

Generative AI in Sociological Research: State of the Discipline

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

a) Cornell University; b) Lehigh University; c) Northwestern University; d) New Mexico State University

Abstract: Generative artificial intelligence (GenAI) has garnered considerable attention for its potential utility in research and scholarship, even among those who typically do not rely on computational tools. However, early commentators have also articulated concerns about how GenAI usage comes with enormous environmental costs, serious social risks, and a tendency to produce low-quality content. In the midst of both excitement and skepticism, it is crucial to take stock of how GenAI is actually being used. Our study focuses on sociological research as our site, and here we present findings from a survey of 433 authors of articles published in 50 sociology journals in the past five years. The survey provides an overview of the state of the discipline with regard to the use of GenAI by providing answers to fundamental questions: how (much) do scholars use the technology for their research; what are their reasons for using it; and how concerned, trustful, and optimistic are they about the technology? Of the approximately one third of respondents who self-report using GenAI at least weekly, the primary uses are for writing assistance and comparatively less so in planning, data collection, or data analysis. In both use and attitudes, there are surprisingly few differences between self-identified computational and non-computational researchers. In general, respondents are very concerned about the social and environmental consequences of GenAI. Trust in GenAI outputs is low, regardless of expertise or frequency of use. Although optimism that GenAI will improve is high, scholars are divided on whether GenAI will have a positive impact on the field.

Keywords: generative AI; sociology of science; large language models; sociological research practices; computational sociology; survey

Reproducibility Package: A replication repository for this article can be found at: https://github.com/Marshall-Soc/genai_sociology. The data for this article are hosted on the Harvard Dataverse (Alvero et al. 2025) and can be accessed through: <https://doi.org/10.7910/DVN/1CXIRP>

Citation: Alvero, AJ, Dustin S. Stoltz, Oscar Stuhler, and Marshall A. Taylor. 2025. "Generative AI in Sociological Research: State of the Discipline" *Sociological Science* 13: 45-62.

Received: August 23, 2025

Accepted: November 8, 2025

Published: January 20, 2026

Editor(s): Arnout van de Rijt, Cristobal Young

DOI: 10.15195/v13.a3

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. 

FOLLOWING the release of the Generative Pre-trained Transformer-3 (GPT-3) in 2020 and the browser-based conversational interface ChatGPT in 2022, academics have been forced to grapple with generative artificial intelligence (GenAI). For better or worse, products and tools that fall under the label "GenAI" are starting to make inroads into the major aspects of academic life: classroom instruction (Xie & Avila 2025), peer review (Li et al. 2024), administration (An, Yu, & James 2025), and not least *research* (Bail 2024; Broska & McFarland 2024; Zhang et al. 2025b).

Thus far, the use of GenAI as a tool for sociological research has mainly been explored from the perspective of research products. Research articles offer "proof of concepts" by identifying important use cases and limitations (see, e.g., Boelaert et al. 2025; Davidson & Karel 2025). Meanwhile, less attention has been given to how and why scholars in our discipline decide to use or avoid GenAI (c.f. Watermeyer

et al. 2024a, 2024b). In addition, while there have been a series of recent publications that showcase high-level applications of GenAI—say, classifying images (Law & Roberto 2025) or extracting information from text (Stuhler, Ton, & Ollion 2025)—there has been comparatively little open discussion about how sociologists may use GenAI for more every-day aspects of research—be it for writing, generating ideas, debugging code, or otherwise.

In this article, we present a survey for which we asked sociologists and their collaborators about how they perceive and use (or do not use) GenAI in their research. Our sampling frame includes the authors of all articles published in 50 sociology journals in the past five years. We contacted a random sample of these authors as well as all authors who published “computational” articles. This allows us to paint the landscape of GenAI usage and attitudes among the population of researchers contributing to recent sociological scholarship. Going beyond speculation about what technology “could” do for sociologists, we offer the first systematic evidence on how GenAI is currently being used and seen by sociologists. Thereby, we provide an empirical baseline for a more grounded discussion on GenAI’s role in sociological research.

To preview our results, the most common usage of GenAI is for writing assistance, especially for grammar, spelling, and paraphrasing sections of one’s own writing. We find that roughly 34 percent of sociologists and their collaborators have used GenAI in this capacity—on par with findings from surveys of scholars in other fields (Kwon 2025; Ng et al. 2025). Respondents use GenAI because they perceive that it saves them time, out of curiosity, and because it is increasingly incorporated into tools they already use (e.g., search engines). However, very few reported feeling pressure to use GenAI from their collaborators, field, or institution. Although we anticipated differences in use and attitudes between those who use computational methods and those who do not, we found very few. Similarly, we find that expertise—in terms of self-reported familiarity and use frequency—is a weak predictor of attitudes. In general, the vast majority of scholars are very concerned about the social and environmental consequences of GenAI and also distrust GenAI outputs. Finally, scholars agreed that GenAI would likely improve in the next few years but were divided about whether it would have a net positive effect on the field.

Background

What Is “Generative” AI?

Language surrounding “artificial intelligence” is often imprecise, with a wide range of technologies grouped under this umbrella (Bender & Hanna 2025, pp. 1–20). Therefore, we begin by briefly explaining how we define GenAI in this study (also see our primer on GenAI in email invitations sent to survey respondents in Appendix A).

The distinguishing characteristic of GenAI from other computational methods is in the name: as opposed to “discriminative” models that find optimal boundaries in data, primarily for classification tasks, generative models are designed to *generate*

text, images, audio, and video. We focus on text generation models, as that is also the focus of the current literature in sociology and represents the most common form of GenAI interaction (Zhang, Xu, & Alvero 2025a). When employing such models, users typically write a *prompt* containing the text that a trained model will take as its point of departure, such as instructions, questions, or an arbitrary query. This prompt is encoded as tokens¹, then the underlying model generates a *probable* continuation of the prompt by identifying the most likely *next token*, defined in probabilistic terms. The model then incorporates that new token into another round of generation and continues until a stopping rule, such as a maximum number of tokens to generate, is satisfied. In other words, the prompt provides a starting point from which the model starts *generating* (for a more thorough introduction and discussion of practical aspects, we recommend Chae & Davidson 2025).

Generative AI in Academic Research

Setting aside quality, many of the tasks that comprise academic research are highly amenable to GenAI precisely because they are text based. However, systematic surveys on the uptake of and attitudes toward GenAI for research have been scarce. Some cross-disciplinary surveys suggest that scholars in the technical and life sciences have embraced GenAI more quickly and are more optimistic about its prospects than those in the humanities or social sciences (Andersen et al. 2025; Hrycyshyn & Eassom 2025; OUP 2024). Most studies report that younger or early career scholars use GenAI more frequently (Andersen et al. 2025; Dorta-González et al. 2024; Hrycyshyn & Eassom 2025; Kwon 2025; Perkowski & Marsal 2024; Van Noorden & Perkel 2023) and consider it more acceptable to do so (Kwon 2025)—though one survey found the opposite (OUP 2024). Some report that men use GenAI more frequently than women (Chakravorti et al. 2025; Dorta-González et al. 2024; Perkowski & Marsal 2024), whereas others find no gender differences (Andersen et al. 2025). A recurring theme noted by several studies is that opinions about GenAI are very heterogeneous and that there is far from a consensus on what kinds of use of GenAI are legitimate (Andersen et al. 2025; Kwon 2025).

