

## Social Scientists on the Role of AI in Research

Tatiana Chakravorti<sup>1</sup>, Xinyu Wang<sup>1</sup>, Pranav Narayanan Venkit<sup>1</sup>, Sai Koneru<sup>1</sup>, Kevin Munger<sup>2</sup>, Sarah Rajtmajer<sup>1</sup>

<sup>1</sup>Pennsylvania State University, USA

<sup>2</sup>European University Institute, Florence

tfc5416@psu.edu, xzw5184@psu.edu, pranav.venkit@psu.edu, sdk96@psu.edu, kevinmunger@gmail.com, smr48@psu.edu

### Abstract

The integration of artificial intelligence (AI) into social science research practices raises significant technological, methodological, and ethical issues. We present a community-centric study drawing on 284 survey responses and 15 semi-structured interviews with social scientists, describing their familiarity with, perceptions of the usefulness of, and ethical concerns about the use of AI in their field. A crucial innovation in study design is to split our survey sample in half, providing the same questions to each – but randomizing whether participants were asked about "AI" or "Machine Learning" (ML). We find that the use of AI in research settings has increased significantly among social scientists in step with the widespread popularity of generative AI (genAI). These tools have been used for a range of tasks, from summarizing literature reviews to drafting research papers. Some respondents used these tools out of curiosity but were dissatisfied with the results, while others have now integrated them into their typical workflows. Participants, however, also reported concerns with the use of AI in research contexts. This is a departure from more traditional ML algorithms which they view as statistically grounded. Participants express greater trust in ML, citing its relative transparency compared to black-box genAI systems. Ethical concerns, particularly around automation bias, deskilling, research misconduct, complex interpretability, and representational harm, are raised in relation to genAI. We situate these findings within broader sociotechnical debates, arguing that responsible integration of AI in social science research requires more than technical solutions: it demands a rethinking of research values, human-centered design, and institutional support structures. To guide this transition, we offer recommendations for AI developers, researchers, educators, and policymakers focusing on explainability, transparency, ethical safeguards, sustainability, and the integration of lived experiences into AI design and evaluation processes.

### 1 Introduction

The rapid advancement of artificial intelligence (AI) and machine learning (ML) technologies is transforming the landscape of social science scholarship and research (Bail 2024). The emergence of large language models (LLMs) such as ChatGPT and Gemini promise new efficiencies but bring challenges as well (Salganik 2019; Ziems et al. 2024; Venkit

et al. 2024b). The education sector and academia have felt the immediate impact of LLMs on teaching practices, experimental workflows, manuscript writing, and editorial processes (Celik et al. 2022; Wu et al. 2024b; Omran Zailuddin et al. 2024). While, researchers have increasingly adopted these tools for a wide range of applications, including data collection, annotation, and analyses, as well as a variety of coding tasks previously inaccessible to researchers without extensive programming backgrounds (Zhang 2023; Dandamudi, Sajja, and Khanna 2025; Sadiku et al. 2021; Fan et al. 2023; Deranty and Corbin 2024; Obreja, Rughinis, and Rosner 2024; Bail 2024). Researchers are exploring entirely new paradigms, such as simulating social interactions, generating synthetic interview data, and prototyping new theories (Grossmann et al. 2023).

However, the adoption of AI into the social sciences is not a simple upgrade, but a fundamental sociotechnical shift, entailing negotiations across practical, methodological, and ethical lines. AI and ML have the potential to produce or reinforce existing biases (O’Neil 2017; Bender et al. 2021). Design choices shape what kinds of knowledge are produced, whose perspectives are prioritized, and how research problems are framed (Selbst et al. 2019). Further, the social sciences are methodologically diverse, including qualitative, quantitative, and mixed methods, e.g., ethnography, interviews, experiments, agent-based modeling, and observational studies (Wulff et al. 2022; Ligo et al. 2021). This diversity presents unique challenges and opportunities when integrating AI, into the epistemological standards and interpretive needs of different sub-fields.

Our work seeks to understand how social science researchers are navigating AI and ML, scaffolded by the following research questions (RQs):

- **RQ1:** How are AI and ML reshaping research practices in the social sciences?
- **RQ2:** What are the common ethical, methodological, and practical concerns among social scientists regarding the use of AI and ML in research?
- **RQ3:** How are social scientists negotiating the role of human judgment in the context of increasingly autonomous tools?

Drawing on a *mixed-methods* approach that includes surveys followed by semi-structured interviews, this study ex-

plores how social science researchers navigate excitement and tensions around automation when working with AI and ML. The same survey was given to all participants, *with one crucial element randomized: half of the participants were asked their perceptions about "AI," and the other half about "ML."* This design allows us to differentiate these two often-conflated developments in social science practice.

We find that the use of AI in research settings is common among social scientists, particularly with the growing popularity of genAI. These tools are employed for a variety of tasks, from summarizing literature to drafting research papers, whereas ML was used for mostly prediction or classification tasks. Participants raised concerns about deskilling, lack of standardization, ethical risks, bias, model interpretability, inadequate training, and over-reliance on AI systems. *Ethical concerns were less pronounced when we asked about ML techniques than when we asked about AI – suggesting evolution in the distribution of attention with the rise of genAI.* These insights underscore the importance of interdisciplinary collaboration and community-centered research paradigms, where AI integration is informed by the experiences and practical needs of social scientists. To support this shift, we propose actionable recommendations for four key stakeholder groups: developers; researchers; educators; and policymakers. Ultimately, we advocate for reimagining AI not as a replacement for human insight, but as a collaborator in processes of knowledge production.

## 2 Related Work

### 2.1 AI and ML in the Social Sciences

Artificial intelligence and machine learning are actively transforming scientific processes across the research lifecycle (Chubb, Cowling, and Reed 2022; Xu et al. 2021). AI integration into the social sciences has facilitated analyses of large scale datasets, uncovering novel insights and opening new research directions (Grimmer 2015; Grossmann et al. 2023). This integration has been accompanied by robust discussions about the significance of the moment and questions about the role these technologies should play in social science research (Iorga, Bocse, and Olar 2023; Williams 2024).

Initially, ML was viewed as a powerful collection of computational tools required for navigating the rapidly expanding availability of “big data” (Salganik 2019). Given that new datasets pose challenges for conventional statistical analyses due to scale or complexity, ML techniques were framed primarily as advanced methods for pattern discovery, prediction, and classification (Varian 2014; Grimmer, Roberts, and Stewart 2021). This positioned ML as an extension of the quantitative toolkit, a required adaptation to deal with new types of evidence (Einav and Levin 2014).

As integration grew, the role of ML began to diversify and sometimes generated discussions centering on the distinction between the predictive strength of ML models and the traditional emphasis of social sciences research on causal explanation (Hofman et al. 2021; Shmueli 2010). Reflecting these discussions, efforts emerged to adapt ML techniques specifically for causal analysis, seeking ways to improve the estimation and understanding of causal relation-

ships, particularly when working with complex observational data (Athey and Imbens 2017; Chernozhukov et al. 2018). Alongside these developments, the language used to describe these approaches often remained fluid and overlapping. Researchers might find the distinctions between “ML”, “AI”, “computational social science,” and advanced statistics fuzzy, with usage varying depending on their particular training or the specific domain (Edelmann et al. 2020; Grimmer, Roberts, and Stewart 2021).

