Prevalence and content of AI guidelines across social science journals and their publishers

Ben Lasse Wolf
German Centre for Higher Education Research and Science Studies (DZHW)
Leibniz University Hannover

Joshua Claassen
German Centre for Higher Education Research and Science Studies (DZHW)
Leibniz University Hannover

Jan Karem Höhne
German Centre for Higher Education Research and Science Studies (DZHW)
Leibniz University Hannover

Abstract

The increasing use of artificial intelligence (AI) in academic research and publishing raises questions about how standards of scientific integrity and research ethics can be adapted to AI-assisted research. Journals and publishers have developed AI guidelines, yet their prevalence and content remain unclear. We investigate the research questions: *How prevalent are AI guidelines across social science journals and their publishers (RQ1), and what is their content (RQ2)?* We examine the five highest-ranked journals of all 41 categories in the "Social Science, General" group in Clarivate's Web of Science database, resulting in a sample of 205 journals and 29 publishers. We code their AI guidelines for general AI provisions and specific AI use cases. AI guidance is common, but unevenly distributed. While 71% of journals provide AI guidelines, only 55% of publishers do so. Content analysis shows consensus on some general AI provisions, particularly regarding human authorship and accountability. However, guidance on specific AI use cases is more heterogeneous. While writing-related AI use cases are frequently addressed and permitted, use cases beyond writing, such as brainstorming and summarization, remain largely unaddressed. Such omissions do not necessarily prohibit these uses but leave their permissibility and disclosure unspecified. Our findings reveal a gap between increasingly well-defined general AI provisions and concrete standards for specific AI use cases. This uncertainty shifts responsibility for interpreting and applying AI guidance to individual authors and highlights the importance of involving the scientific community in developing AI guidelines while considering the scientific and commercial interests shaping such standards.

Keywords: Artificial Intelligence, Scientific integrity, Academic publishing, Research ethics, Author guidelines, Large Language Models

This document is a preprint and thus it may differ from the final version.

Introduction and background

Since the launch of OpenAI's ChatGPT in November 2022, the rapid development of artificial intelligence (AI) technologies, especially Large Language Models (LLMs)[1], has not only shaped society but also academia and the way academics conduct research. There is empirical evidence indicating a fast and growing adoption of AI by academics, with the extent of AI use varying between and within disciplines (Liang et al., 2025). For example, Cao (2026) demonstrates how LLM-assisted academic writing has diffused through sociology and how it affects contemporary sociological writing. Analyzing almost 1.5 million sociology abstracts, the author finds that the proportion of LLM-modified sentences reached almost 25% by mid-2025. Importantly, the adoption of LLMs is particularly concentrated in empirical and applied subfields, among early-career scholars, and in lower-tier journals. These findings indicate that the use of LLM-assisted writing is a socially differentiated practice that appears to be associated with aspects of academic socialization, raising questions with respect to equal opportunities. Beyond academic writing, AI is also becoming increasingly involved in research activities that range from literature search (Walters & Wilder, 2023) to peer review (Goyanes et al., 2025).

AI is capable of supporting researchers both during routine tasks (e.g., analysis code writing; Nejjar et al., 2025; Sebastian et al., 2025) and demanding, theory-driven tasks (e.g., hypothesis generation; Zhou et al., 2024), possibly enhancing productivity and reducing research effort (Goyanes et al., 2025). Even though writing assistance, including grammar correction, language editing, and paraphrasing manuscript sections remains one of the most common applications (Alvero et al., 2026), AI is also increasingly integrated into empirical research processes. During the conceptualization phase, for example, LLMs have been investigated as tools for question and questionnaire pretesting (Lenzner & Hadler, 2026). Within the data collection phase, conversational LLM-based chatbots have been explored for conducting dialogue-driven interviews (Wuttke et al., 2025). During data processing and analysis, Automatic Speech Recognition (ASR) systems have been employed for the automated transcription of audio data (Höhne et al., 2025; Revilla et al., 2026). AI is also increasingly used for text annotation and classification (Claassen et al., 2026; Gilardi et al., 2023; Lin & Zhang, 2026; Pangakis & Wolken, 2024), for information extraction from text or images (Law & Roberto, 2025; Schwitter, 2025; Stuhler et al., 2025), or for imputation of missing data (Holtdirk et al., 2025; Kim & Lee, 2023). Furthermore, AI can be utilized to analyze qualitative data through interactive, dialogue-like processes (e.g., using LLMs to ask targeted questions about recurring patterns in the data or potential interpretations of the material; Hayes, 2025). Recently, researchers also started to generate synthetic survey data through so-called silicon sampling approaches (Argyle et al., 2023) to, for example, predict election outcomes (von der Heyde et al., 2026).

Beyond empirical research, AI has also increasingly been proposed for scientometric and bibliometric tasks. For example, AI tools have been used to predict citation counts, Mendeley readership, and social media engagement related to scientific publications (De Winter, 2024). Relatedly, Pereira et al. (2025) have developed a Python library that integrates AI to combine traditional bibliometric analyses (e.g., identification of research trends and essential researchers) with topic modeling and text summarization. These developments demonstrate that

[1] In the social sciences, the term "AI" often refers to LLMs. In most cases where we refer to AI, we thus are implicitly referring to LLMs. We address this point further in the discussion section.

AI is no longer confined to isolated research tasks but is increasingly integrated into the entire research life cycle.

At the same time, the growing integration of AI raises important questions regarding scientific integrity and knowledge production. While AI promises substantial gains in efficiency and productivity, it also influences how scientific knowledge is produced, curated, and transferred. In this regard, Lindemann (2025) argues that AI, particularly in the form of LLMs, differs fundamentally from prior technologies. Drawing on the concepts of "situated knowledges" (Haraway, 1988) and "sealed surfaces" (Mühlhoff, 2018), the author introduces the concept of "sealed knowledges." More specifically, Lindemann (2025) argues that AI in the form of LLMs tends to compress vast and diverse information into single, plausible outputs concealing its initial complexity and context-dependency. A particular problem is also that AI developers commonly do not disclose the underlying algorithmic processes (also referred to as the "black box" principle; von Eschenbach, 2021), decreasing the traceability and replicability of the outcomes. Importantly, this contradicts basic scientific standards with respect to transparency.