Differences in sampling strategies, fielding time, survey design, and even the assumed definition of “AI” make it difficult to compare results and lead to contrasting findings. For instance, a 2024 survey of researchers by the publisher Wiley reports that 81 percent of respondents used ChatGPT (Hrycyshyn & Eassom 2025). A survey of PhD-holding economists working at European central banks fielded in the same year reports that fewer than half are using “OpenAI’s ChatGPT, Google Gemini, Github Copilot, Meta’s LLaMa, Anthropic’s Claude, or another generative AI tool” (Perkowski & Marsal 2024). Although these studies exhibit heterogeneity in questions and responses, one key takeaway is that researchers across a range of fields have started to incorporate GenAI into their workflow.

Generative AI in Sociological Research

Sociologists have also begun to point out potential use cases of GenAI models for research tasks (Bail 2024; Davidson 2024). Empirical studies have examined whether GenAI can increase research efficiency, for example, by generating survey

questions (Götz et al. 2024); imputing missing data (Kim & Lee 2023); engaging in an exploratory dialog with qualitative data (Hayes 2025; Ibrahim & Voyer 2024); and classifying, annotating, and extracting information from text or images (Gilardi, Alizadeh, & Kubli 2023; Law & Roberto 2025; Lin & Zhang 2025; Maranca et al. 2025; Nelson et al. 2025; Schwitter 2025; Stuhler et al. 2025). Another branch of scholarship has explored GenAI's potential for simulating human behavior (Alvero et al. 2024; Anthis et al. 2025; Broska & McFarland 2024; Kozlowski & Evans 2025). Much of this work has focused on simulating responses to survey items (Boelaert et al. 2025; Broska, Howes, & van Loon 2025; Kim & Lee 2023; Kozlowski, Kwon, & Evans 2024; Zhang et al. 2025a; outside of sociology, see, e.g., Argyle et al. 2023b), but GenAI can also be used to simulate human interaction and communication (Argyle et al. 2023a; Horton 2023; Karell, Sachs, & Barrett 2024).

Sociologists have also pointed out how GenAI complicates sociological analysis. This includes homogenization of outputs (Alvero et al. 2024; Zhang et al. 2025a); injecting new forms of machine bias (Boelaert et al. 2025; Maranca et al. 2025; Stuhler et al. 2025); and large language models (LLMs) parroting of government-approved ideological positions (Waight et al. 2024) or potentially the positions favored by the organizations producing the models (Martin 2023). In some cases, GenAI may perform only on par with, or even worse than, traditional methods that are often more transparent and less computationally costly (Ashwin, Chhabra, & Rao 2025; Mu et al. 2023; Nelson et al. 2025; Stuhler et al. 2025). Furthermore, if sources of empirical data, such as open-ended survey responses (Zhang et al. 2025a) or social media posts (Sourati et al. 2025) are increasingly produced using GenAI, future research will be confronted with difficult questions about sociological explanation and inference.

In the midst of both excitement and skepticism, we are at a critical moment where GenAI is clearly making inroads into research practices, and not just for those who do “computational” work. Yet we have little systematic knowledge about how sociologists and their collaborators currently use GenAI and about the field’s sentiment toward this new technology. Our goal is to help close this gap.

Data and Methods

Sample

We used a bibliometric multistage sampling design to construct a representative sample of both computational and non-computational sociologists and their collaborators. The sample includes 219 non-computational authors and an over-sample of 214 computational authors (defined as those authoring an article with computational terms in the title, abstract, and keywords) for a total of 433 respondents. We used rake weights (DeBell & Krosnick 2009) to adjust for this oversampling as well as selection bias into our sample—namely, gender and location. Our sampling method is detailed in Appendix B. Figure 1 shows a summary schematic of the sampling strategy.

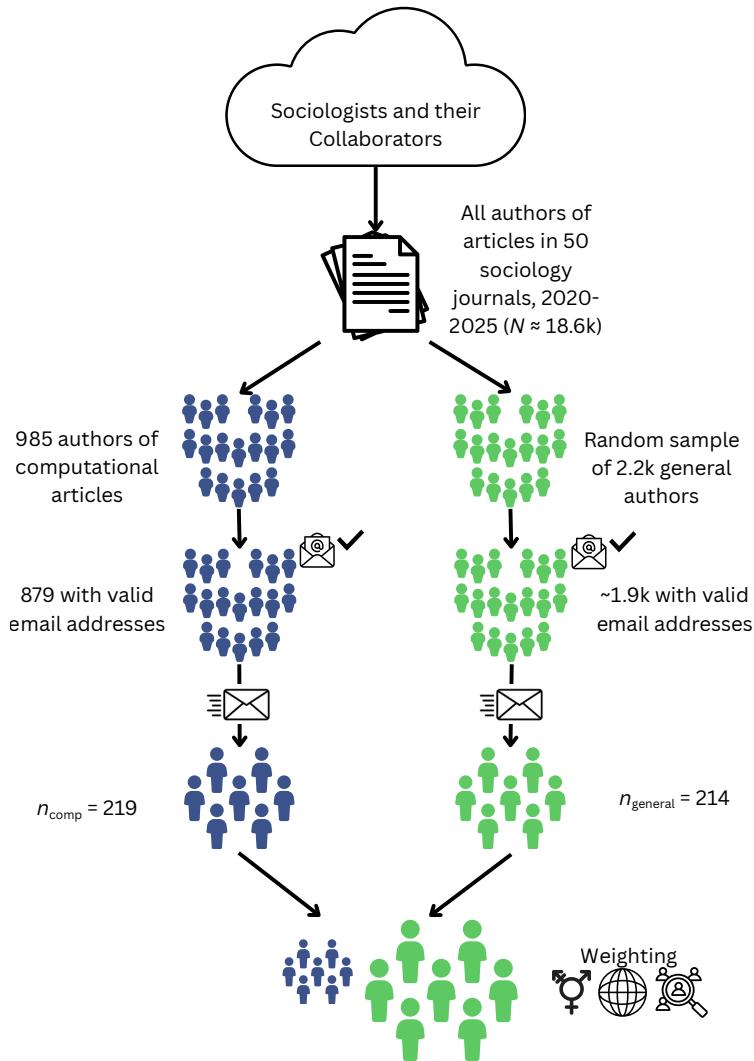


Figure 1: Schematic representation of sampling strategy. *Note:* The survey was fielded from January 2025 to June 2025. The response rates for the computational and general samples were about 24 percent and 11 percent, respectively. The weighting symbols represent, from left to right, gender identity, location, and subsample.

Measures

Our survey focuses on who uses GenAI, its use cases, reasons for and concerns over using GenAI, and sources of optimism and trust in the technology. See the replication repository for information about accessing the survey questionnaire and data, which includes variables not explored in the main text.

Analytic Plan

One particular interest of our work is to describe differences between scholars who self-report using computational methods and those who do not. To do this,

we use an array of univariate, bivariate, and multivariate descriptive data visualizations. We also use regression models to analyze the extent to which trust and optimism are associated with levels of GenAI use and familiarity. All descriptive analyses and regressions are rake-weighted with weight-adjusted standard errors and confidence intervals (Lumley & Scott 2017). See Appendix C for our findings as percentages over individual respondent categories (e.g., X% of respondents stated they “somewhat agree” with Y). Recall that our sample includes sociologists and their collaborators. To simplify, we refer to this population collectively as *scholars*.

