quantitative analysis is more acceptable than its use in qualitative data analysis (Q11). Sharing human subject data with GenAI—whether interviews or survey responses—is unethical (Q8). GenAI may be used for editing, but original drafts should be written by the researcher (Q20). A reviewer sharing other's unpublished work with GenAI is unacceptable (Q7). Many uses of GenAI are not unethical per se, but will result in low-quality research.

We want to stress that these are not consensus positions. In fact, what stood out more than common themes is the *considerable heterogeneity* in the rules and heuristics scholars articulated. For instance, several scholars noted that original ideas should come from researchers, but that GenAI may be a useful tool for executing them (Q13). Meanwhile, others thought that using GenAI to explore ideas at an early stage is one of its prime uses, but that execution requires much human involvement.

Those who find the use of GenAI in a specific scenario acceptable often reason with analogies (Q17, Q19, and Q21): brainstorming with GenAI is no different from talking to a colleague; literature research with GenAI is no different from other search engines; having GenAI generate code is no different from taking code from Stack Overflow; working with GenAI is not much different from working with a research assistant; having GenAI edit text is not different from working with a professional editor or advanced spell checker software. Besides, several respondents found GenAI acceptable because it helps level the playing field for those who do not have English as their first language or the institutional resources to hire research assistants or post-docs (Q10, Q18, and Q24).

Meanwhile, among those who did not categorically reject use of GenAI, many still stated concerns about its negative effects on the discipline. These respondents often expressed worry that the increasing adoption of GenAI will diminish the quality of sociological scholarship (Q9 and Q15), and cause a deprivation of scholars' skills and critical thinking capacities (Q5, Q12, and Q15). Finally, scholars also expressed concern about the homogenizing effects GenAI tools can have on scholarship (Q6), as well as about bias it can introduce into research and ideation at the cost of disadvantaged groups (Q2 and Q14).

Who Accepts GenAI Use?

Are perceptions of GenAI's acceptability associated with researcher characteristics? In Panel A of Figure 3, we present the effects of an ordinary least squares (OLS) regression predicting the standardized GenAI acceptance score (the model table and supplementary analyses are provided in Appendices D through F of the online supplement). Given our small sample size, these analyses should be seen as exploratory, rather than as tests of prespecified hypotheses. The effect estimates carry considerable uncertainty. Note, furthermore, that we use coarse categories for some variables (e.g., a binary notion of gender) to preserve the anonymity of our respondents in the dataset.

The strongest predictor of GenAI acceptance is preferred research methodology. Those who say they typically use quantitative methods are, on average, 0.69 standard deviations more accepting of GenAI use than those who typically use qualitative methods, with mixed-methods users roughly in the middle between these groups.

A second predictor of GenAI acceptance is researchers' native language. Specifically, we estimate that those who report having English as their first spoken language are 0.31 standard deviations less accepting of GenAI use. Supplementary analyses do not indicate that this is driven by our questions about writing (see Appendix E of the online supplement). Instead, non-native English speakers find GenAI use more acceptable in all aspects of the research process.

Third, we find that those who identify as men are more accepting of GenAI use than those who do not, aligning with prior research (Haimson et al. 2025; Chakravorti et al. 2025; Alvero et al. 2026; Cranney, Delecourt, and Koning 2026). Although there is no conclusive mechanism behind these differences, recent work suggests that women are judged more harshly for using GenAI than men (Gai, Hou, and Tu 2025; Chatoo 2026).

Our analyses do not suggest pronounced differences by race. Although our models indicate that those residing in North America perceive GenAI uses as less acceptable than those in the rest of the world, this difference disappears once we control for having English as native language (see Appendix D of the online supplement). We should note that our ability to test regional differences is relatively limited—the vast majority of our respondents come from the United States or Europe. We also find some indication that sociologists are less accepting of GenAI than their collaborators, though the latter is a heterogeneous group, and the effect is not statistically significant in any of our models.