Relatedly, various studies identified challenges associated with the use of AI in research, including homogenization of scientific writing (e.g., AI-generated text may reduce linguistic variation and creativity; Kuteeva & Andersson, 2024), biased or scientifically flawed outputs (e.g., AI-generated summaries may omit limitations and overgeneralize findings; Peters & Chin-Yee, 2025), data privacy concerns (e.g., risks of exposing personally identifiable or other sensitive research data; Ganguly et al., 2025), and hallucinated citations (e.g., AI-generated references that do not correspond to actual scholarly works; Walters & Wilder, 2023). For example, based on an analysis of about 2.5 million scientific articles, Naddaf (2026) reported that the rates of fabricated references have climbed steeply since 2023 (after the launch of OpenAI's ChatGPT). This trend indicates that unmonitored and undisclosed AI use has the potential to negatively affect the integrity of knowledge production, curation, and archiving.

While the responsible use of AI in research has become an increasingly important topic, existing standards on AI use are still in their infancy and considerable uncertainty persists regarding permitted and prohibited forms of AI use throughout the research life cycle. This dynamic is reflected in evolving recommendations by research and funding organizations, as well as in the growing number of AI guidelines by scientific journals and publishers (Ganjavi et al., 2024; Lendvai & Petra, 2026). In the following, we use *AI guidance* to refer to the broader normative orientation regarding permissible and responsible AI use in academic research and publishing, and *AI guidelines* to refer to the formal documents through which such guidance can be articulated.

Several scientific organizations have published overarching recommendations for the responsible use of AI. For example, the Committee on Publication Ethics (COPE) and the International Committee of Medical Journal Editors (ICMJE) have issued position statements emphasizing transparency, accountability, and human responsibility when AI is used in research (COPE Council, 2023; International Committee of Medical Journal Editors, 2025). Additionally, the American Association for Public Opinion Research (AAPOR) recently published a report on the responsible integration of AI in survey research, including recommendations for transparency and reporting (Rothschild et al., 2026). While these recommendations provide important normative guidance across scientific disciplines, they do

not necessarily capture the concrete expectations when it comes to the publication through scientific journals and their respective publishers (Chigwada & Ngulube, 2026). Existing guidelines from research and funding organizations also remain limited in their guidance on AI use throughout the research process. For example, the German Research Foundation (DFG) has issued AI guidelines concerning its funding and assessment procedures, including requirements regarding transparency, confidentiality, quality assurance, and human responsibility (Deutsche Forschungsgemeinschaft, 2023, 2026). These guidelines are of high value and point out important aspects, but they were never developed and issued for general research processes.

At the same time, journals and publishers have started to provide AI guidelines to better manage the existing uncertainty (Goyanes et al., 2025). Yet, these guidelines may not only reflect scientific considerations but also organizational and commercial interests, resulting in a potential conflict of interest. The dominant position of the 'Big 5' academic publishers[2], which control a substantial share of the academic publishing market, further underscores the relevance of this issue (Butler et al., 2023). Their influential role in shaping editorial and authorship guidelines also raises the question of how such stakeholders define permissible AI use (Lendvai & Petra, 2026). Although academic editors may be involved in the formulation of journal-specific AI guidelines, such processes often lack transparency, preventing the broader scientific community from participating in the development of these guidelines.

Chigwada and Ngulube (2026) show that although many publishers uniformly prohibit AI from being listed as author and require disclosure of AI use, existing guidelines appear to differ considerably regarding their scope, specificity, and disclosure requirements. Similarly, Lendvai and Petra (2026) identify a broad consensus regarding authorship and disclosure requirements across the 'Big 5' academic publishers, while highlighting substantial variation concerning permissible AI use, peer review, editorial responsibilities, and the normative framing of AI guidelines. Focusing specifically on social science journals in the broader sense, Goyanes et al. (2025) report that explicit references to AI and its use remain relatively uncommon in the context of journal-specific guidelines. According to the authors, AI guidelines primarily concern writing and transparency requirements (if they exist at all). AI in the context of literature reviews, methodological applications, and peer review receive only little attention. However, none of the AI guidelines analyzed require authors to have specific training in the responsible use of AI that would ensure safe and sound usage. Additionally, there are considerable differences across disciplines, suggesting that expectations regarding responsible AI use remain inconsistent throughout the social sciences (e.g., communication science and sociology appear to be more open to AI use, whereas economics and ethics appear to be less open).

Even though these studies provide valuable insights into AI guidelines within academic publishing, important knowledge gaps remain. Existing research has primarily focused on selected journals or publishers (e.g., AI guidelines of the 'Big 5' academic publishers; Lendvai & Petra, 2026), limited disciplinary samples (e.g., high-ranked journals across a limited number of social science disciplines; Goyanes et al., 2025), or rather abstract dimensions of AI guidelines (e.g., focus on AI authorship, general disclosure requirements, and author

[2] The term 'Big 5' refers to Elsevier, SAGE, Springer Nature, Taylor & Francis, and Wiley, identified as the five largest academic publishers by the number of active, peer-reviewed scholarly journals they publish (Butler et al., 2023; Kim, 2025).

accountability; Ganjavi et al., 2024). To advance the existing literature, it is key to systematically examine both the prevalence and content of AI guidelines across a broad range of social science journals and disciplines. Beyond a high-level focus on AI authorship, general disclosure requirements, and author accountability, this should also include the specific AI use cases addressed throughout different stages of the research process and their respective permission and disclosure requirements. In other words, as authors submitting to these journals potentially struggle with uncertainty regarding permitted and prohibited forms of AI use, it is of utmost importance to shed light on the specific regulations related to AI by conducting an in-depth analysis of the content of existing AI guidelines. In this study, we therefore address the following two research questions:

> ***RQ1:*** *How prevalent are AI guidelines across social science journals and their publishers?*
> ***RQ2:*** *What is the content of AI guidelines issued by social science journals and their publishers?*

To address these two research questions, we systematically examine AI guidelines of the leading social science journals indexed in the Journal Citation Reports (JCR) that is part of Clarivate's Web of Science database, as well as the AI guidelines of their respective publishers. This approach has the great advantage that we shed light on the differences between journal- and publisher-specific AI guidelines, including their potential impact on research practices. Overall, our study makes a three-fold contribution to the existing literature. First, it provides an overview of how leading social science journals and their publishers regulate AI use. By formulating and enforcing guidelines, journals and publishers play an important role in establishing and institutionalizing norms regarding permissible AI use in academic publishing. Examining these guidelines therefore provides insights not only into existing regulatory practices but also into how norms of responsible AI use are formalized and communicated within the scientific community. This contributes to the literature on AI guidance in academic publishing and informs about the state-of-the-art when it comes to responsible AI use. Second, as AI guidelines continue to evolve rapidly, we provide a timely update to the previously discussed research extending contemporary knowledge on publishing in the social sciences. Finally, we expand on the findings of existing studies by applying more detailed and fine-grained coding, shedding further light on specific AI use cases in the context of social science research, such as text production, brainstorming, and literature search.