The advent of generative AI, particularly large language models, has arguably triggered a significant shift in interpretation (Zhang 2023). For many researchers, “AI” is now increasingly understood through the lens of these generative capabilities – systems that can produce human-like text, code, and interact in complex ways (Bail 2023). This interpretation casts AI not merely as an analytical tool operating on existing data, but as an active generator of content and potentially even a collaborator or simulator of human behavior (Anthis et al. 2025; Jansen, Jung, and Salminen 2023; Dominguez-Olmedo, Hardt, and Mendler-Dünner 2024)

### 2.2 Ethical Concerns About AI

As AI and ML technologies become increasingly embedded in research and daily life, conversations about their benefits are frequently accompanied by growing concerns about their potential to cause harm. One of the most pressing and recurring concerns is bias—both in how these systems are developed and how they are deployed (O’Neil 2017; Bender et al. 2021). These issues are exacerbated by the opaque, ‘black-box’ nature of many AI systems, which are often proprietary and lack transparency in terms of architecture, training methodology, and data provenance (O’Neil 2017; Bender and Friedman 2018; Castelveccchi 2016; Von Eschenbach 2021). At the heart of these concerns lies the data that is used to train and develop these systems (Hacker 2021; Sambasivan et al. 2021). GenAI models, especially large language models, are typically trained on vast corpora scraped from the internet and other publicly available sources (Narayanan Venkit et al. 2023; Paullada et al. 2021). While this scale and diversity might suggest objectivity or comprehensiveness, these datasets tend to amplify dominant narratives—often Western, English-speaking, and majority-population-centric—while marginalizing or misrepresenting voices from minoritized communities (Gautam, Venkit, and Ghosh 2024; Das et al. 2024). This results in representational harms that can take the form of stereotyping, cultural misappropriation, and the erasure of difference (Dev et al. 2022; Blodgett et al. 2022). Bias in AI is not limited to model outputs but permeates every layer of these systems, from embedding representations to classification algorithms to generative models (Bolukbasi et al. 2016; Mehrabi et al. 2021; Ntoutsi et al. 2020).

### 2.3 Sociotechnical Concerns About AI

As AI systems become more integrated into decision-making pipelines, the illusion that they are solely technical systems becomes harder to sustain. Researchers have noted how technical design decisions intersect with social contexts, reinforcing existing inequalities or introducing new

forms of algorithmic harm from biased behaviors (Dolata, Feuerriegel, and Schwabe 2022; Narayanan Venkit 2023; Ehsan et al. 2021). For instance, in tasks like recidivism prediction, toxicity detection, and facial recognition, biased outputs can lead to discriminatory outcomes when these systems are deployed in real-world social infrastructures (Bender et al. 2021; O’Neil 2017). Work in HCI and science and technology studies (STS) examines these harms through the lens of sociotechnical entanglement, emphasizing how users, designers, and affected communities experience and respond to algorithmic systems. Research frameworks now increasingly move beyond surface-level performance metrics to interrogate how harm materializes through misinformation, stereotyping, and exoticism—phenomena that not only distort representation but also shape how entire communities are understood, regulated, or excluded (Ghosh et al. 2024; Dev et al. 2022; Blodgett et al. 2022).

*Much of the existing literature has examined harms at the societal or individual user level, yet there remains a gap in understanding how AI technologies impact academic efforts and usage, particularly within the social sciences.* These disciplines, which traditionally rely on interpretive, qualitative, and mixed-method approaches, are being reshaped by tools that offer automation, scale, and new modalities of expression. We respond by offering a detailed, mixed-methods investigation into how social science researchers perceive, adopt, and contend with AI and ML technologies. We explore how scholars navigate the affordances and the ethical entanglements of these tools, focusing on impacts across research design, authorship, epistemology, and judgment.

### 3 Methods

#### 3.1 Study Design

We adopted an *explanatory sequential mixed method technique* (Ivankova, Creswell, and Stick 2006) wherein quantitative (survey) data is collected and analyzed first, followed by qualitative data collection to explore these findings in further depth. *The complete survey instruments (both versions) and the interview topic guide are provided in Supplemental Materials.*

**Survey design.** Our survey followed standard exploratory survey design methodology (Rosen et al. 2013; Baker 2016; Van Noorden and Perkel 2023; Chakravorti, Koneru, and Rajtmajer 2025), with one crucial innovation. *The survey questionnaire was held constant, but the target of inquiry was randomized; one-half of the surveys asked about “Artificial Intelligence,” and the other half of the surveys asked all the same questions about “Machine Learning” instead of AI.* Our aim was to understand how social science researchers’ perceptions, acceptance, and ethical considerations changed over the significant advancement in ML to generative AI.

The survey was organized as follows. Participants were presented with information about the project and consented to participate. **Demographic information:** Participants were asked to provide information related to gender, age, ethnicity, academic position, field of research, and years of experience. The body of the survey contained 4 closed-ended questions and 5 open-ended ques-

tions. **Closed-ended questions:** These asked participants’ to rate their usage of, acceptance of, and familiarity with AI/ML. **Open-ended questions:** These aimed to complement quantitative data with stated perceptions, experiences, and ethical and methodological challenges using AI/ML technologies in their research. A pilot version of the survey (both variants) was created and pretested by 5 social science researchers before deployment.

**Interview design.** All interviews were carried out via Zoom. The length of interviews ranged from 15 to 40 minutes; the majority were completed within 30 minutes.

Interviews were organized as follows. **Understanding AI and ML:** This set of questions is asked to understand what are AI and ML in terms of their understanding and how they differentiate these two from each other as we have not asked for any definition of them during the survey. **Ethical concerns about AI and ML:** With these questions we investigated is the ethical concerns and trust really increase with AI with elaborated examples to validate our survey findings.

**Value of human judgment:** In this section, we asked participants how they are currently using these techniques and how they perceive the role of human judgment in an increasingly autonomous world. We also explored their preferences regarding the types of systems they choose to use.

#### 3.2 Participant Recruitment and Compensation

**Survey recruitment.** We recruited research-active academics across rank, i.e., assistant, associate, and full professors, post-doctorates and research associates, and graduate students working in US universities in the social sciences. We selected 30 universities randomly from the top 100 based on the US News and World Report (2003) rankings.

We directly emailed researchers listed on departmental web pages, targeting the following disciplines: economics; political science; education; psychology; sociology; and marketing. We collected email IDs via web scraping of Universities’ directories. *In total, we emailed 8,000 social science researchers. One-half received the survey on AI, while the other half received a survey on ML. We received 348 responses;* however, some of these responses were incomplete and thus excluded from our analysis. We only considered responses with a 100% completion rate for the closed-ended questions, resulting in **284 valid responses**, or 3.56% of the total surveys sent (143 AI survey; 141 ML survey).

At the end of the survey, participants were given the opportunity to be entered into a draw to win one of ten (10) \$100 gift vouchers.

**Interview recruitment.** Unrelated to the compensation opportunity, survey participants were invited to provide their email IDs for a follow-up interview study. We emailed all the participants with a stated interest in the interview study and conducted **15 semi-structured interviews** (Varanasi et al. 2022; Thakkar et al. 2022) to further validate quantitative findings from the survey study. Each participant was compensated \$30 for their time.

#### 3.3 Participant Demographics

Demographic characteristics of survey participants is presented in Table 1. Responses were received from individuals

across disciplines and academic ranks. Our respondents are not representative of the US population; however, there are no significant differences in demographics across the ML and AI groups. Only one person from the ML survey did not disclose the demographic details but he did provide other responses. Demographic data for the 15 interview participants is provided in Table 3.

### 3.4 Data Analysis

For the *quantitative data analysis*, closed-ended questions were analyzed using **descriptive statistics** (Kaur, Stoltzfus, and Yellapu 2018) and exploratory data visualizations. For the *qualitative data analysis*, first two authors analyzed the semi-structured interview transcripts using **thematic analysis** as outlined by Blandford (Blandford, Furniss, and Makri 2016). This qualitative data analysis method involves thoroughly reading the transcripts to identify patterns across the dataset and derive themes related to the research questions. We employed a collaborative and iterative coding process (Ding et al. 2022; Huang, Vitak, and Tausczik 2020). Initially, the first two authors read the interview transcripts multiple times to become familiar with the data. Open coding was then conducted to identify initial codes by the first and second authors. These codes were then organized into categories based on similarities and relationships. Finally, we refined these categories, assigned names to each theme, and developed a conceptual framework to address our research questions with input from the entire team. A similar approach was taken for the open-ended survey responses.