Age does not have a statistically significant effect on GenAI acceptance, even when excluding tenure status ($r = 0.5$ with age) from our model (see Appendix D of the online supplement for alternative model specifications). Our age variable has some measurement error as we collapsed our bracketed measure (see Table A.1 in the online supplement) into a continuous one, biasing the estimate toward 0. When fitting a model with age as the categorical variable, we do see some indication that older scholars may be less open to GenAI use, but the picture remains ambiguous. With limited sample size, we need to be cautious when interpreting these results, but at least our models do not point to age as a major predictor of GenAI acceptance

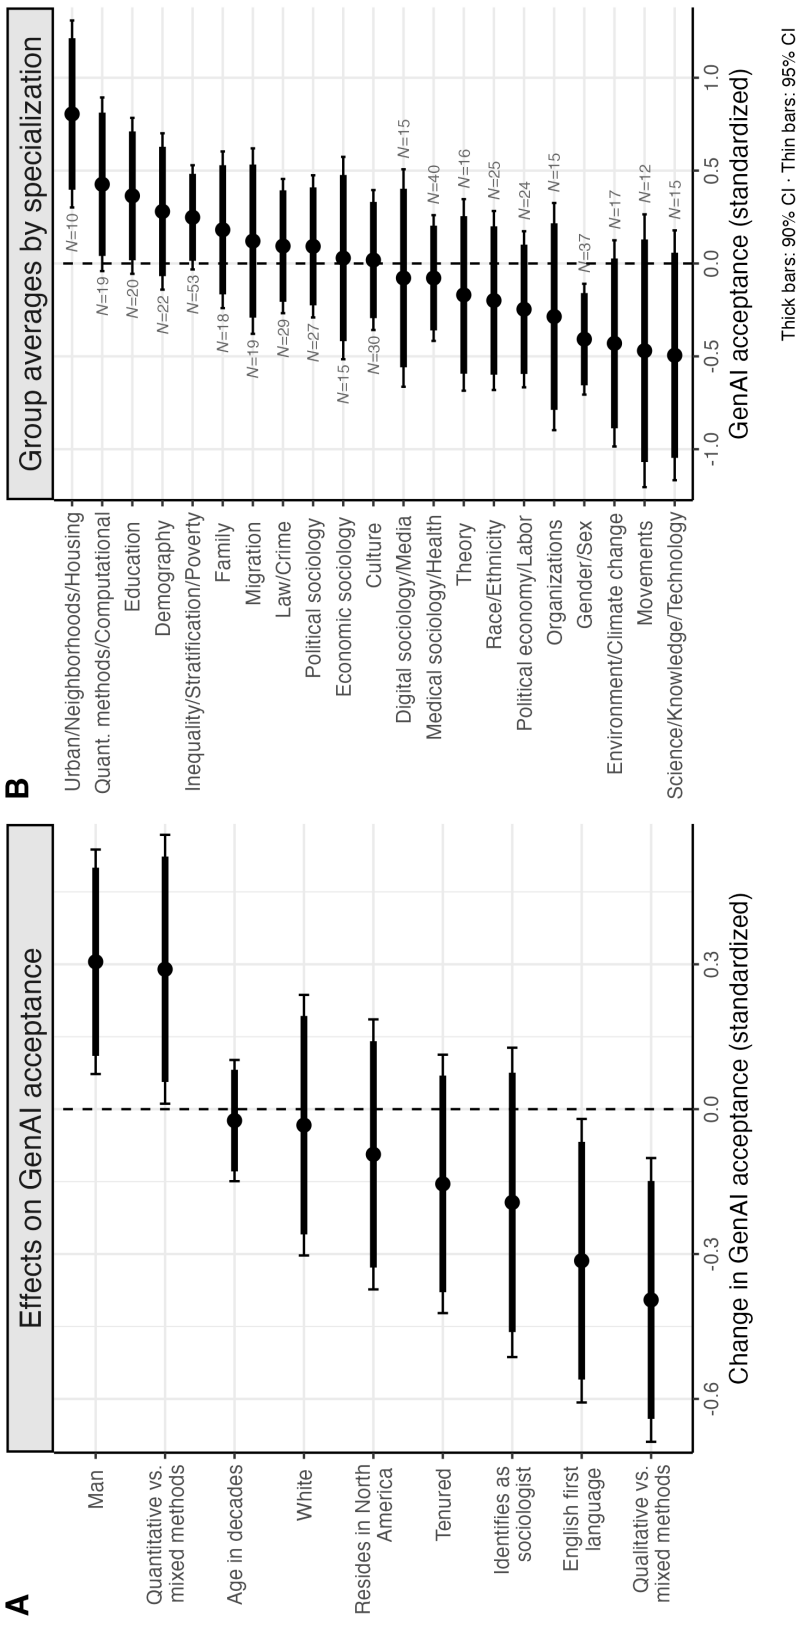


Figure 3: Predictors of GenAI acceptance. *Note:* Panel A shows the effects of an OLS regression ($N = 269$, $R^2 = 0.18$; see Appendix D of the online supplement for details and additional analyses). Panel B shows group averages of GenAI acceptance for researchers' subfields or specializations. Estimates are shown for specializations with at least 10 observations.

among sociologists. Instead, they give some indication that career status may be associated with GenAI acceptance: untenured scholars score 0.16 standard deviations higher on our dependent variable than those with tenure.