Method

Journal selection

To create a sample of leading social science journals, we built on the so-called Journal Citation Reports (JCR; see https://jcr.clarivate.com/) that are part of Clarivate's Web of Science database. The Web of Science is a curated database for scientific publications, and the JCR provides annual, citation-based metrics on scientific journals, including the Journal Impact Factor (JIF). For a given year, the JIF measures the average number of citations received by journal articles (including reviews) published in the preceding two years. This metric is widely

used in the scientific community to assess, rank, and compare the impact of journals. For a comprehensive discussion of the JIF, we refer interested readers to Garfield (2006). Importantly, in this study, we selected journals based on the 2025 JIF that was newly released in June 2026.

In the JCR, journals are organized into 254 categories, which are further structured into 21 overarching categorical groups, such as "Social Sciences, General" and "Arts & Humanities, Interdisciplinary." A journal can be included in multiple categories, and a category can be included in multiple groups. For example, the journal "Sociological Methods & Research" is part of two categories: "Sociology" as well as "Social Sciences, Mathematical Methods." The latter category is part of three groups: "Mathematics," "Multidisciplinary," and "Social Sciences, General."

In this study, we focused on the "Social Sciences, General" group, which includes 41 categories, such as "Sociology," "Demography," and "Law." It is the largest group in terms of the number of categories included, except for "Clinical Medicine" (as of July 15, 2026). Overall, the "Social Sciences, General" group consists of 6,804 journals. Table 1 displays all 41 categories, including the number of journals and citable items (or articles). Considering this large number of journals, a complete survey was not feasible in the context of this article due to budget and time constraints. We therefore decided to select the top five journals (according to their 2025 JIF) in each category, and we further address this point in the discussion and conclusion section. This selection strategy resulted in a total of 205 journals, including 180 unique journals (as some of them are listed in more than one category).

Data collection

For data collection, we used the export function integrated into the Web of Science database to gather various types of information directly from the JCR. Specifically, we gathered the following information for each journal: its category, publisher, ISSN (International Standard Serial Number), and 2025 JIF. In addition, we obtained the five-year JIF from the JCR, which is based on the average number of citations received by journal articles published within the past five years.

The first author manually collected supplementary information from the websites of the journals, including the URL (Uniform Resource Locator) of their respective submission guidelines and the year and volume of their first publication.[3] During this manual assessment, the first author also determined the presence of journal-specific AI guidelines, which was binary coded (yes = existence of journal-specific guidelines; no = nonexistence of journal-specific guidelines), and copied any applicable AI guidelines verbatim.

A three-step search approach was used to manually determine whether and to what extent a journal provides AI guidelines. First, the submission and publishing sections of the journal's website were screened for journal-specific author instructions and/or submission guidelines. As most journals provide such information for various aspects of preparing, formatting, and submitting a manuscript, these sections were considered the central location for information on AI use in the search process. Second, the journal's website was searched for references to, or specific subpages detailing, the journal's editorial policies. This was done because many journals refer to editorial policies in their author instructions and submission guidelines,

[3] The year and volume of the journal's first publication referred to the first publication listed on the journal's website or in the JSTOR database (if no information was given on the journal's website).

Table 1. Overview of all 41 JCR categories within the “Social Sciences, General” group

#	JCR category	Journals	Citable items	#	JCR category	Journals	Citable items
1	Anthropology	143	4,735	22	Language & Linguistics	408	11,910
2	Archaeology	171	5,198	23	Law	443	12,575
3	Area Studies	184	5,337	24	Linguistics	312	11,619
4	Communication	227	10,012	25	Medical Ethics	22	1,177
5	Criminology & Penology	113	5,161	26	Medicine, Legal	24	2,384
6	Demography	51	2,354	27	Nursing	194	18,902
7	Economics	627	3,7591	28	Political Science	332	13,865
8	Education & Educational Research	762	43,083	29	Public Administration	97	4,355
9	Education, Scientific Disciplines	90	9,957	30	Regional & Urban Planning	58	5,739
10	Education, Special	63	2,425	31	Social Issues	72	3,491
11	Environmental Studies	202	33,675	32	Social Sciences, Biomedical	48	5,109
12	Ethics	78	3,592	33	Social Sciences, Interdisciplinary	276	20,334
13	Ethnic Studies	40	2,016	34	Social Sciences, Mathematical Methods	70	3,622
14	Geography	174	9,533	35	Social Work	99	5,725
15	Gerontology	49	6,119	36	Sociology	222	10,493
16	Health Care Sciences & Services	194	30,539	37	Sport Sciences	137	15,439
17	Health Policy & Services	125	14,777	38	Substance Abuse	61	5,391
18	History of Social Sciences	45	1,548	39	Transportation	66	7,483
19	Hospitality, Leisure, Sport & Tourism	141	8,776	40	Urban Studies	86	6,898
20	Industrial Relations & Labor	54	1,716	41	Womens Studies	71	3,421
21	International Relations	173	7,291				

Note. The JCR categories are sorted alphabetically. The JCR dataset was last updated on June 17, 2026, and retrieved on July 15, 2026.

providing further information on AI use. Third, if no information was found in the author instructions, submission guidelines, and/or editorial policies, the journal submission portal was searched for further information on AI use. However, the submission portal itself was not accessed (which would have required a login), but only searched for additional information on the login page. Importantly, the search was conducted from an author's perspective, simulating the process of an author seeking information about a journal's AI guidelines prior to manuscript submission.

Following this approach, the second author manually assessed a random subset of 10% of the unique journals (n = 18), resulting in no substantial disagreement between both authors that would have required a joint inspection. This indicates high quality and reproducibility of the search outcome as well as the robustness of the dataset obtained.

In a further step, to complement our data on journal-specific AI guidelines, the first author manually assessed the websites of all 29 publishers in our sample. As before, the first author determined the presence of publisher-specific AI guidelines, which was binary coded (yes = existence of publisher-specific guidelines; no = nonexistence of publisher-specific guidelines), and copied any applicable AI guidelines verbatim.

This manual approach was chosen over automated web scraping to ensure high data quality. Journal and publisher websites vary substantially in their front-end structure, which makes web scraping error-prone, while many websites also explicitly restrict or prohibit web scraping of any kind. Therefore, our manual approach does not only represent an important quality assurance layer but aligns with and respects legal terms and conditions. The manual data collection was conducted between June 29 and July 9, 2026 (journal-specific guidelines) as well as on August 6, 2026 (publisher-specific guidelines).