## 4 RQ1: AI and ML Reshaping Social Science Research Practices

### 4.1 Understanding of AI and ML

In the survey, participants were asked about their familiarity with AI and ML. Figure 1 shows the differences across our four outcome variables. Among respondents to the ML survey, 41.84% reported being slightly familiar, and 7.09% reported not being familiar at all with ML techniques. In contrast, familiarity was higher in the AI survey: only 22.37% of respondents indicated slight familiarity and just 3.49% reported no familiarity with AI techniques. We also examined how frequently respondents use AI/ML in their research. Notably, 40.43% had never used ML techniques, compared to 27.27% who had never used AI techniques.

Fig 2 displays the results of eight separate regressions, two (one about AI, one about ML) for each of the four outcome variables. The most statistically significant finding emerges in the position-related predictors, each of which is displayed relative to PhD students (the reference category). Here, *there is no significant variation in attitudes towards ML across career stages – but there is significant variation in attitudes towards AI. PhD students report greater acceptance (bottom left panel) than social scientists at any other career stage. They also report a higher frequency of usage and perceive AI to be more useful than either assistant or full professors.*

Gender dynamics present another intriguing dimension. *Female respondents report lower frequency of use and lower*

Characteristic	Category	ML Survey( <i>n</i> )	AI Survey( <i>n</i> )
Gender	Female	71	74
	Male	62	63
	Non-binary	4	3
	Prefer not to say	3	3
	Others	1	0
Ethnicity	White	88	98
	Asian	23	24
	Others	20	12
	Black or African American	10	9
Academic Rank	Graduate Student	54	65
	Assistant Professor	18	14
	Associate Professor	19	22
	Professor	35	34
	Others	15	7
	Postdoc	0	1
Age Group	<25	10	14
	25–35	52	56
	36–45	30	28
	>45	49	45

Table 1: Survey respondent demographics by survey type (ML:  $n = 141$ , AI:  $n = 143$ )

*familiarity with both AI and ML, compared to male respondents.* There are not, however, any differences when we ask about acceptance or perceived usefulness of these technologies. The non-binary or alternative gender category shows additional variation, though the confidence intervals suggest less certainty about the precise nature of this relationship. There are no significant differences between white and non-white respondents; we lack the sample size to break the latter category down into more granular detail.

In the follow-up interviews, we explored how social science researchers perceive AI and ML technologies—and what those terms mean to them. We observed that when participants were asked to explain AI in their own words, they primarily thought of generative AI tools or large language models. They tended to view machine learning as an algorithm or method that helps with classification and prediction for different applications, while they imagined AI as a product or as a tool that uses large language models and can mimic human behavior.

*I think that when we are talking about AI, in the recent discourse, we are talking about the large language models. I use a lot of ChatGPT in particular for all sorts of possible prompts. –P1*

P10 characterized AI deep-learning algorithms leveraging large datasets. P10 had a much deeper understanding of AI and described LLMs as a type of AI working with language. Except for P3, P5, and P7, other participants reported that LLMs come to their minds first when they hear AI. Participant P5 differed from others; as she sees AI as ChatGPT and doesn't have any understanding of ML.

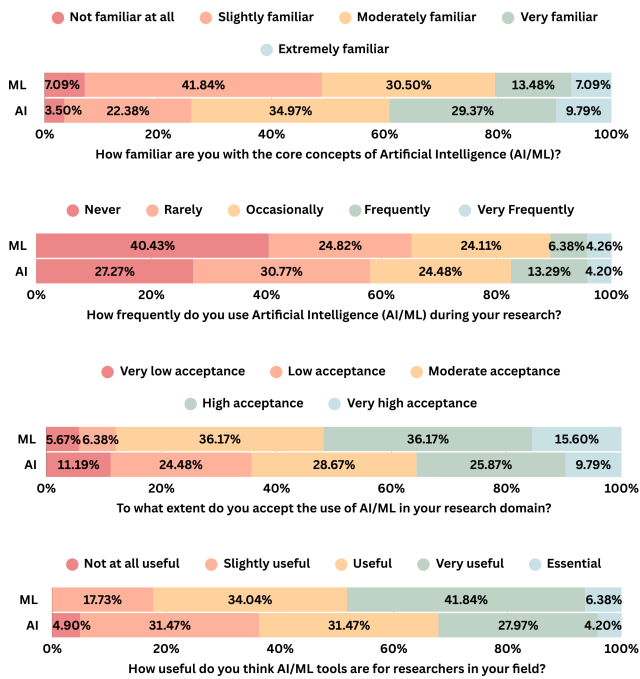


Figure 1: Familiarity, frequency of use, acceptance, usefulness of AI/ML technologies

*It is typically, electronically, some type of system like ChatGPT or another brand to ask questions, and prompts, and have answers spit out at you. I have no idea what is machine learning is. –P5*

All participants except P5 mentioned ML is a part of AI but all AI methods do not belong to ML. They all mentioned AI has much broader and more advanced applications. Some participants have hesitations about their understanding:

*I don't know if I know the nuance difference between artificial intelligence and machine learning. I think of machine learning as being much more predictive or providing output based on a studied corpus of text or a body. I assume artificial intelligence probably goes beyond that in scope. But I'm not sure I know the nuanced difference. –P7*

## 4.2 Usefulness and Acceptance of AI and ML

When participants were asked about their acceptance of AI/ML technologies in their research domain during the survey, we found that ML was better accepted than AI. A combined 51.77% of respondents expressed high or very high acceptance of ML, compared to 35.66% for AI. Similarly, perceptions of usefulness showed that 41.84% of respondents considered ML to be very useful, while only 27.97% said the same for AI. This is described in Figure 1. *These differences were statistically significant ( $p < .01$ ). All statistical results are described in Table 2.*

Survey participants also shared how they use AI tools to support specific research tasks, particularly those related to information processing and literature review. Several responses pointed to the usefulness of AI tools for summarizing research papers:




Figure 2: Regression analysis based on demographics of the frequency, familiarity, acceptance, and usefulness of AI and ML methods

*I have used Consensus to initiate a brief (informal) search — I appreciate being able to quickly pull several relevant papers (with summarization as part of the tool) and then use those as a starting point for a more manual search. –AI Survey*

Others described using AI tools to explore unfamiliar domains, identify key scholars, or clarify concepts:

*Yes, I have used ChatGPT to identify top scholars and seminal works in literature I am not familiar with. –AI Survey*

All interview participants noted that the emergence of generative AI tools has significantly increased the use of AI in social science research—from summarization and literature review, to writing codes, and drafting the text of research papers. Participant P1 uses genAI tools for coding and he reported it helps a lot with debugging. P10 also mentioned that he used ChatGPT specifically to debug code in R. He also mentioned that occasionally he used it for idea generation.

*I try to use ChatGPT pretty conservatively... mostly for helping me debug code or to generate an idea. –P10*

P9 sees potential in genAI tools for tasks like transcription or first-pass annotation. But she also mentioned that she does not think these tools are good for writing papers or conducting qualitative analysis, and emphasizes that it is not yet expected or normalized in her field. On the other hand, P8 mentioned that he uses genAI tools to help automate literature review tasks, like extracting study objectives from pa-

pers, after validating a subset of results manually. He accepts AI performance only after comparing it to a human validated ground truth dataset and allows it to scale after reaching acceptable accuracy (e.g.,99%).

*We do annotation for 50 papers. We do it manually. When we reach 99% accuracy, and then we allow the system to do the other 1,000 papers –P8*

P3 mentioned that AI tools are very helpful in grant writing and paper writing.