Together, these variables account for 18 percent of the variance, indicating that GenAI acceptance is structured along researcher demographics and methodological preferences. Supplementary analyses do not indicate that these effects differ strongly by research stage (see Appendix E of the online supplement). To complement this picture, we present group averages for scholars' specializations (see Panel B of Fig. 3). These were manually coded from open text entries where respondents were asked to name up to three areas they work in. Together, these subfields cover most areas of specialization within sociology. As case numbers are relatively low, any particular estimate should be treated with caution. Still, a pattern emerges: GenAI use appears to be more acceptable in quantitatively oriented areas that typically use standardized data formats like surveys—notably demography, stratification, and education research. Meanwhile, more qualitatively oriented subfields tend to be more critical toward the use of GenAI (e.g., environmental sociology and social movements). With the caveat that there are only 15 observations, some will find it telling that the subfield that might have the most to say about the adoption of new technologies in research—science, knowledge, and technology—is the most skeptical of GenAI use in our data.

Disclosing GenAI Use

The preceding analyses suggest considerable skepticism toward GenAI use in sociological research, and that many uses are unacceptable to a large share of scholars. Nonetheless, many sociologists regularly use GenAI (Table A.1 in the online supplement), and a majority of sociologists consider GenAI either "very useful" or "somewhat useful" for writing, analyzing data, and reviewing literature (see supplementary analyses in Appendix G of the online supplement). This raises the question of whether and how the field can establish norms for its appropriate use and documentation. Several journals have policies on GenAI reporting, and many article submission portals now ask scholars whether and how GenAI was used in their research or writing process. In June 2026, a group of 81 scholars from the behavioral sciences proposed a checklist to standardize such reporting (Feuerriegel et al. 2026).

We asked respondents whether they support journals' adoption of mandatory GenAI disclosure statements (see Fig. 4; our own statement can be found in Appendix K of the online supplement). An overwhelming majority of sociologists support or strongly support such policies (82.2 percent) and few (8.1 percent) explicitly oppose them. However, we also find that trust in the veracity of such statements is low: 22 percent of sociologists trust that GenAI disclosures made by other researchers are truthful. We return to this finding in our discussion.

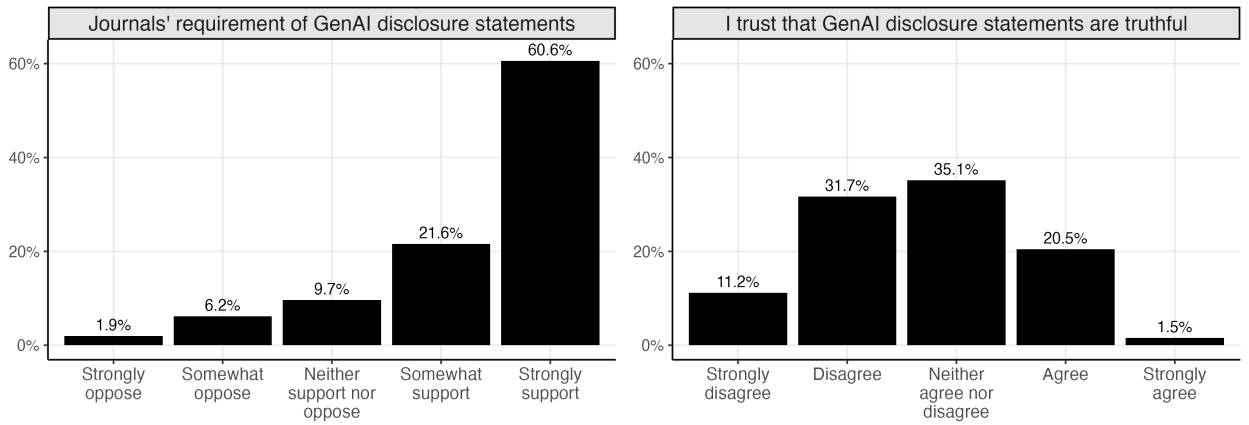


Figure 4: Opinions on journals' adoption of mandatory disclosure statements on GenAI use. *Note:* The figure is based on responses by those identifying as sociologists ($N = 259$).