Data coding

To shed light on the content of the collected AI guidelines, the first author coded all journal- and publisher-specific guidelines. To this end, we inductively developed a coding scheme based on the data rather than using preconceived categories. The final coding scheme consists of 16 categories. Seven of these categories capture general AI provisions, referring to overarching rules and requirements that apply to AI use in general (e.g., AI authorship and accountability for AI use in the research process) rather than to specific use cases. The remaining nine categories capture the permission and disclosure of specific AI use cases (e.g., text production or data generation). After the first author applied this coding scheme to all AI guidelines, the second author independently coded a randomly selected subset of 10% of the journal-specific AI guidelines (n = 15) using the same coding scheme, resulting in an average agreement rate of 89% across the 16 categories. Given the small number of publisher-specific guidelines in the sample (n = 16), the second author independently coded all publisher-specific guidelines rather than a 10% subset, resulting in an average agreement rate of 79%.

Analytical strategy

In order to address RQ1, we examine the prevalence of AI guidelines. Before conducting any analysis, we separate the journals into three categories: (1) those with journal-specific AI guidelines, (2) those that only refer to publisher-specific guidelines, and (3) those with no

guidelines at all. Then, we report the overall share of each category across all journals under investigation.

In the next step, and to investigate potential differences in the prevalence of AI guidelines between JCR categories, we create an additive index. Specifically, we calculate a score for each JCR category by adding 1 point for journals with journal-specific AI guidelines, 0.5 points for journals that only refer to publisher-specific guidelines, and 0 points for journals with no guidelines at all. For journals providing both journal- and publisher-specific AI guidelines, only the journal-specific guidelines are considered for this index and the journal receives one point, as journal-specific guidelines represent the highest level of guidance captured by the index. Since each JCR category in our sample contains five journals, the resulting index ranges from 0 to 5. We then visualize this index by creating a bar chart, with the 41 JCR categories ranked according to the index score.

With respect to RQ2 we investigate the actual content of the AI guidelines by drawing on the manual coding. To provide a comprehensive overview of both journal- and publisher-specific guidelines, we first report the category shares referring to general AI provisions, which cover overarching rules and requirements that apply to AI use throughout the research process. Second, we report the category shares that refer to the permission and disclosure of specific AI use cases. This is done for journal- and publisher-specific guidelines, respectively. Given the dominant position of the five largest academic publishers in academic publishing ('Big 5'), we additionally provide the category shares for these publishers separately in the Appendix.

For replication purposes and transparency reasons, a Scientific Use File, including research data and analysis script, is available through Harvard Dataverse (see **[BLINDED]**). This includes a complete list of all (unique) journals in our sample including further information.

Results

Research question 1

Regarding RQ1, which focuses on the prevalence of AI guidelines, approx. 71% of the 180 unique journals under investigation have journal-specific AI guidelines. More specifically, approx. 57% have both journal- and publisher-specific AI guidelines, while approx. 13% have journal-specific guidelines only. Another 22% exclusively refer to publisher-specific guidelines, whereas approx. 8% have no AI guidelines at all. Prevalence is lower when it comes to the publishers: Only approx. 55% of the 29 publishers in our sample have AI guidelines.

The prevalence of AI guidelines also varies across JCR categories. To capture these differences, we calculated an additive index for each category, ranging from 0 to 5. A score of 1 point was assigned for journals with journal-specific AI guidelines, 0.5 points for journals that only refer to publisher-specific guidelines, and 0 points for journals without any AI guidelines. As each JCR category comprises five journals, higher (lower) index values indicate a higher (lower) prevalence of AI guidelines. Figure 1 presents the index scores across all 41 JCR categories, ranked from highest to lowest scores.

It can be observed that the index scores are relatively high across most JCR categories. More specifically, in 29 of the 41 categories, the index has a value of at least 4.0, implying that four or more of the five sampled journals provide some form of AI guidance, either through journal-specific or publisher-specific guidelines. Nine categories even reach the maximum

score of 5.0, indicating that all five journals have at least journal-specific AI guidelines. With respect to the lower end of the distribution, seven categories score 3.5, while five categories show scores of 3.0 or below. We also observe some disciplinary patterns. For example, health- and medicine-related JCR categories (e.g., "Health Care Sciences & Services," "Medical Ethics," and "Medicine, Legal") generally show relatively high index scores, with most of these categories having a score ranging between 4.0 and 4.5. Education-related categories (i.e., "Education, Scientific Disciplines," "Education, Special," and "Education & Educational Research") consistently score between 3.5 and 4.0. Other disciplinary fields, such as "Demography," "Communication," "Ethnic Studies," "Geography," and "Social Work," reach the maximum score of 5.0. By contrast, several broader social science fields, including "Economics," "Political Science," and "Public Administration," are among the categories with the lowest scores (3.0 and lower). Overall, the distribution suggests that AI guidelines are widespread across most disciplinary fields within the JCR "Social Sciences, General" group, but their prevalence varies considerably across JCR categories, with some notable differences even among related disciplinary fields (e.g., between "Political Science," with an index score of 2.5, and "International Relations," with an index score of 5.0).

Research question 2

Turning to RQ2, which focuses on the content of AI guidelines, the results reveal a consensus regarding some general AI provisions, but considerable variation in how specific AI use cases are permitted (or prohibited) and whether these must be disclosed. Table 2 provides an overview of the results on general AI provisions, considering both journal- and publisher-specific AI guidelines. These provisions cover overarching rules and requirements for AI use throughout the research process. As shown in Table 2, the guidelines show a strong consensus on *Authorship*: AI authorship is explicitly prohibited in approx. 77% of journal-specific as well as approx. 88% of publisher-specific guidelines, while no guideline under investigation permits AI to be listed as an author. Wherever *Accountability* is addressed in the guidelines, it is fully assigned to the human author. This applies to approx. 66% of the journal-specific and approx. 88% of the publisher-specific guidelines. Higher variation emerges when it comes to the *Extent of permission* and the *Extent of disclosure*. Journal-specific guidelines more frequently permit all AI use cases (approx. 59%; e.g., the guidelines do not restrict AI use to specific use cases), whereas publisher-specific guidelines more often permit only some use cases (approx. 50%; e.g., the guidelines restrict AI use to specific use cases by permitting language editing and grammar correction but prohibiting AI use for data generation or analysis code production). The extent of disclosure requirements show the opposite pattern, with approx. 63% of publisher-specific guidelines requiring disclosure of all use cases, compared to approx. 37% of journal-specific guidelines, which often exempt specific use cases, such as language editing and grammar correction, from disclosure. The *Place of disclosure* is mostly predetermined, meaning that the guidelines explicitly specify where authors should disclose their AI use rather than leaving this decision to the authors. In particular, journal-specific guidelines more frequently specify a predetermined place (e.g., the acknowledgments section), with approx. 89% of guidelines doing so compared to approx. 63% of publisher-specific guidelines.