*it has made things slightly easier in terms of like, organizing your thoughts, writing your papers, writing your grants, and other stuff. It helps you pretty well to restructure your grant in a way that would be useful for us. –P3*

In sum, the ways in which these tools, particularly genAI, are currently depend to a great extent on individual researchers' trust and adoption.

## 5 RQ2: Ethical, Methodological and Practical Concerns

### 5.1 Technological, Operational and Ethical Concerns

Both ML and AI raise ethical issues. However, in the ML survey, we observe that there are generally fewer ethical concerns. Most respondents to the AI survey expressed some level of concern or hesitation about the use of AI in social science, centering around potentially threatening foundational aspects of human-centered research.

*I am okay with machine learning in which the research trains the models, but I am deeply skeptical of generative AI built on non-transparent and proprietary training data – AI Survey*

**Limited Understanding, Training, and Resources** In both surveys, several researchers acknowledged that they do not have sufficient knowledge about AI and ML to effectively evaluate their usefulness in social science research. This lack of knowledge extends not just to how these systems work technically but also to their appropriate application in research design, methodology, and interpretation. A dominant concern emerging from the responses is that individuals—particularly students and early-career researchers—may adopt these tools without a proper conceptual understanding of how these systems generate outputs.

*Lack of understanding of relationship between ML and statistics, inadequate graduate training – ML Survey*

In educational settings, this concern is amplified with the use of genAI. Students use AI to “cheat” rather than learn according to many respondents. Students who lean too heavily on AI tools might never learn to troubleshoot code, critically assess data, or synthesize information independently. Instead of gaining fluency in research methods, they risk becoming passive users of AI outputs.

*My hesitation is students not understanding how to critically think for themselves. –AI survey*

According to one participant, these tools can artificially enhance the quality of a CV, which may not accurately reflect the candidate's actual qualifications.

*One of my students used ChatGPT to create her resume. She never reviewed the resume, so she just assumed that it was fine, and then asked me to write a recommendation letter for her; I found that the resume vastly oversold what she actually accomplished. She said she was writing papers, running her own analyses, and coming up with the theoretical models. Obviously, these things were not true. The embedded things here have huge ethical challenges for me. She trusts it so much that it's actually leading her to abandon critical thinking of the outputs intentionally. –P10*

**Ethical Risks and Bias** Many researchers expressed serious ethical concerns regarding the use of AI in social science research. Although ethical concerns are present in both responses they appeared more frequently in the AI survey.

*The field is rapidly expanding and I worry about ethics informing that expansion. I also worry about bias and the diversity and quality of the training data. – AI Survey*

Respondents highlighted that AI models are trained on existing data, which often reflects the historical and societal inequalities embedded in those datasets. For instance, they mentioned these tools used in facial recognition or predictive modeling have been shown to reproduce and even amplify racial, gender, and socioeconomic biases. This is particularly troubling in social science contexts, where research outcomes can influence policy decisions, social interventions, and public opinion. Several researchers pointed out that AI systems may inadvertently perpetuate discriminatory practices if applied uncritically, especially in sensitive areas like policing, surveillance, and healthcare.

*It's a black box that transmutes bias in the training data into bias in the outputs in a difficult to audit/detect way. It's also incredibly environmentally destructive. – AI Survey*

Concerns were also raised about the ethical sourcing of training data. Some researchers questioned whether data used to train AI models were obtained with proper consent, particularly when derived from personal or publicly available information mentioned by P12.

*“There are IRB concerns when it comes to what you would put into some of these models. . . run the risk of data sharing and privacy issues related to where that information is stored, how it is used, and how it is going to be used to train other models.” –P12*

Other growing concerns mentioned during the interviews included: plagiarism; academic integrity; and misconduct in peer review. Interviewees noted ethical gray areas where it is unclear how to appropriately cite AI-generated content or how to ensure the originality of work. Some researchers worried this could lead to a dilution of scholarly standards, especially among students and less experienced scholars who might misuse AI-generated content without a clear understanding of academic norms.

*We haven't reached the point of having clear norms about when LLMs are acceptable and not acceptable, including things like peer review. For example, I've spoken with colleagues and editors of journals who are getting, like AI-generated reviews. The use case that we just mutually agree that you know we should not use LLMs, is write a peer review –P14*

Survey Question	Type	Mean	Sd	N	p-value
How familiar are you with the core concepts of AI/ML?	ML	2.72	1.02	141	<b>0.00012</b>
	AI	3.19	1.01	143	
How frequently do you use AI/ML during your research?	ML	2.09	1.13	141	<b>0.046</b>
	AI	2.36	1.14	143	
To what extent do you accept the use of AI/ML in your research domain?	ML	3.49	1.02	141	<b>0.0001</b>
	AI	2.98	1.16	143	
How useful do you think AI/ML tools are for researchers in your field?	ML	3.37	0.85	141	<b>0.00014</b>
	AI	2.95	0.98	143	

Table 2: Perceptions, usage, and acceptance of AI and ML among social science researchers

ID	Gender	Profession	Research field
P1	Male	PhD student	Economics
P2	Male	PhD student	Political Science
P3	Female	PhD student	Psychology
P4	Male	PhD student	Political Science
P5	Female	Associate Prof	Education
P6	Male	PhD student	Sociology
P7	Female	Assistant Prof	Psychology
P8	Male	PhD student	Psychology
P9	Female	PhD student	Psychology
P10	Male	Assistant Prof	Sociology
P11	Male	PhD student	Economics
P12	Female	PhD student	Psychology
P13	Male	PhD student	Economics
P14	Female	Assistant Prof	Political Science
P15	Female	PhD Student	Sociology

Table 3: Gender, rank, and field of interview participants

The cost of subscriptions are also seen as a barrier to equitable adoption:

*Amongst people who use the tools, there's a possibility of inequity... there's a subscription model for the latest version of ChatGPT... \$40 a month, which is a lot of money. –P15*

Fabricating data, generating false citations, or creating papers that appear rigorous but lack substantive analysis could flood academic publication channels with low-quality or misleading research, undermining the credibility of scholarly communication.

*Hype chasing is a major issue because it "sucks the oxygen" out of the room in terms of resources and time. Furthermore, it enables bad actors to plug up publication pipelines with garbage. Not to mention generative text and images are changing the incentive structures of producing content online – it will be very difficult to find content actually produced by people who care about the content next to a mountain of botshit. – AI Survey*

**Lack of Standardization and Best Practice** A prominent concern voiced by researchers was the lack of standardization and established best practices in the use of AI tools for social science research. This issue was raised both explicitly and implicitly in multiple responses, reflecting a widespread unease about how AI is currently being integrated into academic workflows without clear guidelines or benchmarks. New tools, models, and capabilities are being released continuously, often without peer-reviewed validation or clarity

around their appropriate application. This leaves researchers unsure about which tools are reliable, what constitutes appropriate use, and how to evaluate AI-generated outputs within the standards of rigorous social science inquiry.

*Yes, I agree that there is currently a lack of standardization in terms of what can be used and what cannot be used – AI Survey*

This also leads to concerns about transparency and documentation. Without clear standards, researchers may omit or inconsistently report how AI tools were used in their methods sections, making it difficult for peers to evaluate the integrity of the work or to reproduce the findings mentioned by 60% of the interview participants as well.

**Over-Reliance and Overuse** Over-reliance and overuse are major concerns reported across both surveys and interviews. Specific technologies such as AI-based decision support systems, autonomous systems, or recommendation systems show high rates of over-reliance compared to others. As users grew accustomed to the convenience and efficiency offered by automated tools, they started to begin to trust these systems excessively, often without verifying outputs or questioning decisions. This automation bias can reduce critical thinking skills, diminish situational awareness, and reduce the ability to detect errors, especially in complex or dynamic environments.