Findings

Who Uses GenAI?

How many scholars use GenAI and for what purpose(s)? The top panel of Figure 2 shows the distribution of the frequency with which GenAI is used in research practices. Computational scholars use GenAI with higher frequency than non-computational scholars, but the differences are generally small. For computational scholars, “weekly” is the most frequent response (27.8 percent; $18.8\% \leq \mu \leq 39$ percent), but with “never” a close second (23.7 percent; $15.2\% \leq \mu \leq 34.9$ percent), whereas non-computational say “at least once” the most (26.7 percent compared to 10.8 percent) with “weekly” (25.3 percent; $17.5\% \leq \mu \leq 36.5$ percent) a close second.

GenAI Use Cases

Of those who have used GenAI in their research at least once, how do they use it? The second panel of Figure 2 shows the distribution of GenAI use at various stages of research—planning, writing, data collection, and analysis—separated by non-computational and computational scholars. The percentages sum to 100 percent per research stage, but with the never-users included in the calculation and removed from the visualization. As the bar chart shows, computational scholars are slightly more likely to use GenAI in analysis tasks (21.9 percent; $16\% \leq \mu \leq 29.2$ percent) than non-computational scholars (13.6 percent; $8.4\% \leq \mu \leq 21.4$ percent). Outside this research stage, however, the differences are marginal. Importantly, while the never-users are not shown, they make up the largest group for each research stage—in other words, most scholars do not use GenAI in any of the research stages as defined by the survey. Nonetheless, among those who have used GenAI, writing is the most common research stage where GenAI comes into play for both non-computational and computational scholars. In the Appendix (Figure C.5 in Appendix C), we break down each research stage into specific tasks, such as translating one’s own writing or explaining statistical outputs.²

Of course, the options provided do not encompass all possible tasks. Some noted in an open-ended response field that they use GenAI for generating article titles, creating standardized images for survey stimuli, or automating emails related to conducting research. Furthermore, some respondents indicated that groupings of tasks could be disaggregated, for example, annotating text may be distinguished from classification.

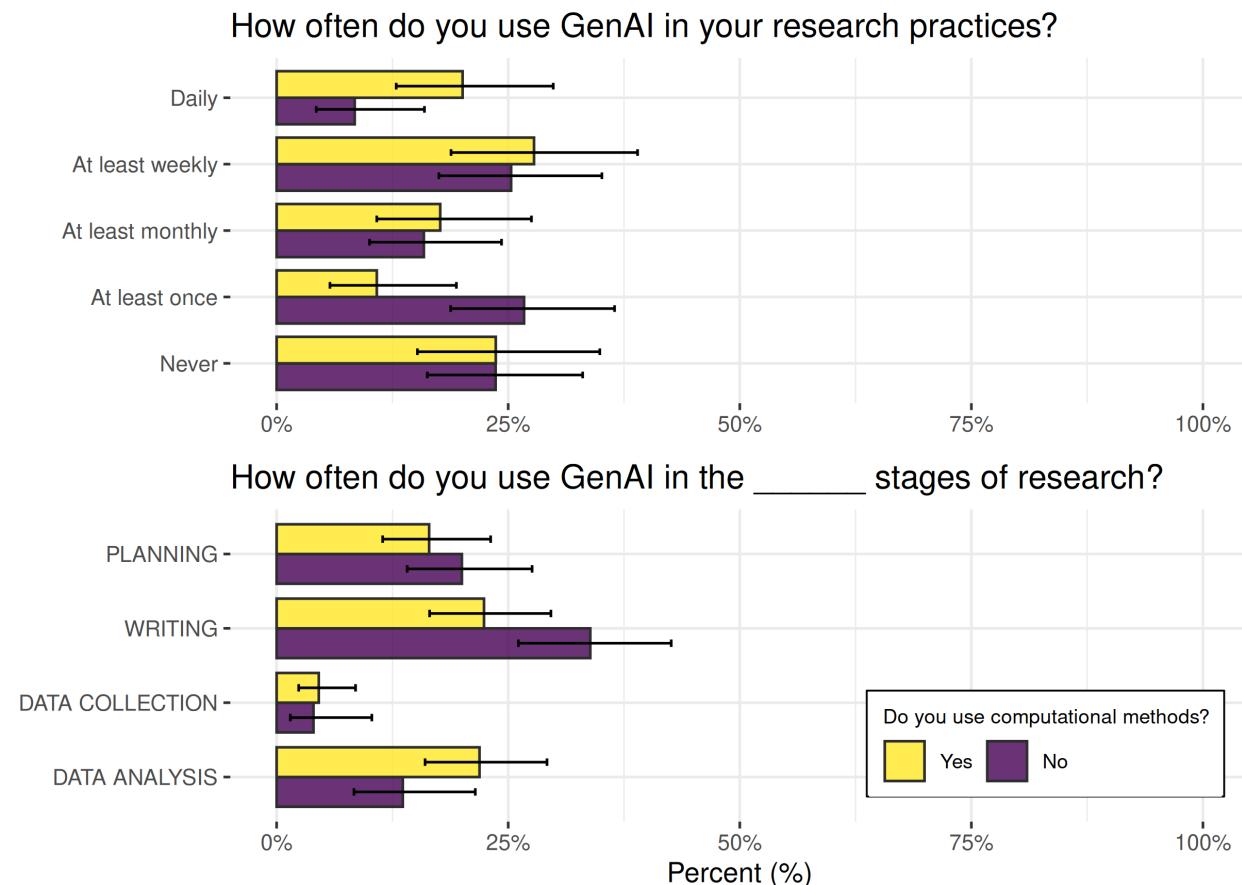


Figure 2: Distribution of GenAI usage frequency. *Note:* Error bars are 95 percent confidence intervals. Top panel $n = 394$ after listwise deletion and bottom panel $n = 334$ after removing respondents who responded having never used GenAI and listwise deletion.

Reasons for Using GenAI

Of those who use GenAI in their research, what are the reasons for doing so? To answer this question, we rely on a combination of closed-ended items and representative selections from open-ended responses. The top panel of Figure 3 visualizes the mean responses for possible reasons, split between computational and non-computational scholars. There was general agreement among scholars that use GenAI because they say it saves them time and satisfies their curiosity. Scholars also report using GenAI because tools they typically use are now incorporating GenAI. As one respondent remarked:

I only use GenAI because Google has normalized AI summaries in their standard searches. I otherwise do not use GenAI in my work.

Scholars are more ambivalent on whether it allows them to focus on more meaningful aspects of research or whether it saves them money, and do not think it enables otherwise impossible research.