However, the specific place of disclosure differs between journal- and publisher-specific guidelines. Among guidelines specifying a *Predetermined place of disclosure*, journal-specific

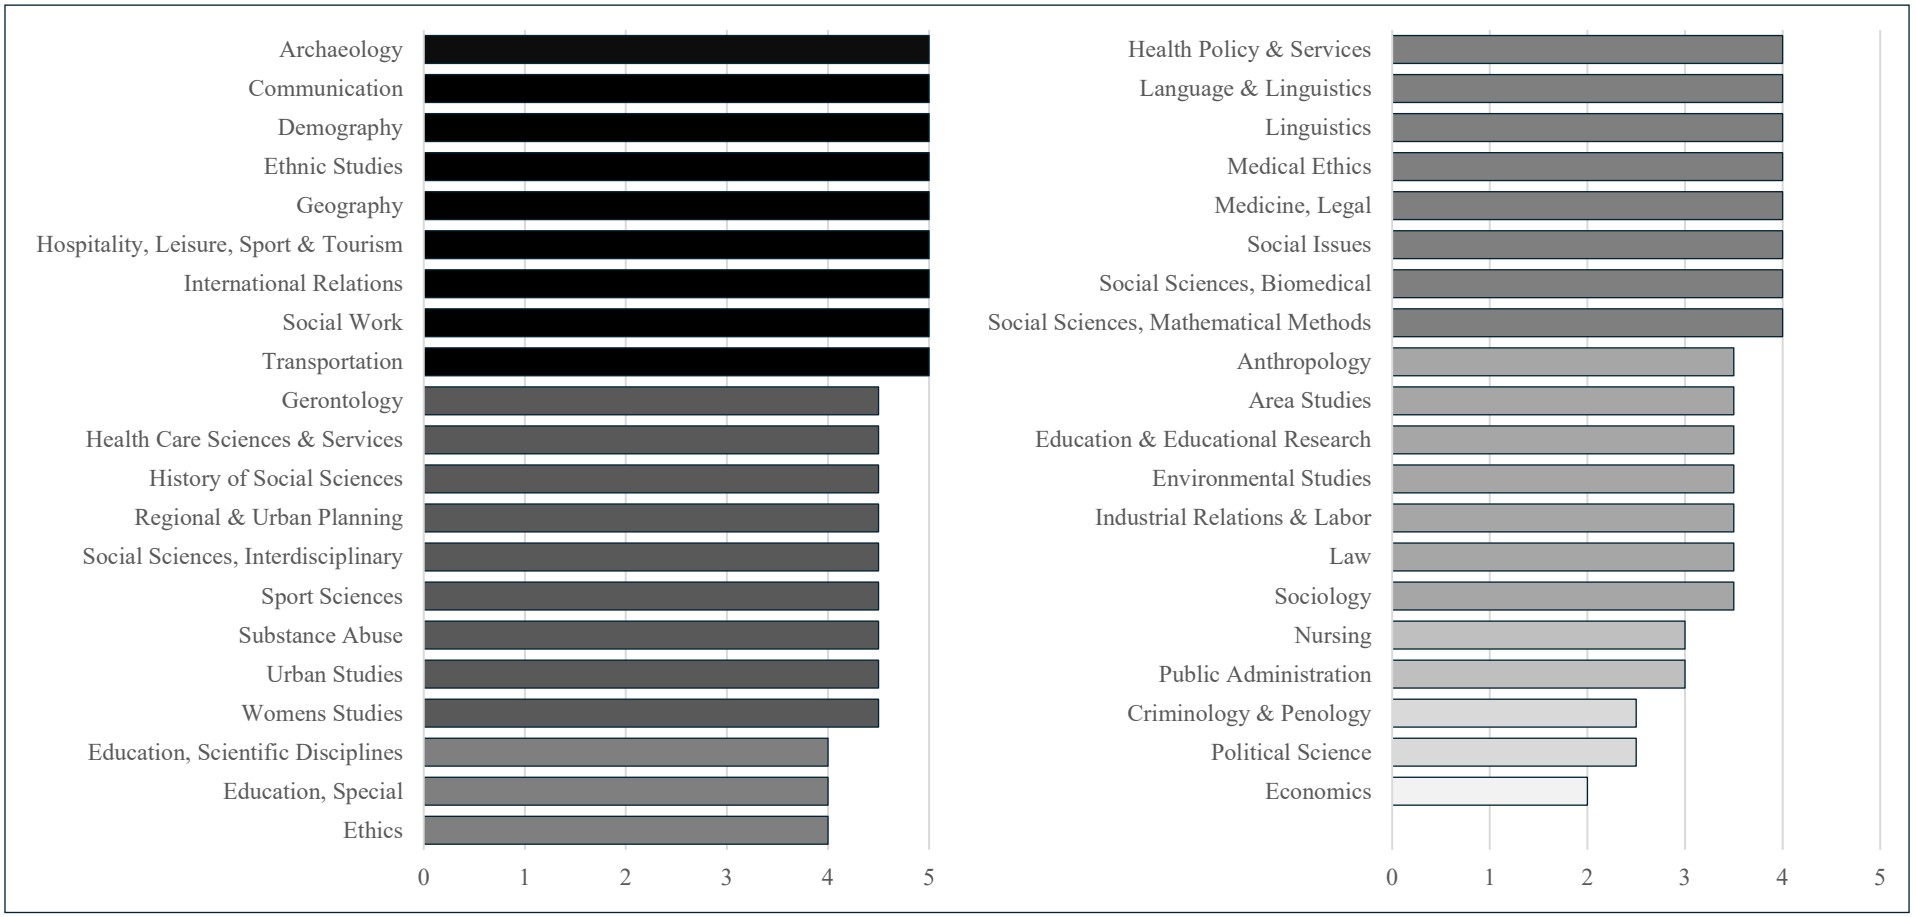


Figure 1. Additive index across the JCR categories in the "Social Sciences, General" group.

Note. The index ranges from 0 to 5, with higher (lower) values indicating a higher (lower) prevalence of AI guidelines among the five highest-ranked journals in each JCR category.