*I think the overuse of ML has become more prevalent in research, even outside my field. As the ML technique is evolving so fast, some researchers might use it without understanding it, or use it when it is not necessary. –ML survey*

**Difficulty in Model Interpretability** Difficulty in model interpretability and high complexity are critical challenges highlighted in the data. As AI models, particularly deep learning systems, become more complex, they often operate as "black boxes," producing outputs without offering clear insights into how decisions are made. This lack of transparency can undermine user trust, hinder effective oversight, and complicate error diagnosis.

*Interpretability is certainly key to much social scientific research, but ML models are increasingly complex and difficult to interpret. This is probably the biggest challenge. –ML Survey*

**Environmental Harm** Environmental harm is a significant concern associated with the development and deployment of complex AI models, as reflected in the data. The intensive energy consumption required for model training and

operation contributes to a large carbon footprint, exacerbating environmental degradation. As AI models grow larger and more complex, these costs—both monetary and ecological—continue to escalate, raising urgent ethical questions about sustainability.

*I have felt that there's been a lack of acknowledgment of the environmental impact of using generative AI models. Not everyone needs to have an AI-generated image for their presentation. –P9*

## 6 RQ3: Role of Human Judgment and Interpretation in Social Science Research

We asked survey participants to envision scenarios where AI/ML complements or collaborates with human researchers in social science research. We followed up with a similar question during interviews.

### 6.1 AI/ML As Collaborators

Many respondents expressed that AI-powered tools are increasingly being viewed as valuable collaborators in social science research, particularly for tasks that involve handling large volumes of data or text. A majority indicated optimism about AI's potential to assist with data coding, statistical interpretation, literature summarization, and generating outlines or code skeletons for research projects. These tools were especially noted for their utility in navigating vast amounts of academic literature and supporting non-native English speakers with writing and proofreading.

*Yes I think that its ability to expedite complicated, lengthy processes like literature reviews or brainstorming paper titles, etc. can allow researchers time to focus more on things like how to write in a way that allows for communication to a wide audience, not just scientific communities. –AI Survey*

### 6.2 Human Judgments Are Irreplaceable

**Critical Thinking and Analytical Reasoning** Participants consistently emphasized that critical thinking and analytical reasoning are fundamental human skills that remain irreplaceable in research. They highlighted that while AI and ML tools can process large amounts of data and identify patterns, they fundamentally (still) lack the capacity for deep reasoning, complex causal inference, and nuanced interpretation. Many noted that research demands not just pattern recognition, but the ability to critically assess evidence, weigh conflicting information and develop sophisticated theoretical insights capabilities that current AI systems cannot replicate.

*Humans will be hardest to replace when normative judgments—about which questions are important and which policy solutions are appropriate—must be made. –AI Survey*

**Creativity, Innovation, and Theory Development** While AI is seen as a potentially useful tool for tasks like measurement, data analysis, and hypothesis generation, participants believe that the core intellectual and theoretical contributions—such as theory building, interpreting political implications, and engaging in value discourse are domains where human insight cannot be replaced.

*Absolutely. Historical and social context are absolutely needed. Groundtruthing needs to be done by people. And theory development and writing can't be automated, even if some ML tools might enhance our ability to do this work effectively and at scale. –ML Survey*

### Ethical Reasoning and Human-Centered Judgment

Participants consistently highlighted the irreplaceable role of human ethical reasoning and human-centered judgment in research. They noted that AI and machine learning systems currently lack the ability to navigate complex moral considerations, social norms, and the ethical dilemmas that often arise in research practice. Many stressed that ethical decision-making such as ensuring participant welfare, mitigating bias, interpreting sensitive contexts, and maintaining integrity requires human intuition, empathy, and moral responsibility. There was a strong sense that delegating ethical reasoning entirely to AI would risk undermining the foundational human values that guide responsible research.

*Areas where human morality and direct interaction with other humans are needed. Unless AIs become extremely sophisticated, there will always be areas where takes humans to know humans. –AI Survey*

### Validation, Verification, and Quality Control

Validation, verification, and quality control of research processes and outcomes have been viewed as the most important steps where human judgments are irreplaceable. Activities such as ground truth, critical review, and making final judgments were seen as areas where human oversight remains indispensable. Many participants expressed deep skepticism about the idea of fully trusting AI systems without human verification, pointing out that even highly advanced AI models can generate errors, overlook context, or reinforce biases.

*I don't have a problem with AI. But we need the validation by humans. And I think if we are doing qualitative research, it becomes even more important to validate by humans. –P6*

### Motivation, Passion, and Human Values

Motivation, passion, and human values are fundamental drivers of meaningful research, qualities that AI and ML cannot replicate. Several responses highlighted that passion fuels perseverance through complex, ambiguous, and challenging research problems, often inspiring creative breakthroughs that AI lacks the capacity to achieve.

*Yes, human creativity, human passion, and judgment and interpretation are key. AI isn't able to be creative and passionate. –AI Survey*

## 7 Discussion and Recommendations

AI and ML are reshaping research practices in the social sciences. While these technologies expand the technological toolkit available to researchers, they also simultaneously introduce new tensions around human judgment, automation, bias, and research integrity. The rapid adoption of generative AI tools, particularly large language models like ChatGPT, has enabled researchers to automate or accelerate tasks such as literature summarization, idea generation, coding support, and exploratory analysis. However, our findings show that this integration is not uniform in social science, while some

researchers view genAI as a productivity tool, others remain cautious due to concerns about accuracy, deskilling, and ethical implications. ML techniques, on the other hand, are perceived as more statistically rigorous and better aligned with established research norms in fields that emphasize prediction and data modeling. Yet, they remain less accessible to researchers without technical training. This divergence reveals a dual dynamic: AI is seen as more usable but less trustworthy; whereas, ML is more trusted but less frequently used. Importantly, social scientists are not merely adopting these technologies, they are critically engaging with them, reflecting on how these tools fit within their disciplinary values of interpretation, transparency, and social accountability.

### 7.1 Human-Centered Integration of AI/ML in Social Science Research

Social scientists' ambivalence toward AI and ML reflects long-standing concerns about the agency of technical systems within knowledge production processes (Star 1999; Suchman 2007). Rather than just functioning as neutral instruments, AI and ML models actively shape what counts as knowledge, whose perspectives are prioritized, and which research questions are considered tractable or valuable. As prior works on algorithmic mediation highlights (Gillespie 2014; Sandvig et al. 2014), technical systems encode values, assumptions, and limitations that become naturalized through repeated use. Our results showcase that these concerns need to be explored and attended to, to ensure that tools' inherent weaknesses and biases are not overlooked.

Our participants' worry about 'theoretical misalignment' and 'de-skilling' point to deep epistemic shifts: *AI tools privilege data-driven, pattern-based reasoning may crowd out theory-driven, critical, and contextual approaches essential to social science inquiry*. This finding resonates with work in critical algorithm studies (Seaver 2017; Boyd 2021), which argues that machine learning systems risk promoting empiricism without explanation, undermining traditional scholarly norms of causal reasoning and reflective critique.

Our results also demonstrate that while social scientists increasingly view AI/ML tools as valuable collaborators—especially for tasks involving data management, statistical analysis, and literature summarization—there is strong consensus that these technologies must remain subordinate to human judgment. This aligns with prior HCI work emphasizing human-in-the-loop approaches (Ehsan et al. 2021), which argue that AI systems should augment rather than automate human decision-making, especially in domains requiring interpretation, ethical reasoning, and creativity. We highlight the persistent emphasis by participants on maintaining critical thinking, ethical oversight, and theoretical engagement.