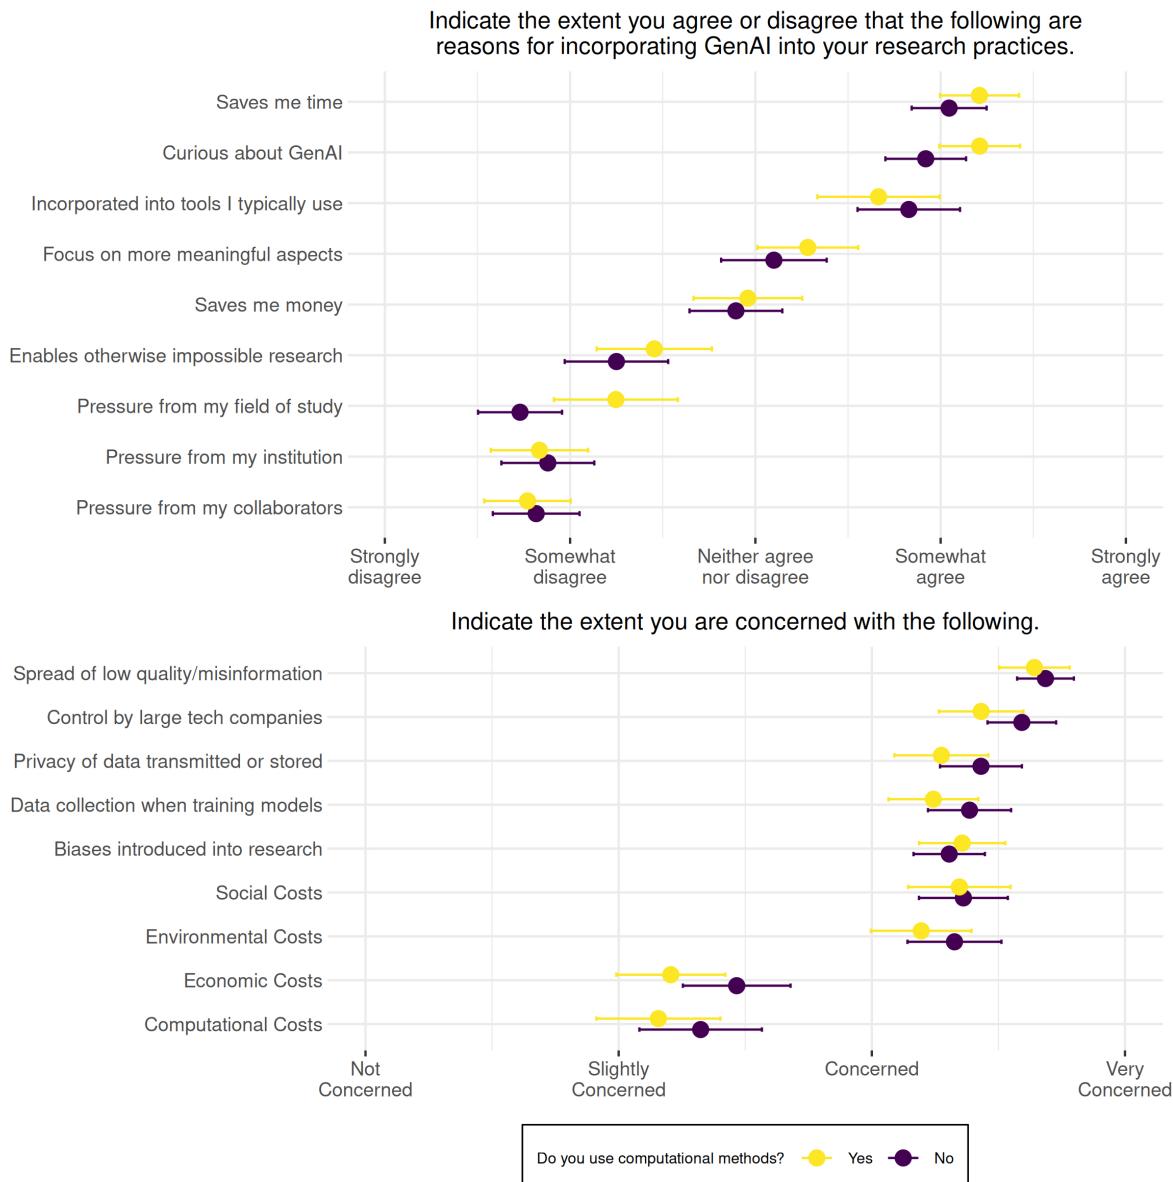


Figure 3: Reasons for, and concerns about, using GenAI in research. Note: Points are means. Error bars are 95 percent confidence intervals. Top panel: $n = 318$ after removing respondents who responded having never used GenAI in their research, as well as listwise deletion. Bottom panel: $n = 376$ after listwise deletion.

External pressures may also be driving reasons to adopt GenAI in research. As one respondent noted:

I feel like people are going to be using it more and more, and I do not want to be left behind. I am worried about that my research process may become too 'slow' or obsolete.

Although this respondent was not alone, the top panel of Figure 3 suggests that pressures from collaborators, the field, or institutions are generally low for computational and non-computational scholars alike. In addition, several open-ended responses from scholars who were not native English speakers noted that GenAI helps them be more confident in their writing (though not for direct translation of their work). As before, there are not many meaningful distinctions between non-computational and computational scholars.

Concerns about GenAI

Among all respondents, what possible costs or risks of GenAI are concerning? When asked about general “costs” of GenAI, scholars express the most concern about its environmental and social impacts (bottom panel of Figure 3). Although non-computational scholars were, on average, more worried about nearly all costs, the differences with computational scholars are not significant. When asked about a series of more specific consequences, respondents expressed the most concern about information quality and the level of control held by large tech companies over GenAI tools. Still, scholars were also highly concerned about the privacy of their data, data collection practices used in training GenAI, and GenAI biasing research practices. On all questions, there are, again, no strong differences between non-computational and computational scholars.

Overall, our items indicate that there is a high level of concern about various aspects of GenAI; of the nine questions we asked, seven were rated as being between “concerned” and “very concerned” on average. To capture possible nuance regarding respondent’s perceptions of risks, we also asked about their GenAI concerns in an open-ended question. For instance, one respondent wrote that they were “wholly against the use of GenAI tools. They have no place in social science practice.” Below, we complement the aggregate findings with further examples that we found to be representative of general themes in the open-ended responses.

One pattern we found is that scholars do not just worry about the spread of low quality content/misinformation, but that GenAI may lead to a general reduction in critical thinking. One respondent summarizes this common sentiment:

I am deeply concerned that the use of GenAI and other AI tools will result in significant compromises in research, including false information, researchers following false leads, reduction in human critical thinking, and reduction in the value of human knowledge/critical thinking/thought. At this stage, GenAI and AI tools pose a threat to the integrity of knowledge that is factual, thoughtful, reflective, nuanced, and critical.

Respondents also noted that GenAI may undermine the development of tacit knowledge that is important for sound research—for example, one wrote that “[t]here are key cognitive scaffolds that come from practice which might be undermined by easy use [of GenAI] for simple tasks.” They also voiced concerns that GenAI may exacerbate the current trend to publish more (Warren 2019). As a respondent explains,

Ideally, GenAI would be used to relieve some of the workload that comes with the ‘publish or perish’ culture. However, if the faster output it helps to generate just increases expectations for that much more output, then it may increase the number of studies done but won’t be of much personal benefit to stressed and overworked researchers/students.