Table 2. General AI provisions in journal- and publisher-specific AI guidelines

Categories	Journal-specific	Publisher-specific
Authorship		
AI can be author	0	0
AI cannot be author	77.2	87.5
Not mentioned	22.8	12.5
Accountability		
Lies fully with author	65.5	87.5
Lies partially with author	0	0
Lies not at all with author	0	0
Not mentioned	34.5	12.5
Extent of permission		
All use cases	59.3	12.5
Some use cases	15.9	50
Not at all	0	0
Not mentioned	24.8	37.5
Extent of disclosure		
All use cases	37.2	62.5
Some use cases	51.7	25.0
Not at all	0	0
Not mentioned	11.0	12.5
Place of disclosure		
Predetermined	89	62.5
Author's choice	0	0
Not mentioned	11	37.5
Predetermined place of disclosure: Manuscript sections (e.g., methods)		
Yes	43.4	70
Predetermined place of disclosure: Acknowledgments		
Yes	22.5	50
Predetermined place of disclosure: Cover letter		
Yes	7	0
Predetermined place of disclosure: Disclosure statement		
Yes	58.9	30
Predetermined place of disclosure: Other place		
Yes	0.8	30
Reference to journal- or publisher-specific AI guidelines		
Yes	81.4	25
Reference to external AI guidelines (such as COPE and ICMJE)		
Yes	23.4	31.2

Note. For every category (e.g., "authorship"), each AI guideline was assigned one code only (e.g., "AI can be author" or "AI cannot be author"). Shares are reported in percentages. The *Place of disclosure* category is split into five dichotomous variables (i.e., "yes" or "no"), following the notion of check-all-that-apply categories, as some journals or publishers may list several options for the disclosure of AI use in their guidelines. For the five *Predetermined place of disclosure variables*, only "yes" percentages are reported (with "no" percentages inferred as complements to 100%), based only on guidelines specifying a predetermined place of disclosure. Due to rounding, percentages may not add up to 100%. Journal-specific AI guidelines: n = 145. Publisher-specific AI guidelines: n = 16. Journals and publishers without AI guidelines are not included in the table.

guidelines most frequently refer to a disclosure statement (approx. 59%; e.g., separate document that has to be submitted alongside the manuscript), whereas publisher- specific guidelines most often specify a manuscript section (approx. 70%; e.g., methods section). Journal-specific guidelines also frequently include a *Reference to publisher-specific AI guidelines*, with approx. 81% directing authors to consult such guidelines in addition to the journal-specific ones (e.g., the publishing ethics of the respective publisher). Conversely, publisher-specific guidelines less frequently include a *Reference to journal-specific AI guidelines*, with 25% instructing authors to additionally review journal-specific guidelines (e.g., the editorial policies of the respective journal).

References to external AI guidelines are less common, only occurring in approx. 23% of journal-specific and approx. 31% of publisher-specific guidelines, respectively. Among these references, the vast majority refers to the guidelines issued by COPE or ICMJE.

The regulation of specific AI use cases is also inconsistent. Table 3 provides an overview of the provisions on specific AI use cases in journal- and publisher-specific AI guidelines, respectively. As shown, *Language editing and grammar correction* is the most frequently permitted use case. This similarly applies to journal- (61%) and publisher-specific guidelines (75%). However, disclosure requirements differ substantially: only approx. 10% of journal-specific guidelines require disclosure of language editing and grammar correction, compared to approx. 56% of publisher-specific guidelines. *Figure creation* is also relatively frequently permitted, with approx. 44% of journal-specific and approx. 56% of publisher-specific guidelines permitting it. In contrast to language editing and grammar correction, the permission of creating figures through AI is generally accompanied by disclosure requirements. Importantly, the absence of explicit permission should not be interpreted as an explicit prohibition. For example, although language editing and grammar correction is frequently permitted, the remaining 39% of journal-specific and 25% of publisher-specific guidelines simply do not mention this use case. Neither journal- nor publisher-specific guidelines explicitly prohibit this use case. Thus, the absence of a provision does not necessarily indicate that a use case is prohibited, but rather that its permissibility remains unspecified. Beyond these use cases, which were most frequently addressed, substantial differences emerge in the extent to which other specific AI use cases are addressed at the journal- and publisher-specific level, respectively. Several specific AI use cases remain largely unaddressed in journal-specific guidelines. For example, *Brainstorming and summarization* as well as *Literature search* are not mentioned in any journal-specific guideline. Similarly, more than 80% of journal-specific guidelines do not mention *Classification and translation*, *Analysis code production*, or *Data generation. Peer-review assistance* is somewhat more frequently addressed at the journal-specific level but still remains largely unregulated: approx. 31% of journal-specific guidelines prohibit its use, while approx. 59% do not mention whether it is permitted. Disclosure requirements are even less frequently specified, with approx. 92% of journal-specific guidelines not addressing the disclosure of peer-review assistance. By contrast, several of these use cases are addressed more frequently at the level of publisher-specific guidelines. For example, classification and translation is permitted in approx. 56% of publisher-specific guidelines, compared to approx. 15% of journal-specific guidelines. A similar pattern can be observed for peer-review assistance, which is permitted in approx. 50% of publisher-specific guidelines, compared to approx. 10% of journal-specific guidelines. This difference is even more

pronounced with regard to disclosure requirements for peer-review assistance, with approx. 50% of publisher-specific guidelines requiring disclosure, whereas only approx. 3% of journal-specific guidelines do so. Thus, while several AI use cases remain unaddressed at the journal-specific level, publisher-specific guidelines address some of these uses more frequently, with particularly notable differences in permission and disclosure requirements.

Discussion

The increasing use of AI in academia has been accompanied by the development of guidelines by journals and publishers to govern its use. Against this background, this study examined the prevalence and content of AI guidelines of journals (and their publishers) indexed in the JCR, which is part of Clarivate's Web of Science database. More specifically, we investigated journals in the "Social Science, General" group.

Regarding the prevalence of AI guidelines (RQ1), our findings indicate that AI guidance in the form of AI guidelines has become relatively widespread in social science publishing. This development, however, is not uniform: the prevalence of AI guidelines is higher for journals than for publishers, and we observe no clear disciplinary pattern across the scientific disciplines in our sample. Our findings indicate a higher prevalence of journal-specific guidelines than reported by Goyanes et al. (2025), although differences in the research designs limit direct comparison. However, the comparatively high prevalence of AI guidelines does not necessarily imply comprehensive guidance on permitted and prohibited AI use in research and publishing.