Limitations

Probably the most important limitation of our survey is that there is likely to be some degree of self-selection bias in our sample. In the absence of ground-truth data, we can only speculate on the nature of such bias. Two dynamics strike us as worthy of highlighting.

First, it appears likely that those with *formed* opinions about GenAI were more inclined to take part in our survey than those without. Consequently, we probably underestimate the prevalence of nonattitudes among sociologists (i.e., "don't know/can't tell"), which currently account for only 3.4 percent of responses across the 19 items. This implies that sociologists might be less divided over individual GenAI use scenarios than Figure 1 suggests, as many have no formed opinion (yet). In other words, there are likely sociologists who do not care or do not know much about GenAI, who are underrepresented in our descriptive results. Second, it is possible that those with *stronger* stances—be it for or against GenAI—were more eager to participate in the survey and voice their opinions. In that case, we underestimate the share of sociologists who consider scenarios "acceptable if disclosed" or have heterogeneous views across scenarios. If true, the lack of unimodality in our acceptance measure would, at least partially, reflect selection, rather than genuine division in the discipline.

Other forms of bias are conceivable: the authors of this article identify as men, work in the United States, use primarily quantitative methods, and have used language models in published research, albeit mainly in the context of explicitly computational text analysis. We cannot exclude the possibility that this may have biased participation along related lines. That said, all invitees were given the opportunity to take the survey anonymously, responses do not suggest that respondents held back criticism of GenAI, and our sociodemographic variables and the information we collected on career status, methods preferences, and subfield suggest that our sample includes authors from diverse segments of the discipline (see Table A.1 in the online supplement).

Overall, then, our own view is that the results broadly reflect stances held in the discipline, but that they are likely somewhat biased toward those invested in the issue and those holding stronger views. Readers may, of course, reasonably differ on this point. What we can confidently say is that the results presented are currently the best available evidence on sociologists' stances toward GenAI use, as existing surveys do not focus on normative questions (Alvero et al. 2026) and/or were distributed in ways likely to induce even greater bias (Lyttelton, Massenkoff, and Wilmers 2026). Although self-selection is hard to eliminate, institutional backing by scholarly associations for future surveys could mitigate these concerns.

Another limitation relates to the fact that the GenAI landscape is complex and quickly evolving, producing ever new tools and terminology. Our goal was to survey the discipline by constructing items that would not alienate respondents with unfamiliar technical terms, and that could be answered by all sociologists, rather than a few experts. We therefore framed scenarios around chat-style, proprietary models—the mode of using GenAI that is still most widespread and likely to be familiar to all respondents. But this choice comes with downsides.

First, it is possible that some respondents reject this particular use modality of GenAI, but are more open to others. For instance, some may reject proprietary models that require data sharing with a third party, but be enthusiastic about the use of open-weight models that can be run locally; others might be encouraged by recent efforts to train models using permissively licensed materials while condemning approaches that disregard copyright protections entirely. Although we asked about some such conditions in additional questions (see Appendix H of the online supplement), our main results reflect stances toward a particular mode of usage, rather than GenAI writ large. Open source enthusiasts could fairly point out that our framing ties GenAI inseparably to questionable organizational and technological practices. Second, we did not ask about some of the more recent developments around GenAI. Most notably, we did not specifically ask about tools designed to integrate with LLMs, such as retrieval systems (Lewis et al. 2020), or "agents" (Engzell and Wilmers 2026; Zhang 2026). Therefore, the acceptability judgments presented here point toward typical use cases (chat-style and proprietary), but likely overlook many other uses—some of which may soon become more pervasive among sociologists. Our results are very much a snapshot of a moving target.

Conclusion

Together, our findings highlight a few points of consensus in the field. A majority find that generating a first draft of an article is unacceptable, as is using GenAI to conduct an interview without disclosing this to the interviewee or making substantive use of GenAI tools when reviewing others' manuscripts. Meanwhile, most believe it is acceptable to use these technologies to search for literature, polish a paper's prose, organize fieldnotes, or write code for visualizations. Beyond these tasks, the path is more divided. Mapping this heterogeneity in opinions is an important step toward establishing common ground, as it gives scholars the chance to put their own stances into perspective.