Finally, respondents noted that current training, policies, and ethical standards regarding GenAI are lagging in the field. For instance, one respondent stated that “[p]eople are not trained, or are poorly trained, in GenAI and do not understand how to use it appropriately. This leads to widespread misuse of this tool.” In fact, some suggest that GenAI itself may be a hindrance to such training:

GenAI makes computational analytical tools accessible to more researchers. However, I’m concerned that it might undermine the training in computational methods, as students may no longer feel motivated to develop solid programming skills or a deep understanding of computational theories.

This lack of a deeper understanding as to how GenAI works may create fertile ground for unscrupulous behavior, as noted by one respondent:

I think there are a lot of overly credulous academics, the ones getting the most money for generative AI in research right now, who genuinely think LLMs are magic... they take a totally unscientific approach to what these objects are ontologically and how they work mechanically.

Optimism and Trust

Finally, we measure the optimism and trust about the future of GenAI for research. Scholars are generally more optimistic that GenAI will continue to improve over the next two to three years (40.2 percent of all scholars “somewhat agree” or “strongly agree”)-though there is ambivalence about whether this will have a net positive effect on the field or whether the current advantages of GenAI outweigh the drawbacks for the field. Furthermore, scholars generally believed that GenAI outputs cannot be trusted (only 4.5 percent of all scholars “somewhat agree” or “strongly agree” that GenAI can be trusted, see also Figure C.10 in Appendix C).

Skepticism was also a common theme in the open-ended responses:

I treat GenAI as an inexperienced, mediocre RA. I don’t trust it, but use it to do simple things I can easily check” and “I think we’re at the early stages here. I think it can sometimes be useful to bounce ideas off, but often it returns quite derivative feedback, and I have not found these tools very useful for searching for articles.

We used rake-weighted generalized linear regression models to further probe heterogeneity in trust and optimism.³ Specifically, we examine whether scholars who are more familiar with GenAI are more likely to trust GenAI, believe that GenAI will improve in the next two to three years, or believe that GenAI will have a net positive impact on their field in the next two to three years, and if those effects might be moderated by their use level.

Our outcome variables were the Likert-type “future improvement,” “net positive,” and “trust” questions (see Figure C.10), respectively. All three variables range from “strongly disagree” = 1 to “strongly agree” = 5. The two primary covariates are (1) a binary variable equal to 1 if the respondent reports using GenAI in their research at least weekly and a 0 otherwise, and (2) a “familiarity” composite score computed as the sum across three variables: a question asking respondents to indicate the extent they are familiar with GenAI, understand how GenAI works, and are confident in their ability to use GenAI.⁴ We control for gender identity using a binary variable (1 = not a cis-man and 0 = cis-man).

The regression results are shown in Table D.5 in Appendix D with unstandardized estimates. A model with only familiarity as a predictor of trust (not shown in D.5) had a very small positive effect ($\hat{\beta} = 0.078$ with $t = 3.660$). However, use level is not a significant moderator of this effect (see model 1 in Table D.5). Use level is also not a significant moderator for the net positive outcome variable but is for the future improvement variable. The regression results are summarized with adjusted predictions in Figure 4, where gender identity is held at its modal value (cis-man).

As predictions show in the furthest left panel, trust is generally low regardless of use-level or familiarity. Optimism that GenAI will improve, on the other hand, is generally high regardless of use level or familiarity—with regular users showing a bit more variation by familiarity level (as suggested by the significant interaction term). However, beliefs that this will result in positive effects for sociology are somewhat higher for regular users (about neutral to somewhat agree) than for non-regular users (somewhat disagree). Use level again does not moderate the relationship between familiarity and perceptions of net positive effects, but both active and non-active users tend to be somewhat ambivalent on this issue. Part of this ambivalence can be decoupled from how useful or accurate GenAI may actually be, as it may lead to the delegation of key research tasks. Respondents commented that this arrangement might, in the long term, “produce increasingly poorer researchers,” or undermine critical thinking leading to a “very slippery slope.”

We note that across the three models (see Table D.5), there is a consistent negative association with gender identity. Controlling for familiarity and use levels, non cis-men are less likely than cis-men to trust GenAI output, think that it will improve in the short term, or agree that GenAI will have a net positive effect on their field in the near future. Although not statistically significant in these models, the negative estimates are in line with surveys of the general U.S. adult population, which found that women, nonbinary, and transgender groups hold more negative attitudes about AI (Haimson et al. 2025).

Discussion

We presented findings from a survey asking sociologists and their collaborators how they use (or do not use) GenAI in their research and their attitudes toward the technology more broadly. Thereby, we provided an assessment of generative AI’s role in contemporary sociological research. This assessment is not conclusive, but instead a snapshot that captures a specific moment in a rapidly evolving context.