Particularly, regarding the content of AI guidelines (RQ2), our findings reveal a distinction between relatively consistent general AI provisions and more heterogeneous guidance on specific AI use cases. We observed a broad consensus on human authorship and accountability in both journal- and publisher-specific guidelines, which is consistent with previous research on AI guidelines in academic publishing for journals (Ganjavi et al., 2024; Goyanes et al., 2025) and publishers (Chigwada & Ngulube, 2026). In contrast, guidance becomes less consistent when addressing specific AI use cases. Although publisher-specific guidelines generally provide broader guidance than journal-specific guidelines, several specific AI use cases remain unaddressed. Thus, the high prevalence of AI guidelines does not necessarily translate into clear and comprehensive standards for specific AI use cases and their disclosure.

Taken together, these findings highlight that effective AI guidance involves more than simply increasing the mere prevalence of AI guidelines. While some general AI provisions are becoming established, gaps in guidance on specific AI use cases leave researchers with substantial uncertainty about what constitutes permissible AI use and when, how, and where AI use should be disclosed. This uncertainty is also reflected in the vague terminology frequently used in AI guidelines. For example, the term "AI" often refers primarily to "LLMs," even though there are many other types of AI, such as machine learning. Unfortunately, these uncertainties and ambiguities place part of the responsibility for interpreting and applying existing AI guidance on individual authors, with potentially adverse consequences. If authors, journals, and publishers interpret a provision differently, authors may face criticism or scrutiny regarding their use of AI possibly only after publication. Uncertainty may be further increased when guidelines change during the submission and peer-review process, raising questions about which version of the guidelines applies. These issues highlight the need for clear, consistent,

and transparent guidance that not only specifies permitted AI use and disclosure requirements but also clarifies how changes in guidelines are handled and how disagreements regarding their interpretation are resolved.

At the same time, the rapid development of AI technologies raises the question of whether, and to what extent, guidelines can or should keep pace with the evolution of AI capabilities and the emergence of new AI use cases in academic research and publishing. Given the pace of development, it is questionable whether guidelines can, or should, attempt to regulate every emerging use case individually. Rather, AI guidance may need sufficiently concrete provisions that remain applicable as technologies and their uses evolve. For example, provisions requiring independent human oversight throughout the entire research process and verification of AI-generated outputs could apply across a wide range of existing and emerging use cases, without requiring separate rules for each application. To make such provisions more actionable, guidelines could require authors to describe in the methods section how human oversight was ensured throughout the research process, and how AI-generated outputs were reviewed and verified (e.g., in a dedicated subsection).

This, in turn, raises questions about who should define such standards, how their development should be conducted, and whose perspectives, expertise, and interests should be represented. Transparency in this process is particularly important. Given ongoing concerns about transparency in the development of AI technologies, the processes through which standards and guidelines for AI use within academic research are defined should not themselves become a 'black box.' In this process, the scientific community may have an important role in two respects: first, researchers as users of AI technologies can identify the practical requirements they face when disclosing and reporting AI use; second, researchers with expertise in AI, scientific integrity, research ethics, and related fields can contribute the expertise needed to define appropriate standards. At the same time, the development of AI guidance should also account for the perspectives and interests of other actors involved in academic publishing, including publishers and journals. Their perspectives are important for ensuring that guidelines can be implemented within existing publishing structures and practices. AI guidance should therefore be developed through a multi-stakeholder process that brings together those who use AI, those with relevant scientific, technological, and ethical expertise, and those responsible for implementing and applying these standards in academic publishing.

Contributions of the 'Big 5' publishers

Given their dominant position in academic publishing, the 'Big 5' publishers (i.e., Elsevier, SAGE, Springer Nature, Taylor & Francis, and Wiley) warrant particular attention. As these publishers account for a substantial share of high-ranked journals (e.g., 70% of the journals in our sample), their approach to AI guidance may influence how standards for permissible AI use are established. In particular, they may serve as a point of reference for other (smaller) publishers developing their own AI guidelines. As shown in the Appendix, the 'Big 5' do not substantially differ from other publishers regarding general AI provisions (see Table A1). This finding is consistent with previous research on the AI guidelines issued by the 'Big 5' (Lendvai & Petra, 2026) and publishers more broadly (Chigwada & Ngulube, 2026). Differences become more apparent in the regulation of specific AI use cases. Compared with the other publishers, the 'Big 5' address a broader range of use cases and more frequently provide explicit

permissions, including for use cases such as classification and translation or analysis code production (see Table A2). However, even these guidelines leave a substantial number of use cases unaddressed. Given their dominant position in academic publishing, these gaps further underscore the broader challenges identified above: incomplete guidance can place responsibility for interpreting and applying standards on individual researchers, while the development of such standards requires transparency and involvement of different stakeholders in the context of rapidly evolving AI technologies.

However, by serving as potential points of reference for smaller publishers, the guidelines of the 'Big 5' may not only shape the formal standards adopted by other publishers but also contribute to the institutionalization of norms regarding permissible AI use across academic publishing. Improving the clarity and comprehensiveness of the 'Big 5' guidelines, particularly regarding specific AI use cases, could therefore be a first step towards significant improvements in AI guidance for academic research.

Limitations and future research avenues

Our study has some (methodological) limitations that provide avenues for future research. First, our sample focuses on the five highest-ranked journals according to the 2025 JIF across 41 JCR categories within the "Social Sciences, General" group. The findings therefore provide insights into AI guidance among highly ranked social science journals but cannot be generalized to all social science journals or to journals in other scientific disciplines. We urge future studies to extend the analysis to a broader range of journals and scientific disciplines. In doing so, it can be examined whether the patterns identified in our study also hold beyond highly ranked social science journals, and whether the prevalence and content of AI guidelines vary systematically across disciplines. Such research could build on existing cross-disciplinary studies of AI guidance, including Ganjavi et al. (2024), who examined AI guidelines among leading publishers and journals across scientific fields, and Wang and Gong (2026), who compared AI policies for peer review across disciplines, as well as studies focusing on other scientific disciplines, such as Yin et al. (2025), who analyzed AI guidelines among medical journals.

Second, our analysis provides a cross-sectional assessment of AI guidelines at a time when AI technologies and their uses are developing rapidly. We therefore cannot determine how guidelines will evolve in response to technological developments or the emergence of new AI use cases. Longitudinal research could examine how quickly guidelines are revised in response to such developments, which emerging AI use cases become subject to explicit regulation, and whether guidelines follow or anticipate changes in AI use. Ganjavi et al. (2024) provide an initial indication of such dynamics: their repeated searches over a five-month period documented both an increase in the prevalence of AI guidelines and changes in their content. Future longitudinal research could build on this approach by systematically tracking changes in AI guidance over longer periods and relating them to developments in AI technologies and their use in academic research and publishing. This could help assess whether general AI provisions remain applicable as AI use expands beyond writing to other stages of the research process.