Thus, adopting AI in social science is not just a technical decision—it is a political and epistemological one that needs to be navigated carefully to preserve the disciplines' interpretive traditions. Stemming from the findings of STS theory, we argue that a potential way to move forward would rather be to focus on giving power back to the community by allowing organizations and technologies to be designed together, recognizing the interdependent relationship

between technical and social systems (Walker et al. 2008; Narayanan Venkit 2023). The tenets of HCXAI (Human-centered Explainable AI) (Ehsan et al. 2021; Ehsan and Riedl 2020; Chakravorti et al. 2023a) can also be used to drive this initiative, putting the relevant community that interacts with the system at the center of technology design.

### 7.2 Standardization, Transparency, and Accountability

Another major theme identified through our work was the *risk of automation bias*—an over-reliance on AI outputs at the expense of double-checking and validation. Participants' concerns reflect broader findings in human-AI interaction research, where trust calibration remains a core challenge (Dietvorst, Simmons, and Massey 2015; Zhang, Liao, and Bellamy 2020). Over time, users tend to defer to technological and algorithmic recommendations even when they conflict with domain knowledge, particularly when systems are black boxes (O'Neil 2017).

In the context of academic research, automation bias threatens not only individual judgment but the collective development of knowledge. As noted by (Green and Chen 2019; Bender et al. 2021), reliance on algorithmic outputs without contestation can institutionalize existing biases, marginalizing alternative epistemologies and reinforcing dominant narratives. Additional concerns around bias, environmental harm, and data exploitation surfaced throughout our study. Prior research has shown that these issues are often treated as externalities to be *managed* after technical deployment (Selbst et al. 2019; Sambasivan et al. 2021). However, our participants' narratives confirm that ethical risks are infrastructural: they are embedded into how AI systems are built, trained, deployed, and normalized.

For instance, concerns about dataset opacity, algorithmic bias, and representational harms mirror findings from (Paullada et al. 2021; Dev et al. 2022; Blodgett et al. 2022), who argue that data provenance and curation practices must be treated as first-order ethical issues. Similarly, worries about carbon costs echo the Green AI movement (Strubell, Ganesh, and McCallum 2020; Schwartz et al. 2020), emphasizing that technical efficiency should be weighed against environmental sustainability.

We also argue that future AI tools for research need to be explicitly designed to foreground uncertainty, invite interrogation, and support counterfactual exploration—design principles that align with calls for contestable AI (Binns 2018; Alfrink et al. 2023; Reckinger and Wille 2018). Embedding friction into AI systems—through transparency prompts, interpretability features, or required human-in-the-loop validation—can help sustain critical rather than erode it.

Finally, the absence of best practices for documenting and evaluating AI/ML-assisted research emerged as another barrier to trust among social science researchers. This adds to the concerns about the reproducibility crisis in computational fields (Sandve et al. 2013; Chakravorti, Koneru, and Rajtmajer 2025; Wu et al. 2024a; Chakravorti et al. 2023b) and points to the need for standardized methodologies. Drawing from frameworks proposed in explainable AI

(XAI) (Miller 2019) and HCI for ML systems (Amershi et al. 2019), we argue that social science research domains need to develop AI documentation protocols that include: [1] model transparency (e.g., architecture, training data, limitations); [2] usage disclosure (e.g., which parts of the analysis involved AI assistance); [3] evaluation criteria (e.g., audits for fairness, bias, epistemic validity); and [4] reproducibility artifacts (e.g., code, data, prompt logs).

### 7.3 Recommendations: A Way Towards Responsible Integration of AI in the Social Sciences

Taking inspiration from the findings provided by our participants, as well as motivated by prior technological audits and recommendation frameworks (Blodgett et al. 2022; Bender and Friedman 2018; Venkit et al. 2024a), we provide *research-driven recommendations* for different stakeholder groups in this social science space involved in the integration of AI and ML into social science research ecosystems.

#### For AI Developers and Tool Designers

- **Prioritize Explainability and Transparency:** Incorporate interpretable models, visual explanations, and clear documentation of training data and model behavior.
- **Incorporate Lived Experiences Into the Design Process:** Involve end users, especially social science researchers—in ongoing, real-world testing and iterative feedback loops to ensure AI tools reflect actual needs, values, and research practices.
- **Enable Critical Interventions:** Provide users with editable outputs, override mechanisms, and configurable prompt logs for agency and accountability.
- **Embed Ethical Safeguards:** Integrate prompts, alerts, or design constraints that flag potentially biased, unsafe, or misleading outputs, if possible.
- **Support Sustainability:** Disclose environmental impacts (e.g., carbon emissions) and develop lightweight alternatives where possible.

#### For Researchers and Social Science Practitioners

- **Develop Critical AI Literacy:** Engage with foundational technical knowledge and ethical frameworks to understand the capabilities and limitations of AI/ML systems.
- **Adopt Human-AI Collaboration Mindsets:** Use AI as an assistant rather than a substitute; maintain human judgment in interpretation, validation, and theorizing.
- **Engage in Participatory, Lived Experimentation:** Treat the implementation of AI/ML tools not as static adoption but as dynamic, co-constructed processes, using lived experiments to surface tacit knowledge, contextual concerns, and ethical challenges.
- **Document AI Usage Transparently:** Clearly describe when, where, and how AI tools are used in research to enhance reproducibility and trust.
- **Maintain Ethical Oversight:** Continuously assess representational harms, dataset bias, and implications for marginalized populations in research contexts.

#### For Educators and Academic Institutions

- **Integrate Critical AI Pedagogy:** Incorporate algorithmic ethics, bias analysis, and sociotechnical systems thinking into curricula for students in both technical and non-technical disciplines.
- **Support Interdisciplinary Skill-Building:** Offer workshops, boot camps, and joint seminars that bring together social scientists, data scientists, and designers.
- **Revise Research Training Guidelines:** Update review board (eg: IRB) protocols and academic writing guidelines to include AI-specific considerations.

#### For Policymakers and Funding Bodies

- **Support Development of Ethical Standards:** Fund the creation of domain-specific AI governance guidelines, audit frameworks, and participatory design protocols.
- **Mandate Sustainability Reporting:** Require AI research proposals to include environmental impact assessments and encourage the adoption of carbon-conscious modeling practices.
- **Promote Equitable Access:** Ensure equitable distribution of AI infrastructure, funding, and training opportunities across institutions and researcher demographics.
- **Encourage Community-Led Governance:** Involve diverse stakeholders, including marginalized communities, in decisions about AI deployment and oversight in research.

## 8 Conclusion

Our work highlights the ways in which the integration of AI and ML into social science research marks a significant sociotechnical shift raising important ethical and methodological questions. Through a combination of surveys followed by semi-structured interviews, we found that AI especially in its generative AI form, is increasingly accessible and attractive to everyone. Its rapid adoption raises serious concerns. These include over-reliance, erosion of theoretical grounding, risks of bias, threats to research integrity, transparency, and environmental impacts. Meanwhile, machine learning, although more rigorously aligned with statistical reasoning, remains less accessible to practitioners due to complex methodologies. This divergence in accessibility and perception underscores the need for targeted capacity building, clearer standards, and ongoing ethical reflection. Critically, we have argued for the importance of human-centered designs—not just a matter of interface design but a deeper commitment to preserving the agency, expertise, and interpretive traditions of social science. Ethical engagement, methodological transparency, and inclusive design must be core to how AI and ML systems are built and adopted. This means moving beyond technological determinism toward a more participatory, reflective, and justice-oriented framework of research practice. In response, we offer actionable recommendations for developers, researchers, educators, and policymakers focused on explainability, ethics, sustainability, and equitable access.