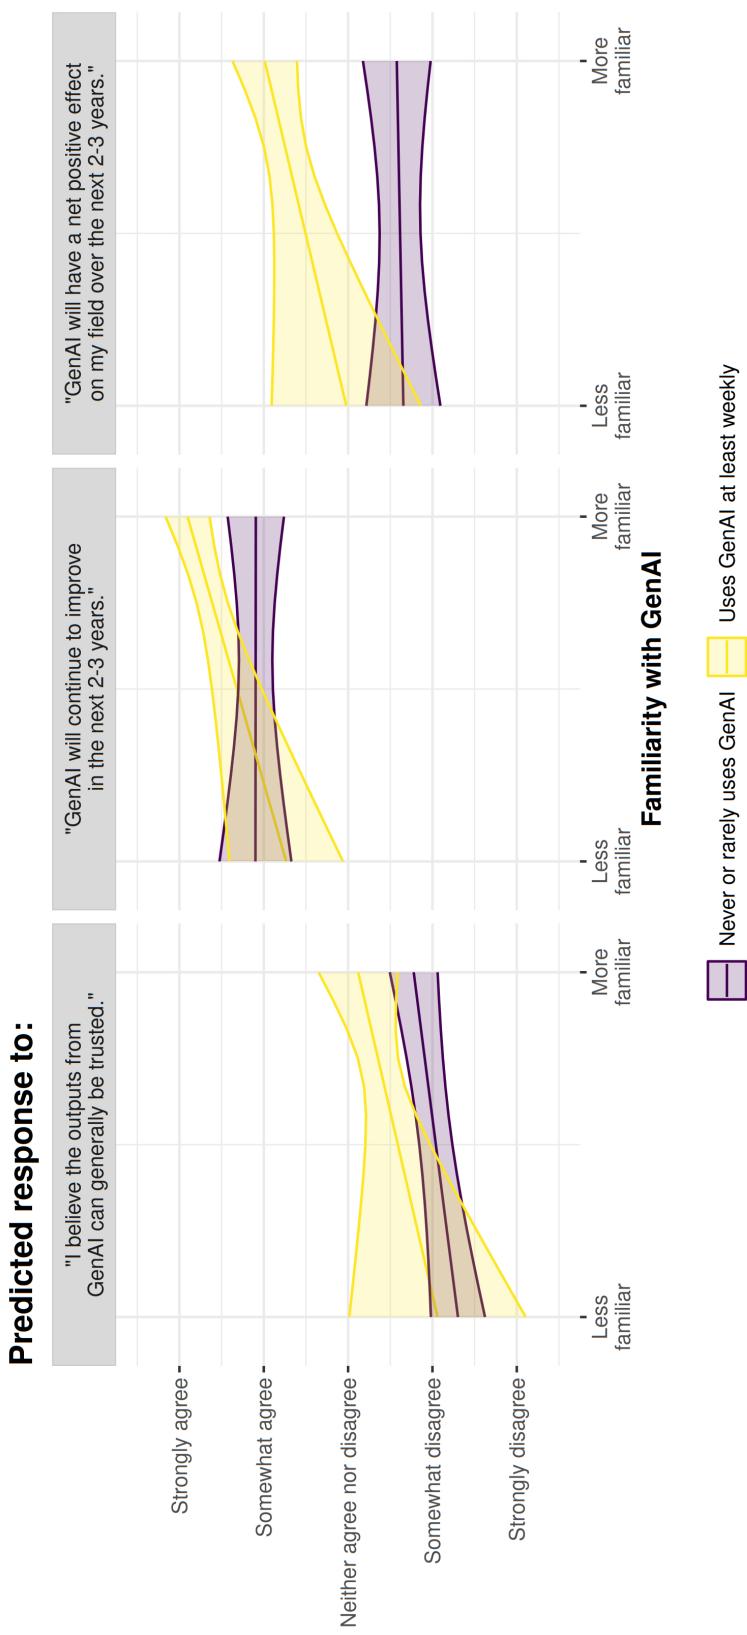


Figure 4: Adjusted predictions of GenAI trust, future improvement, and net positive for research. Note: Error ribbons represent 95 percent confidence intervals. Predictions based on $n = 411$, 407 , and 409 , respectively, after listwise deletion. Gender identity is held at the modal value ("cis-man").

We hope that our survey will serve as a baseline for continued efforts to monitor how and whether GenAI diffuses across the discipline. In closing, we point out important fault lines and questions for the future.

First, our survey suggests that scholars are currently not feeling pressures to adopt GenAI in their own research. The confluence of capital, labor, technology, and politics that formed the GenAI ecosystem has already caused disruption and pressures in a variety of other fields. As the landscape of availability changes—that is, GenAI being incorporated into more applications, such as email, word processors, qualitative data analysis suites, or integrated development environments—and as universities and other organizations implement GenAI policies, it will be important to monitor the extent to which sociologists perceive being expected to use these new tools.

Second, we do not yet know to what extent sociologists might see GenAI as a threat to their sociological expertise. Sociological work may be challenged by GenAI as even the specialties most removed from computational methodology, namely sociological theory, are implicated.⁵ More generally, GenAI threatens knowledge work by broadening the availability of required expertise, weakening occupational closure mechanisms, and reducing bargaining power. Such tensions are further exacerbated by the increased dependence on science and technology in tandem with a decreasing trust in scientists (Eyal 2019). This routinization of knowledge production may eventually allow for the displacement of well-paid professionals for “workers paid a fraction of their wage, flood[ing] the market with inferior products, and impos[ing] unreasonable and punishing working conditions on those” using GenAI (Bender & Hanna 2025, p. 39, c.f., Nelson et al. 2025). Social scientists might come to feel these tensions more acutely than they do now, given that GenAI can present *convincing or even compelling* arguments about human society and social behavior, even if these are inaccurate or entirely fabricated.

Third, attitudes and concerns about GenAI may change considerably if scholars become more familiar with and knowledgeable about the technology. Our study does not allow us to draw any causal inferences, but some of our results could at least be interpreted as hinting in this direction: familiarity was (weakly) associated with trust in GenAI outputs, and those who use GenAI more frequently are generally more optimistic about its effects on the discipline. Sociologists using GenAI usually have interactional expertise as opposed to contributory expertise (Collins & Evans 2019), that is, they know how to use these tools and how to trigger desired outputs through carefully crafted prompts. If more sociologists find ways to make GenAI useful for their work, this will likely have an effect on attitudes and raise new or exacerbate existing concerns. However, use in this sense is not necessary for gaining a deeper understanding of GenAI. That is, there may be sociologists with rich contributory expertise, but who remain skeptical of this technology and refuse to use it in their research.

Finally, our study suggests that there is wide heterogeneity in both usage and attitudes. Some scholars use GenAI daily for a variety of tasks and are optimistic about its effects on the discipline. Many others have never used it and see no role or at most a very limited one for it in sociological research. Surprisingly, these differences did not fall along the lines of division that we had anticipated

(computational/non-computational scholars). Nevertheless, our results point to a significant challenge for the discipline: establishing norms on GenAI's proper use (if any) in the research process. We hope that our study makes a small contribution to addressing this challenge by mapping out how GenAI is currently being used. Future studies might focus more directly on the normative considerations surrounding GenAI use than we did. This could help us locate a consensus on what we should consider appropriate and inappropriate uses of GenAI in sociological research.

Notes

- 1 Tokens are the basic inputs and outputs in natural language processing. In LLMs, tokens are words, combinations of words (including complete sentences), parts of words (e.g., word stems), and/or common non-word characters such as spaces or dashes. See Mielke et al. (2021) for additional details.
- 2 In general, we find very few differences, the most significant of which is in the analysis stage. Specifically and unsurprisingly, computational scholars are using GenAI to help with coding-writing (20.8 percent) and debugging (21.8 percent). Non-computational scholars are also somewhat more likely to use GenAI for grammar and spelling assistance (32.6 percent) and translation (24.8 percent). That said, both computational and non-computational scholars most often use GenAI to help with writing tasks. Though analysis is tied with writing for computational scholars.
- 3 We also ran a series of rake-weighted ordered logistic regression models as a robustness check given the ordinal nature of the outcome variables. They are consistent with the simpler linear model results reported here. These alternative models are shown in Table E.6 in Appendix E. Adjusted predictions for these models are visualized in Figure E.11.
- 4 Each of these three questions was measured using the same Likert scale as the one used for the outcome variables. The three items have high reliability (Cronbach's $\alpha = 0.87$) and are unidimensional according to a principal component analysis ($\lambda_1 = 2.37$).
- 5 Incidentally, the most elite sociology departments in the United States specialize in theory (Elder & Kozlowski 2025).

References

Alvero, AJ, Jinsook Lee, Alejandra Regla-Vargas, René Kizilcec, Thorsten Joachims, anthony lising antonio. 2024. "Large Language Models, Social Demography, and Hegemony: Comparing Authorship in Human and Synthetic Text." *Journal of Big Data*, 11(1). <https://doi.org/10.31235/osf.io/qfx4a>

Alvero, AJ, Dustin Stoltz, Oscar Stuhler, & Marshall Taylor. 2025. *Replication Data for: "Generative AI in Sociological Research: State of the Discipline"* (Version V1). Harvard Dataverse. <https://doi.org/10.7910/DVN/ICXIRP>

An, Yunjo, Ji Hyun Yu, and Shadarra James. 2025. "Investigating the Higher Education Institutions' Guidelines and Policies Regarding the Use of Generative AI in Teaching, Learning, Research, and Administration." *International Journal of Educational Technology in Higher Education* 22(1):1-23. <https://doi.org/10.1186/s41239-025-00507-3>

Andersen, Jens P., Lise Degn, Rachel Fishberg, Ebbe K. Graversen, Serge P. J. M. Horbach, Eyanthia K. Schmidt, Jesper W. Schneider, and Mads P. & Sørensen. 2025. "Generative

Artificial Intelligence (Genai) in the Research Process-A Survey of Researchers' Practices and Perceptions." *Technology in Society* 81:102813. <https://doi.org/10.1016/j.techsoc.2025.102813>