Our qualitative analysis revealed uncertainty, skepticism, and concerns even among those who are more accepting of using GenAI in research. Some of the more accepting scholars suggest that GenAI provides equal footing for researchers at resource-constrained institutions and equate it to the kind of work that a research assistant might do or the feedback a colleague might give. Others emphasize that it is a tool like any other and is acceptable so long as the researcher maintains creative control and ultimate responsibility over the output. Nevertheless, concerns about "slop," privacy, environmental costs, and cognitive decline abound. Regardless of one's stance on the acceptability of GenAI use, the open-ended responses support the sentiment that "the status quo is both undesirable and unsustainable" (Munger et al. 2026:3). Establishing new institutional norms will be central to the conversation in the coming years.

Disclosing how one has used GenAI tools in producing peer-reviewed materials is likely to be one such norm. Our results indicate that the vast majority of sociologists would like to see journals implement mandatory GenAI disclosure statements, and we commend recent efforts to standardize such practices (Feuerriegel et al. 2026). However, our findings reveal a dilemma: Many uses of GenAI are deemed generally unacceptable by large parts of the discipline. This incentivizes downplaying use of GenAI. Thus, it is unsurprising that few researchers trust GenAI disclosure statements. Solutions to this dilemma are unlikely to be effortless. Still, we echo the concerns of Munger et al. that "uncertainty about what kinds of AI use is acceptable can give short-term advantages to the least conscientious actors in the system" (2026:3). Whatever institutional norms the field establishes in response to GenAI are unlikely to placate everyone, but the way forward is likely through greater transparency and accountability, and encouraging open conversation about the proper role of technology in our discipline.

Notes

1 We share some of the concerns regarding this terminology (e.g., Bender and Hanna 2025), but adopt it here given its wide use. We also provided a brief explainer in our survey materials to orient respondents (see Appendix B of the online supplement).

2 The authors note that "quality varies" among the 33 articles produced in their experiment, but "many contain discoveries that, with little extra tweaking, would fit well in mid-tier social science journals." (ibid., p. 13).

3 This is slightly lower than a 2025 survey using the same sampling strategy (Alvero et al. 2026), which may suggest some fatigue among sociologists with regard to surveys on GenAI.

References

Abramson, Corey M., Tara Prendergast, Zhuofan Li, and Daniel Dohan. 2026. "Qualitative Research in an Era of Artificial Intelligence: A Pragmatic Approach to Data Analysis, Workflow, and Computation." *Annual Review of Sociology* 52. https://doi.org/10.1146/annurev-soc-011824-104836

Alvero, A. J., Dustin S. Stoltz, Oscar Stuhler, and Marshall A. Taylor. 2026. "Generative AI in Sociological Research: State of the Discipline." *Sociological Science* 13:45–62. https://doi.org/10.15195/v13.a3

Argyle, Lisa P., Christopher A. Bail, Ethan C. Busby, Joshua R. Gubler, Thomas Howe, Christopher Rytting, Taylor Sorensen, and David Wingate. 2023. "Leveraging AI for Democratic Discourse: Chat Interventions Can Improve Online Political Conversations at Scale." *Proceedings of the National Academy of Sciences* 120(41):e2311627120. https://doi.org/10.1073/pnas.2311627120

Bail, Christopher A. 2024. "Can Generative AI Improve Social Science?." *Proceedings of the National Academy of Sciences of the United States of America* 121(21):e2314021121. https://doi.org/10.1073/pnas.2314021121

Bender, Emily M. and Alex Hanna. 2025. *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want.* New York, NY: Harper.