## References

- Alfrink, K.; Keller, I.; Kortuem, G.; and Doorn, N. 2023. Contestable AI by design: Towards a framework. *Minds and Machines*, 33(4): 613–639.
- Amershi, S.; Weld, D.; Vorvoreanu, M.; Fourney, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P. N.; Inkpen, K.; et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*, 1–13.
- Anthis, J. R.; Liu, R.; Richardson, S. M.; Kozlowski, A. C.; Koch, B.; Evans, J.; Brynjolfsson, E.; and Bernstein, M. 2025. LLM Social Simulations Are a Promising Research Method. *arXiv preprint arXiv:2504.02234*.
- Athey, S.; and Imbens, G. W. 2017. The state of applied econometrics: Causality and policy evaluation. *Journal of Economic perspectives*, 31(2): 3–32.
- Bail, C. 2023. Can generative artificial intelligence improve social science. *SocArXiv*.
- Bail, C. A. 2024. Can Generative AI improve social science? *Proceedings of the National Academy of Sciences*, 121(21): e2314021121.
- Baker, M. 2016. Reproducibility crisis. *nature*, 533(26): 353–66.
- Bender, E. M.; and Friedman, B. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6: 587–604.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610–623.
- Binns, R. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency*, 149–159. PMLR.
- Blandford, A.; Furniss, D.; and Makri, S. 2016. *Qualitative HCI research: Going behind the scenes*. Morgan & Claypool Publishers.
- Blodgett, S. L.; Liao, Q. V.; Olteanu, A.; Mihalcea, R.; Muller, M.; Scheuerman, M. K.; Tan, C.; and Yang, Q. 2022. Responsible language technologies: Foreseeing and mitigating harms. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–3.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29.
- Boyd, D. 2021. Data, Algorithms, Fairness, Accountability. *TRT Akademi*, 6(11): 202–209.
- Castelvecchi, D. 2016. Can we open the black box of AI? *Nature News*, 538(7623): 20.
- Celik, I.; Dindar, M.; Muukkonen, H.; and Järvelä, S. 2022. The promises and challenges of artificial intelligence for teachers: A systematic review of research. *TechTrends*, 66(4): 616–630.
- Chakravorti, T.; Fraleigh, R.; Fritton, T.; McLaughlin, M.; Singh, V.; Griffin, C.; Kwasnica, A.; Pennock, D.; Giles, C. L.; and Rajtmajer, S. 2023a. A prototype hybrid prediction market for estimating replicability of published work. In *HHAI 2023: Augmenting Human Intellect*, 300–309. IOS Press.
- Chakravorti, T.; Koneru, S.; and Rajtmajer, S. 2025. Reproducibility and replicability in research: What 452 professors think in Universities across the USA and India. *PloS one*, 20(3): e0319334.
- Chakravorti, T.; Wu, C.; Koneru, S.; and Rajtmajer, S. 2023b. Perspectives from India: Opportunities and Challenges for AI Replication Prediction to Improve Confidence in Published Research. *arXiv preprint arXiv:2310.19158*.
- Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; and Robins, J. 2018. Double/debiased machine learning for treatment and structural parameters.
- Chubb, J.; Cowling, P.; and Reed, D. 2022. Speeding up to keep up: exploring the use of AI in the research process. *AI & society*, 37(4): 1439–1457.
- Dandamudi, S. R. P.; Sajja, J.; and Khanna, A. 2025. Advancing Cybersecurity and Data Networking Through Machine Learning-Driven Prediction Models. *International Journal of Innovative Research in Computer Science and Technology*, 13(1): 26–33.
- Das, D.; Guha, S.; Brubaker, J.; and Semaan, B. 2024. The “Colonial Impulse” of Natural Language Processing: An Audit of Bengali Sentiment Analysis Tools and Their Identity-based Biases. *arXiv preprint arXiv:2401.10535*.
- Deranty, J.-P.; and Corbin, T. 2024. Artificial intelligence and work: a critical review of recent research from the social sciences. *Ai & Society*, 39(2): 675–691.
- Dev, S.; Sheng, E.; Zhao, J.; Amstutz, A.; Sun, J.; Hou, Y.; Sanseverino, M.; Kim, J.; Nishi, A.; Peng, N.; et al. 2022. On Measures of Biases and Harms in NLP. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, 246–267.
- Dietvorst, B. J.; Simmons, J. P.; and Massey, C. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of experimental psychology: General*, 144(1): 114.
- Ding, X.; Kou, Y.; Xu, Y.; and Zhang, P. 2022. “As Uploaders, We Have the Responsibility”: Individualized Professionalization of Bilibili Uploaders. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–14.
- Dolata, M.; Feuerriegel, S.; and Schwabe, G. 2022. A sociotechnical view of algorithmic fairness. *Information Systems Journal*, 32(4): 754–818.
- Dominguez-Olmedo, R.; Hardt, M.; and Mendler-Dünner, C. 2024. Questioning the survey responses of large language models. *Advances in Neural Information Processing Systems*, 37: 45850–45878.
- Edelmann, A.; Wolff, T.; Montagne, D.; and Bail, C. A. 2020. Computational social science and sociology. *Annual review of sociology*, 46(1): 61–81.

- Ehsan, U.; Liao, Q. V.; Muller, M.; Riedl, M. O.; and Weisz, J. D. 2021. Expanding explainability: Towards social transparency in ai systems. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, 1–19.
- Ehsan, U.; and Riedl, M. O. 2020. Human-centered explainable ai: Towards a reflective sociotechnical approach. In *HCI International 2020-Late Breaking Papers: Multimodality and Intelligence: 22nd HCI International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings 22*, 449–466. Springer.
- Einav, L.; and Levin, J. 2014. Economics in the age of big data. *Science*, 346(6210): 1243089.
- Fan, C.; Xu, J.; Natarajan, B. Y.; and Mostafavi, A. 2023. Interpretable machine learning learns complex interactions of urban features to understand socio-economic inequality. *Computer-Aided Civil and Infrastructure Engineering*, 38(14): 2013–2029.
- Gautam, S.; Venkit, P. N.; and Ghosh, S. 2024. From melting pots to misrepresentations: Exploring harms in generative ai. *arXiv preprint arXiv:2403.10776*.
- Ghosh, S.; Venkit, P. N.; Gautam, S.; Wilson, S.; and Caliskan, A. 2024. Do Generative AI Models Output Harm while Representing Non-Western Cultures: Evidence from A Community-Centered Approach. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 476–489.
- Gillespie, T. 2014. The relevance of algorithms. *Media technologies: Essays on communication, materiality, and society*, 167(2014): 167.
- Green, B.; and Chen, Y. 2019. Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the conference on fairness, accountability, and transparency*, 90–99.
- Grimmer, J. 2015. We are all social scientists now: How big data, machine learning, and causal inference work together. *PS: Political Science & Politics*, 48(1): 80–83.
- Grimmer, J.; Roberts, M. E.; and Stewart, B. M. 2021. Machine learning for social science: An agnostic approach. *Annual Review of Political Science*, 24(1): 395–419.
- Grossmann, I.; Feinberg, M.; Parker, D. C.; Christakis, N. A.; Tetlock, P. E.; and Cunningham, W. A. 2023. AI and the transformation of social science research. *Science*, 380(6650): 1108–1109.
- Hacker, P. 2021. A legal framework for AI training data—from first principles to the Artificial Intelligence Act. *Law, innovation and technology*, 13(2): 257–301.
- Hofman, J. M.; Watts, D. J.; Athey, S.; Garip, F.; Griffiths, T. L.; Kleinberg, J.; Margetts, H.; Mullainathan, S.; Salganik, M. J.; Vazire, S.; et al. 2021. Integrating explanation and prediction in computational social science. *Nature*, 595(7866): 181–188.
- Huang, X.; Vitak, J.; and Tausczik, Y. 2020. "You Don't Have To Know My Past": How WeChat Moments Users Manage Their Evolving Self-Presentation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing systems*, 1–13.
- Iorga, D.; Bocse, B.; and Olar, D. 2023. SOCIOLOGICAL ARTIFICIAL INTELLIGENCE RESEARCH (SAIR). A REVIEW AND CONCEPTUAL FRAMEWORK. *Revista Romana de Sociologie*.
- Ivankova, N. V.; Creswell, J. W.; and Stick, S. L. 2006. Using mixed-methods sequential explanatory design: From theory to practice. *Field methods*, 18(1): 3–20.
- Jansen, B. J.; Jung, S.-g.; and Salminen, J. 2023. Employing large language models in survey research. *Natural Language Processing Journal*, 4: 100020.
- Kaur, P.; Stoltzfus, J.; and Yellapu, V. 2018. Descriptive statistics. *International Journal of Academic Medicine*, 4(1): 60–63.
- Ligo, A. K.; Rand, K.; Bassett, J.; Galaitsi, S. E.; Trump, B. D.; Jayabalasingham, B.; Collins, T.; and Linkov, I. 2021. Comparing the emergence of technical and social sciences research in artificial intelligence. *Frontiers in Computer Science*, 3: 653235.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.
- Miller, T. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267: 1–38.
- Narayanan Venkit, P. 2023. Towards a holistic approach: Understanding sociodemographic biases in nlp models using an interdisciplinary lens. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 1004–1005.
- Narayanan Venkit, P.; Gautam, S.; Panchanadikar, R.; Huang, T.-H.; and Wilson, S. 2023. Unmasking nationality bias: A study of human perception of nationalities in ai-generated articles. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 554–565.
- Ntoutsis, E.; Fafalios, P.; Gadiraju, U.; Iosifidis, V.; Nejdil, W.; Vidal, M.-E.; Ruggieri, S.; Turini, F.; Papadopoulos, S.; Krasanakis, E.; et al. 2020. Bias in data-driven artificial intelligence systems—An introductory survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(3): e1356.
- Obreja, D. M.; Rughinis, R.; and Rosner, D. 2024. Mapping the conceptual structure of innovation in artificial intelligence research: A bibliometric analysis and systematic literature review. *Journal of Innovation & Knowledge*, 9(1): 100465.
- Omran Zailuddin, M. F. N.; Nik Harun, N. A.; Abdul Rahim, H. A.; Kamaruzaman, A. F.; Berahim, M. H.; Harun, M. H.; and Ibrahim, Y. 2024. Redefining creative education: a case study analysis of AI in design courses. *Journal of Research in Innovative Teaching & Learning*, 17(2): 282–296.
- O'Neil, C. 2017. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- Paullada, A.; Raji, I. D.; Bender, E. M.; Denton, E.; and Hanna, A. 2021. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11).