Anthis, Jacy R., Ryan Liu, Sean M. Richardson, Austin C. Kozlowski, Bernard Koch, James Evans, Erik Brynjolfsson, and Michael Bernstein. 2025. "LLM Social Simulations Are a Promising Research Method. *Human-Computer Interaction*. <https://doi.org/10.48550/arXiv.2504.02234>

Argyle, Lisa P., Christopher A. Bail, Ethan C. Busby, Joshua R. Gubler, Thomas Howe, T., Christopher Ryttig, 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>

Argyle, Lisa P., Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Ryttig and David Wingate. 2023. "Out of One, Many: Using Language Models to Simulate Human Samples." *Political Analysis* 31(3): 337–51. <https://doi.org/10.1017/pan.2023.2>

Ashwin, Julian, Aditya Chhabra, and Vijayendra Rao. 2025. "Using Large Language Models for Qualitative Analysis Can Introduce Serious Bias." *Sociological Methods & Research* 00491241251338246. <https://doi.org/10.1177/00491241251338246>

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*. Collins.

Boelaert, Julien, Etienne Ollion, Samuel Coavoux, Ivaylo D. Petev, Patrick Präg. 2025. "Machine Bias. How Do Generative Language Models Answer Opinion Polls?" *Sociological Methods & Research*. https://doi.org/10.31235/osf.io/r2pnb_v2

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>

Broska, David and Daniel A. McFarland. 2024. "Engineering of Inquiry: The 'Transformation' of Social Science Through Generative AI." *SocArXiv*. <https://doi.org/10.31219/osf.io/uy4c7>

Chae, Youngjin and Thomas Davidson. 2025. Large Language Models for Text Classification: From Zero-Shot Learning to Instruction-Tuning. *Sociological Methods & Research*. <https://doi.org/10.1177/00491241251325243>

Chakravorti, Tatiana, Xinyu Wang, Pranav N. Venkit, Sai Koneru, Kevin Munger and Sarah Rajtmajer. 2025. "Social Scientists on the Role of AI in Research." <https://arxiv.org/abs/2506.11255>

Collins, Harry and Robert Evans. 2019. *Rethinking Expertise*. University of Chicago Press.

Davidson, Thomas. 2024. "Start Generating: Harnessing Generative Artificial Intelligence for Sociological Research." *Socius* 10. <https://doi.org/10.1177/23780231241259651>

Davidson, Thomas and Daniel Karel. 2025. "Integrating Generative Artificial Intelligence into Social Science Research." *Sociological Methods & Research* 00491241251339184. <https://doi.org/10.1177/00491241251339184>

DeBell, Matthew and Jon A. Krosnick. 2009. "Computing Weights for American National Election Study Survey Data." *nes012427*. Ann Arbor, MI, Palo Alto, CA: ANES Technical Report Series.

Dorta-González, Pablo, Alexis J. López-Puig, María Dorta-González, and Sara M. González-Betancor. 2024. "Generative Artificial Intelligence Usage by Researchers at Work: Effects of Gender, Career Stage, Type of Workplace, and Perceived Barriers." *Telematics and Informatics* 94:102187. <https://doi.org/10.1016/j.tele.2024.102187>

Elder, Timothy B. and Austin C. Kozlowski. 2025. "Status and Subfield: The Distribution of Sociological Specializations across Departments." *Poetics* 109:101973. <https://doi.org/10.1016/j.poetic.2025.101973>

Eyal, Gil. 2019. *The Crisis of Expertise*. John Wiley & Sons.

Gilardi, Fabrizio, Meysam Alizadeh, and Maël Kubli. 2023. "ChatGPT Outperforms Crowd Workers for Text-Annotation Tasks." *Proceedings of the National Academy of Sciences of the United States of America* 120(30):e2305016120. <https://doi.org/10.1073/pnas.2305016120>

Götz, F. M., R. Maertens, S. Loomba, and S. van der Linden. 2024. "Let the Algorithm Speak: How to Use Neural Networks for Automatic Item Generation in Psychological Scale Development." *Psychological Methods* 29(3):494–518. <https://doi.org/10.1037/met0000540>

Haimson, Oliver. L., Samuel R. Mayworm, Alexis S. Ingber, and Nazanin Andalibi. 2025. "AI Attitudes among Marginalized Populations in the US: Nonbinary, Transgender, and Disabled Individuals Report More Negative AI Attitudes." *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 1224–1237. <https://doi.org/10.1145/3715275.3732081>

Hayes, Adam S. 2025. "'Conversing' with Qualitative Data: Enhancing Qualitative Research Through Large Language Models (LLMs)." *International Journal of Qualitative Methods* 24. <https://doi.org/10.1177/16094069251322346>

Horton, John J. 2023. "Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?" *Nber*. <https://doi.org/10.3386/w31122>

Hrycyna, A., and H. Eassom. 2025. *ExplanAItions: An AI Study* by Wiley (Research report). John Wiley & Sons, Inc. <https://www.wiley.com/content/dam/wiley-com/en/pdfs/about/wiley-explanAItions-ai-study-february-2025vers3.pdf>

Ibrahim, Elida I. and Andrea Voyer. 2024. "The Augmented Qualitative Researcher: Using Generative AI in Qualitative Text Analysis." *Preprints SocArXiv*. <https://doi.org/10.31235/osf.io/gkc8w>

Karell, Daniel, Jeffrey Sachs, and Ryan Barrett. 2024. "Synthetic Duality: A Framework for Analyzing Natural Language Generation's Representation of Social Reality." *SocArXiv*. <https://doi.org/10.31235/osf.io/htxa4>

Kim, Junsol and Byungkyu Lee. 2023. "AI-Augmented Surveys: Leveraging Large Language Models and Surveys for Opinion Prediction." *arXiv[cs.CL]*. <https://doi.org/10.48550/arXiv.2305.09620>

Kozlowski, Austin C. and James Evans. 2025. "Simulating Subjects: The Promise and Peril of AI Stand-ins for Social Agents and Interactions." *SocArXiv*. https://doi.org/10.31235/osf.io/vp3j2_v3

Kozlowski, Austin C., Hyunku Kwon, and James A. Evans. 2024. "In Silico Sociology: Forecasting COVID-19 Polarization with Large Language Models." *arXiv [cs.CY]*. <https://doi.org/10.31235/osf.io/7dfbc>

Kwon, Diana. 2025. "Is It OK for AI to Write Science Papers? Nature Survey Shows Researchers are Split." *Nature* 641(8063):574–8. <https://doi.org/10.1038/d41586-025-01463-8>

Law, Tina and Elizabeth Roberto. 2025. "Generative Multimodal Models for Social Science: An Application with Satellite and Streetscape Imagery." *Sociological Methods & Research* 00491241251339673. <https://doi.org/10.1177/00491241251339673>

Li, Zhi-Qiang, Hui-Lin Xu, Hui-Juan Cao, Zhao-Lan Liu, Yu-Tong Fei, and Jian-Ping Liu. 2024. "Use of Artificial Intelligence in Peer Review among Top 100 Medical Journals." *JAMA Network Open* 7(12):e2448609. <https://doi.org/10.1001/jamanetworkopen.2024.48609>