Boelaert, Julien, Samuel Coavoux, Etienne Ollion, Ivaylo Petev, and Patrick Präg. 2025. "Machine Bias. How Do Generative Language Models Answer Opinion Polls?." *Sociological Methods & Research* 54(3):1156–96. https://doi.org/10.1177/00491241251330582

Bonikowski, Bart, Yuchen Luo, and Oscar Stuhler. 2022. "Politics as Usual? Measuring Populism, Nationalism, and Authoritarianism in U.S. Presidential Campaigns (1952-2020) with Neural Language Models." *Sociological Methods & Research* 51(4):1721–87. https://doi.org/10.1177/00491241221122317

Broska, David, Michael Howes, and Austin van Loon. 2025. "The Mixed Subjects Design: Treating Large Language Models as Potentially Informative Observations." *Sociological Methods & Research* 54(3):1074–109. https://doi.org/10.1177/00491241251326865

Cao, Likun. 2026. "Visualizing the Adoption of Large Language Models across Sociology Subfields." *Socius* 12. https://doi.org/10.1177/23780231261460657

Chakravorti, Tatiana, Xinyu Wang, Pranav Narayanan Venkit, Sai Koneru, Kevin Munger, and Sarah Rajtmajer. 2025. "Social Scientists on the Role of AI in Research." https://doi.org/10.1609/aies.v8i1.36568

Chatoo, Zehra Chatoo. 2026. "The AI Judgment Penalty." https://codeforgoodnow.com/wp-content/uploads/2026/06/Research-Report-The-AI-Judgment-Penalty.pdf.

Cohen, Philip N. 2025. "SocArXiv moratorium on AI-topic papers, policy in formation." November 27, 2025. https://socopen.org/2025/11/27/socarxiv-moratorium-on-ai-topic-papers-policy-in-formation/.

Cohen, Philip N. 2026. "SocArXiv AI Policy." August 26, 2026. https://socopen.org/ai-policy/.

Cort, David, Joya Misra, Laurel Smith-Doerr, and Donald Tomaskovic-Devey. 2026. "ASR Annual Reports for 2025." https://www.asanet.org/publications/journals/annual-reports/.

Cranney, Katelyn, Solène Delecourt, and Rembrand Koning. 2026. "Global Evidence on Gender Gaps and Generative AI Over Time." *Harvard Business School Entrepreneurial Management Working*. https://doi.org/10.2139/ssrn.6880085

Davidson, Thomas. 2024. "Start Generating: Harnessing Generative Artificial Intelligence for Sociological Research." *Socius* 10. https://doi.org/10.1177/23780231241259

Davidson, Thomas. 2026. "Multimodal Large Language Models Can Make Context-Sensitive Hate Speech Evaluations Aligned with Human Judgement." *Nature Human Behaviour* 10:514–30. https://doi.org/10.1038/s41562-025-02360-w

Davidson, Thomas and Daniel Karell. 2025. "Integrating Generative Artificial Intelligence into Social Science Research: Measurement, Prompting, and Simulation." *Sociological Methods & Research* 54(3):775–93. https://doi.org/10.1177/00491241251339184

Do, Salome, Etienne Ollion, and Rubing Shen. 2022. "The Augmented Social Scientist: Using Sequential Transfer Learning to Annotate Millions of Texts with Human-Level Accuracy." OSF Preprints. https://doi.org/10.31235/osf.io/3fkzc

Dunivin, Zackary Okun. 2025. "Scaling Hermeneutics: A Guide to Qualitative Coding with LLMs for Reflexive Content Analysis." *EPJ Data Science* 14(28). https://doi.org/10.1140/epjds/s13688-025-00548-8

Engzell, Per and Nathan Wilmers. 2026. "The Paper Factory." https://doi.org/10.31235/osf.io/24xfq_v1

Feuerriegel, Stefan, Christopher Barrie, M. J. Crockett, Laura K. Globig, Killian L. McLoughlin, Dan-Mircea Mirea, Arthur Spirling, Diyi Yang, Tim Althoff, Maria Antoniak, Lisa P. Argyle, and Ashwini Ashokkumar. 2026. "A Reporting Checklist for Large Language Models in Behavioural Science." *Nature Human Behaviour* 10:1182–6. https://doi.org/10.1038/s41562-026-02492-7

Gai, Phyliss Jia, Jiayi Hou, and Yanping Tu. 2025. "Competence Penalty Is a Barrier to the Adoption of New Technology." *SSRN* (5255039). https://doi.org/10.2139/ssrn.5255039

Haimson, Oliver L., Samuel Reiji Mayworm, Alexis Shore Ingber, and Nazanin Andalibi. 2025. "AI Attitudes Among Marginalized Populations in the US: Nonbinary, Transgender, and Disabled Individuals Report More Negative AI Attitudes." Pp. 1224–37 in *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency.* https://doi.org/10.1145/3715275.3732081