- Reckinger, R.; and Wille, C. 2018. Situative interdisciplinarity: empirical reflections on ten years of cross-disciplinary research. *Academic Journal of Interdisciplinary Studies*, 7(3).
- Rosen, L. D.; Whaling, K.; Carrier, L. M.; Cheever, N. A.; and Rokkum, J. 2013. The media and technology usage and attitudes scale: An empirical investigation. *Computers in human behavior*, 29(6): 2501–2511.
- Sadiku, M. N.; Ashaolu, T. J.; Ajayi-Majebi, A.; and Musa, S. M. 2021. Artificial intelligence in social media. *International Journal of Scientific Advances*, 2(1): 15–20.
- Salganik, M. J. 2019. *Bit by bit: Social research in the digital age*. Princeton University Press.
- Sambasivan, N.; Kapania, S.; Highfill, H.; Akrong, D.; Paritosh, P.; and Aroyo, L. M. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15.
- Sandve, G. K.; Nekrutenko, A.; Taylor, J.; and Hovig, E. 2013. Ten simple rules for reproducible computational research. *PLoS computational biology*, 9(10): e1003285.
- Sandvig, C.; Hamilton, K.; Karahalios, K.; and Langbort, C. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry*, 22(2014): 4349–4357.
- Schwartz, R.; Dodge, J.; Smith, N. A.; and Etzioni, O. 2020. Green ai. *Communications of the ACM*, 63(12): 54–63.
- Seaver, N. 2017. Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big data & society*, 4(2): 2053951717738104.
- Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, 59–68.
- Shmueli, G. 2010. To explain or to predict? *Statistical science*, 289–310.
- Star, S. L. 1999. The ethnography of infrastructure. *American behavioral scientist*, 43(3): 377–391.
- Strubell, E.; Ganesh, A.; and McCallum, A. 2020. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 13693–13696.
- Suchman, L. A. 2007. *Human-machine reconfigurations: Plans and situated actions*. Cambridge university press.
- Thakkar, D.; Ismail, A.; Kumar, P.; Hanna, A.; Sambasivan, N.; and Kumar, N. 2022. When is machine learning data good?: Valuing in public health datafication. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–16.
- Van Noorden, R.; and Perkel, J. M. 2023. AI and science: what 1,600 researchers think. *Nature*, 621(7980): 672–675.
- Varanasi, R. A.; Siddarth, D.; Seshadri, V.; Bali, K.; and Vashistha, A. 2022. Feeling Proud, Feeling Embarrassed: Experiences of Low-income Women with Crowd Work. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Varian, H. R. 2014. Big data: New tricks for econometrics. *Journal of economic perspectives*, 28(2): 3–28.
- Venkit, P. N.; Chakravorti, T.; Gupta, V.; Biggs, H.; Srinath, M.; Goswami, K.; Rajtmajer, S.; and Wilson, S. 2024a. “Confidently Nonsensical?”: A Critical Survey on the Perspectives and Challenges of ‘Hallucinations’ in NLP. *arXiv preprint arXiv:2404.07461*.
- Venkit, P. N.; Chakravorti, T.; Gupta, V.; Biggs, H.; Srinath, M.; Goswami, K.; Rajtmajer, S.; and Wilson, S. 2024b. An Audit on the Perspectives and Challenges of Hallucinations in NLP. *arXiv preprint arXiv:2404.07461*.
- Von Eschenbach, W. J. 2021. Transparency and the black box problem: Why we do not trust AI. *Philosophy & Technology*, 34(4): 1607–1622.
- Walker, G. H.; Stanton, N. A.; Salmon, P. M.; and Jenkins, D. P. 2008. A review of sociotechnical systems theory: a classic concept for new command and control paradigms. *Theoretical issues in ergonomics science*, 9(6): 479–499.
- Williams, R. T. 2024. Paradigm shifts: exploring AI’s influence on qualitative inquiry and analysis. *Frontiers in Research Metrics and Analytics*, 9: 1331589.
- Wu, C.; Chakravorti, T.; Carroll, J. M.; and Rajtmajer, S. 2024a. Integrating measures of replicability into scholarly search: Challenges and opportunities. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Wu, C.; Wang, X.; Carroll, J.; and Rajtmajer, S. 2024b. Reacting to generative AI: Insights from student and faculty discussions on Reddit. In *Proceedings of the 16th ACM Web Science Conference*, 103–113.
- Wulff, P.; Buschhüter, D.; Westphal, A.; Mientus, L.; Nowak, A.; and Borowski, A. 2022. Bridging the gap between qualitative and quantitative assessment in science education research with machine learning—A case for pre-trained language models-based clustering. *Journal of Science Education and Technology*, 31(4): 490–513.
- Xu, Y.; Liu, X.; Cao, X.; Huang, C.; Liu, E.; Qian, S.; Liu, X.; Wu, Y.; Dong, F.; Qiu, C.-W.; et al. 2021. Artificial intelligence: A powerful paradigm for scientific research. *The Innovation*, 2(4).
- Zhang, Y. 2023. Generative AI has lowered the barriers to computational social sciences. *arXiv preprint arXiv:2311.10833*.
- Zhang, Y.; Liao, Q. V.; and Bellamy, R. K. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 295–305.
- Ziems, C.; Held, W.; Shaikh, O.; Chen, J.; Zhang, Z.; and Yang, D. 2024. Can large language models transform computational social science? *Computational Linguistics*, 50(1): 237–291.