Lin, Hao and Yongjun Zhang. 2025. "The Risks of Using Large Language Models for Text Annotation in Social Science Research." *arXiv [cs.CL]*. <https://doi.org/10.1177/08944393251366243>

Lumley, Thomas and Alastair Scott. 2017. "Fitting Regression Models to Survey Data." *Statistical Science* 32(2):265–78. <https://doi.org/10.1214/16-STS605>

Maranca, Alessandra R. P., Jihoon Chung, Musashi Hinck, Adam D. Wolsky, Naoki Egami, and Brandon M. Stewart. 2025. "Correcting the Measurement Errors of AI-Assisted Labeling in Image Analysis Using Design-Based Supervised Learning." *Sociological Methods & Research* 54(3). <https://doi.org/10.1177/00491241251333372>

Martin, John L. 2023. "The Ethico-Political Universe of ChatGPT." *Journal of Social Computing* 4(1):1–11. <https://doi.org/10.23919/JSC.2023.0003>

Mielke, Sabrina J., Zaid Alyafeai, Elizabeth Salesky, Colin Raffel, Manan Dey, Matthias Gallé, Arun Raja, Changlei Si, Wilson Y. Lee, Benoît Sagot, and Samson Tan. 2021. "Between Words and Characters: A Brief History of Open-Vocabulary Modeling and Tokenization in NLP." *arXiv preprint arXiv:2112.10508*. <https://doi.org/10.48550/arXiv.2112.10508>

Mu, Y., Ben P. Wu, William Thorne, Ambrose Robinson, Nikolaos Aletras, Carolina Scarton, Kalina Bontcheva, and Xingyi Song. 2023. "Navigating Prompt Complexity for Zero-Shot Classification: A Study of Large Language Models in Computational Social Science." *arXiv [cs.CL]*. <https://doi.org/10.48550/arXiv.2305.14310>

Nelson, Laura K., Nga Than, Leanne Fan, Tina Law, and Leslie McCall. 2025. "Updating 'The Future of Coding': Qualitative Coding with Generative Large Language Models." https://doi.org/10.31235/osf.io/wg82k_v2

Ng, Jeremy Y., Sharleen Maduranayagam, Nirekah Suthakar, Amy Li, Cynthia Lokker, Alfonso Iorio, R. Brian Haynes, and David Moher. 2025. "Attitudes and Perceptions of Medical Researchers towards the Use of Artificial Intelligence Chatbots In the Scientific Process: An International Cross-Sectional Survey." *The Lancet. Digital Health* 7(1):e94–102. [https://doi.org/10.1016/S2589-7500\(24\)00202-4](https://doi.org/10.1016/S2589-7500(24)00202-4)

OUP. (2024). *Researchers and AI: Survey Findings*. <https://fdslive.oup.com/www.oup.com/academic/pdf/Researchers-and-AI-survey-findings.pdf>

Perkowski, Patryk and Ales Marsal. 2024. "Generative AI at Work: Survey Evidence from Three Central Banks." Available at SSRN. <https://doi.org/10.2139/ssrn.4957562>

Schwitter, Nicole. 2025. "Using Large Language Models for Preprocessing and Information Extraction from Unstructured Text: A Proof-Of-Concept Application in the Social Sciences." *Methodological Innovations* 18(1). <https://doi.org/10.1177/20597991251313876>

Sourati, Zhivar, Farzan Karimi-Malekabadi, Meltem Ozcan, Colin McDaniel, Alireza Ziabari, Jackson Trager, Ala Tak, Meng Chen, Fred Morstatter, and Morteza Dehghani. 2025. "The Shrinking Landscape of Linguistic Diversity in the Age of Large Language Models." *arXiv Preprint arXiv:2502.11266*. <https://doi.org/10.48550/arXiv.2502.11266>

Stuhler, Oscar, Cat D. Ton, and Etienne Ollion. 2025. "From Codebooks to Prompt-books: Extracting Information from Text with Generative Large Language Models."

Sociological Methods & Research 54(3):00491241251336794. <https://doi.org/10.1177/00491241251336794>

Van Noorden, Richard and Jeffrey M. Perkel. 2023. "AI and Science: What 1,600 Researchers Think." *Nature* 621(7980):672–5. <https://doi.org/10.1038/d41586-023-02980-0>

Waight, Hannah C., Brandon M. Stewart, Margaret E. Roberts, Solomon Messing, Eddie Yang, Joshua A. Tucker, and Yin Yuan. 2024. Propaganda BIas.

Warren, John R. 2019. "How Much Do You Have to Publish to Get a Job in a Top Sociology Department? Or to Get Tenure? Trends over a Generation." *Sociological Science* 6:172–96. <https://doi.org/10.15195/v6.a7>

Watermeyer, Richard, Donna Lanclos, Lawrie Phipps, Hanne Shapiro, Danielle Guizzo, and Cathryn Knight. 2024. Academics' Weak(ening) Resistance to Generative AI: The Cause and Cost of Prestige? *Postdigital Science and Education*. <https://doi.org/10.1007/s42438-024-00524-x>

Watermeyer, Richard, Lawrie Phipps, Donna Lanclos, and Cathryn Knight. 2024. "Generative AI and the Automating of Academia." *Postdigital Science and Education* 6(2):446–66. <https://doi.org/10.1007/s42438-023-00440-6>

Xie, Yu and Sofia Avila. 2025. "The Social Impact of Generative LLM-Based AI." *Chinese Journal of Sociology* 11(1):31–57. <https://doi.org/10.1177/2057150X251315997>

Zhang, Simone, Janet Xu, and A.J. Alvero. 2025. "Generative AI Meets Open-Ended Survey Responses: Research Participant Use of AI and Homogenization." *Sociological Methods & Research* 54(3). <https://doi.org/10.1177/00491241251327130>

Zhang, Yanbo, Sumeer A. Khan, Adnan Mahmud, Huck Yang, Alexander Lavin, Michael Levin, Jeremy Frey, Jared Dunnmon, James Evans, Alan Bundy, Saso Dzeroski, Jesper Tegner, and Hector Zenil. 2025. "Exploring the Role of Large Language Models in the Scientific Method: From Hypothesis to Discovery." *npj Artificial Intelligence* 1(1):14. <https://doi.org/10.1038/s44387-025-00019-5>

Acknowledgments: All authors contributed equally. We thank all the respondents of our survey for being generous with their time. We are indebted to Kim Weeden and Cat Dang Ton for giving us crucial comments on an early version of this article. We are also grateful for the important feedback we received at the *ASA Session on Culture and Computational Social Science*, the *Sociological Science Conference* at Cornell University, the *International Network of Analytical Sociology Conference* at Columbia University, the *Institute for Analytical Sociology Symposium* at Linköping University, and the *Culture and Action Network* at the University of Chicago.

AJ Alvero: Center for Data Science for Enterprise and Society, Cornell University.
E-mail: ajalvero@cornell.edu.

Dustin S. Stoltz: Department of Sociology and Anthropology, Lehigh University.
E-mail: dss219@lehigh.edu.

Oscar Stuhler: Department of Sociology, Northwestern University.
E-mail: oms@northwestern.edu.

Marshall A. Taylor: Department of Sociology, New Mexico State University.
E-mail: mtaylor2@nmsu.edu.