Ibrahim, Elida Izani, Laura K. Nelson, and Andrea Voyer. 2026. "Should Generative AI be Used in Reflexive Qualitative Research?." *Qualitative Inquiry*. https://doi.org/10.1177/1077800426145620

Ibrahim, Elida Izani and Andrea Voyer. 2026. "Qualitative Research with LLM Chatbots: Technological Reflexivity for Interpretative Technology." *Qualitative Research* 26(1):133–59. https://doi.org/10.1177/14687941251390794

Jowsey, Tanisha, Virginia Brain, Victoria Clarke, Deborah Lupton, and Michelle Fine. 2025. "We Reject the Use of Generative Artificial Intelligence for Reflexive Qualitative Research." *Qualitative Inquiry*. https://doi.org/10.1177/10778004251401851

Joyce, Kelly and Taylor M. Cruz. 2024. "A Sociology of Artificial Intelligence: Inequalities, Power, and Data Justice." *Socius* 10. https://doi.org/10.1177/23780231241275393

Joyce, Kelly, Laurel Smith-Doerr, Sharla Alegria, Susan Bell, Taylor Cruz, Steve G. Hoffman, Safiya Umoja Noble, and Benjamin Shestakofsky. 2021. "Toward a Sociology of Artificial Intelligence: A Call for Research on Inequalities and Structural Change." *Socius* 7:2378023121999581. https://doi.org/10.1177/2378023121999581

Karell, Daniel, Jeffrey Sachs, and Ryan Barrett. 2025. "Synthetic Duality: A Framework for Analyzing Generative Artificial Intelligence's Representation of Social Reality." *Poetics* 108:101966. https://doi.org/10.1016/j.poetic.2024.101966

Kim, Junsol and Byungkyu Lee. 2026. "AI-Augmented Surveys: Leveraging Large Language Models and Surveys for Opinion Prediction." *arXiv preprint* (2305.09620v4). https://doi.org/10.48550/arXiv.2305.09620

Kozlowski, Austin. 2026. "AutoInterp Agent Framework." GitHub. https://github.com/akozlo/AutoInterp.

Kozlowski, Austin and James Evans. 2025. "Simulating Subjects: The Promise and Peril of Artificial Intelligence Stand-Ins for Social Agents and Interactions." *Sociological Methods & Research* 54(3):1017–73. https://doi.org/10.1177/00491241251337316

Law, Tina and Leslie McCall. 2024. "Artificial Intelligence Policymaking: An Agenda for Sociological Research." *Socius* 10:23780231241261596. https://doi.org/10.1177/23780231241261596

Law, Tina and Elizabeth Roberto. 2025. "Generative Multimodal Models for Social Science: An Application with Satellite and Streetscape Imagery." *Sociological Methods & Research* 54(3):889–932. https://doi.org/10.1177/00491241251339673

Lee, Byungkyu. 2026. "Claude Code + Cursor Setup for Academic Researchers." GitHub. https://github.com/letitbk/claude-academic-setup.

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel. 2020. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *Advances in Neural Information Processing Systems* 33:9459–74. https://doi.org/10.48550/arXiv.2005.11401

Lu, Chris, Cong Lu, Robert Tjarko Lange, Yutaro Yamada, Shengran Hu, Jakob Foerster, David Ha, and Jeff Clune. 2026. "Towards End-to-End Automation of AI Research." *Nature* 651:914–9. https://doi.org/10.1038/s41586-026-10265-5

Lupton, Deborah. 2026. "Generative AI's Impacts on the Environment: What Can Critical Humanities Scholarship Offer Sociological Inquiry?." *Sociology Compass* 20(5):e70195. https://doi.org/10.1111/soc4.70195

Lyttelton, Thomas, Maxim Massenkoff, and Nathan Wilmers. 2026. "Coding Agents in the Social Sciences." May 27, 2026. https://www.anthropic.com/research/coding-agents-social-sciences.

Munger, Kevin, Bert Bakker, Adam Berinsky, Nathalie Giger, Andrew Guess, Natascha Just, Regina Lawrence, Keren Tenenboim-Weinblatt, and Arnout van de Rijt. 2026. "Peer Review 2027: Scenarios for Academic Publishing in the Age of AI." https://doi.org/10.31235/osf.io/594zj_v1

Nguyen, Duc Cuong and Catherine Welch. 2026. "Engaged and Responsible Scholarship: Why Qualitative Researchers Should Not Embrace GenAI." *Business & Society* 65(7):1787–92. https://doi.org/10.1177/00076503251386539

Stoltz, Dustin S. 2025. "The Duality of Class and Love: Homogamy Spaces and the New York Social Elite, 1970-2020." *Poetics* 113:102060. https://doi.org/10.1016/j.poetic.2025.102060

Stuhler, Oscar, Cat Dang Ton, and Etienne Ollion. 2025. "From Codebooks to Promptbooks: Extracting Information from Text with Generative Large Language Models." *Sociological Methods & Research* 54(3):794–848. https://doi.org/10.1177/00491241251336794

Stuhler, Oscar, Dustin S. Stoltz, A. J. Alvero, and Marshall A. Taylor. 2026. "Replication Data for: "Is it 'Okay' to Use Generative AI in Sociological Research?"." Version V1. Harvard Dataverse. https://doi.org/10.7910/DVN/HNKOSO.

Than, Nga, Leanne Fan, Tina Law, Laura K. Nelson, and Leslie McCall. 2025. "Updating 'The Future of Coding': Qualitative Coding with Generative Large Language Models." *Sociological Methods & Research* 54(3):849–88. https://doi.org/10.1177/00491241251339188

Timmermans, Stefan and Iddo Tavory. 2026. "AI and Qualitative Data Analysis: On the Vulnerability of Thematic Analysis and the Importance of Theorizing." *Qualitative Inquiry*. https://doi.org/10.1177/10778004261470128

Vertesi, Janet, Danah Boyd, Alex S. Taylor, and Benjamin Shestakofsky. 2026. "Reckoning with the Political Economy of AI: Avoiding Decoys in Pursuit of Accountability."

Pp. 2186–205 in *The 2026 ACM Conference on Fairness, Accountability, and Transparency*. https://doi.org/10.1145/3805689.3806739

Waight, Hannah, Solomon Messing, Anton Shirikov, Margaret E. Roberts, Jonathan Nagler, Jason Greenfield, Megan A. Brown, Kevin Aslett, and Joshua A. Tucker. 2025. "Quantifying Narrative Similarity Across Languages." *Sociological Methods & Research* 54(3):933–83. https://doi.org/10.1177/00491241251340080

Wankmüller, Sandra. 2024. "Introduction to neural transfer learning with transformers for social science text analysis." *Sociological Methods & Research* 53(4):1676–752. https://doi.org/10.1177/00491241221134527

Xiao, Jian and Pieter Verdegem. 2025. "Hype, Resistance, Power and Inequalities: Why Synthesizing Critical Perspectives Is Essential to AI Research." *Sociology Compass* 19(12):e70155. https://doi.org/10.1111/soc4.70155

Zajko, Mike. 2022. "Artificial Intelligence, Algorithms, and Social Inequality: Sociological Contributions to Contemporary Debates." *Sociology Compass* 16(3):e12962. https://doi.org/10.1111/soc4.12962

Zhang, Yongjun. 2026. "Vibe Researching As Wolf Coming: Can AI Agents with Skills Replace or Augment Social Scientists?." *arXiv preprint arXiv:2602.22401*. https://doi.org/10.48550/arXiv.2602.22401

Acknowledgments: First and foremost, we want to thank all respondents of our survey who were willing to take the time to share their thoughts. This article exists because of their generosity. For feedback on survey design and prior drafts, we are indebted to Zackary Dunivin, Jan Fuhse, Patrick Kaminski, Ryan Leung, Santiago J. Molina, Laura K. Nelson, Étienne Ollion, Christine Percheski, and Cat Dang Ton. We are also thankful for comments received when presenting this work at the MZES speaker series at Mannheim University and the session on Implications of AI at the 2026 ASA Annual Meeting.

Corresponding author: Oscar Stuhler, E-mail: oms@northwestern.edu

Oscar Stuhler: Department of Sociology, Northwestern University.
E-mail: oms@northwestern.edu
Dustin S. Stoltz: Department of Sociology and Anthropology, Lehigh University.
E-mail: dss219@lehigh.edu
AJ Alvero: Center for Data Science for Enterprise and Society, Cornell University.
E-mail: aja326@cornell.edu
Marshall A. Taylor: CREST, ENSAE, Institut Polytechnique de Paris.
E-mail: marshall.taylor@ensae.fr