Third, our analysis focuses on formal AI guidelines rather than their interpretation and implementation in practice. We therefore cannot determine whether researchers are aware of the guidelines, how they interpret them, whether they follow them, or how editors apply them in the editorial process. Future research should therefore investigate how researchers and editors

perceive, interpret, and apply AI guidance, and whether differences in guidelines are associated with differences in AI use, disclosure practices, or editorial decision-making. Studies on researchers' actual AI use and perceptions, such as the one by Alvero et al. (2026), provide a useful starting point, but further research could explicitly connect such AI use and perceptions to the guidelines governing academic publishing.

Conclusion

In sum, our findings suggest that AI guidance in social science publishing is already relatively widespread and well established for general AI provisions concerning human authorship and accountability. However, considerable heterogeneity and uncertainty remain regarding the regulation and disclosure of specific AI use cases, especially beyond writing-related tasks. These findings also highlight the importance of considering AI guidance across journal and publisher levels, which may differ in their scope and specificity. Overall, current AI governance appears more settled in defining fundamental principles than in specifying the concrete conditions under which different AI use cases may be used and disclosed. The increasing prevalence of AI guidelines should therefore not be equated with the development of clear and comprehensive AI guidance. While guidelines can provide important orientation for researchers, their usefulness also depends on how clearly they address specific AI use cases and how their development responds to the rapid evolution of AI technologies and their uses. Thus, to move forward as a field, we need to discuss who should define standards for AI use in research and publishing, and how scientific, commercial, and institutional interests shape the development of standards for AI use in academic research.

References

Alvero, A. J., Stoltz, D. S., Stuhler, O., & Taylor, M. A. (2026). Generative AI in Sociological Research: State of the Discipline. *Sociological Science, 13*, 45-62. https://doi.org/10.15195/v13.a3

Argyle, L. P., Busby, E. C., Fulda, N., Gubler, J. R., Rytting, C., & Wingate, D. (2023). Out of one, many: Using language models to simulate human samples. *Political Analysis, 31*, 337–351. https://doi.org/10.1017/pan.2023.2

Butler, L. A., Matthias, L., Simard, M. A., Mongeon, P., & Haustein, S. (2023). The oligopoly's shift to open access: How the big five academic publishers profit from article processing charges. *Quantitative Science Studies, 4*(4), 778-799. https://doi.org/10.1162/qss_a_00272

Cao, L. (2026). Visualizing the adoption of Large Language Models across sociology subfields. *Socius: Sociological Research for a Dynamic World, 12*, 1–3. https://doi.org/10.1177/23780231261460657

Chigwada, J., & Ngulube, P. (2026). Use of artificial intelligence tools in the publishing process: expectations from publishers through author guidelines. *Frontiers in Research Metrics and Analytics, 11*, 1740510. https://doi.org/10.3389/frma.2026.1740510

Claassen, J., Revilla, M., & Höhne, J.K. (2026). Epistemology of traces: Human coding, LLM classification, and the production of knowledge. *ResearchGate*. https://doi.org/10.13140/RG.2.2.29363.80168

COPE Council. (2023) COPE position - Authorship and AI - English. https://doi.org/10.24318/cCVRZBms

de Winter, J. (2024). Can ChatGPT be used to predict citation counts, readership, and social media interaction? An exploration among 2222 scientific abstracts. *Scientometrics, 129*(4), 2469-2487. https://doi.org/10.1007/s11192-024-04939-y

Deutsche Forschungsgemeinschaft (2023). Stellungnahme des Präsidiums der Deutschen Forschungsgemeinschaft (DFG) zum Einfluss generativer Modelle für die Text- und Bilderstellung auf die Wissenschaften und das Förderhandeln der DFG. https://www.dfg.de/resource/blob/289674/ff57cf46c5ca109cb18533b21fba49bd/230921-stellungnahme-praesidium-ki-ai-data.pdf

Deutsche Forschungsgemeinschaft (2026). Leitlinie zur Nutzung von Künstlicher Intelligenz in der Begutachtung. https://www.dfg.de/resource/blob/387664/4-04-de.pdf

Ganguly, A., Johri, A., Ali, A., & McDonald, N. (2025). Generative artificial intelligence for academic research: evidence from guidance issued for researchers by higher education institutions in the United States. *AI and Ethics, 5*(4), 3917-3933. https://doi.org/10.1007/s43681-025-00688-7

Ganjavi, C., Eppler, M. B., Pekcan, A., Biedermann, B., Abreu, A., Collins, G. S., Gill, I. S. & Cacciamani, G. E. (2024). Publishers' and journals' instructions to authors on use of generative artificial intelligence in academic and scientific publishing: bibliometric analysis. *bmj, 384*. https://doi.org/10.1136/bmj-2023-077192

Garfield, E. (2006). The history and meaning of the journal impact factor. *JAMA, 295*(1), 90-93. https://doi.org/10.1001/jama.295.1.90

Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences, 120*(30), e2305016120. https://doi.org/10.1073/pnas.2305016120

Goyanes, M., Lopezosa, C., & Piñeiro-Naval, V. (2025). The use of artificial intelligence (AI) in research: a review of author guidelines in leading journals across eight social science disciplines. *Scientometrics, 130*(7), 3725-3741. https://doi.org/10.1007/s11192-025-05377-0

Haraway, D. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies, 14*(3), 575-599. https://doi.org/10.2307/3178066

Hayes, A. 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

Holtdirk, T., Assenmacher, D., Bleier, A., & Wagner, C. (2025). Learning from Convenience Samples: A Case Study on Fine-Tuning LLMs for Survey Non-response in the German Longitudinal Election Study. *arXiv*. https://doi.org/10.48550/arXiv.2509.25063

Höhne, J. K., Lenzner, T., & Claassen, J. (2025). Automatic speech-to-text transcription: Evidence from a smartphone survey with voice answers. *International Journal of Social Research Methodology, 28*, 625-632. https://doi.org/10.1080/13645579.2024.2443633

International Committee of Medical Journal Editors (2025). Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals. https://icmje.org/icmje-recommendations.pdf

Sebastian, J., Moon, J.-M., & Camburn, E. (2025). Enhancing Quantitative Analysis in Social Sciences With Large Language Models (LLMs): A Methodological Case Study in Educational Research. *AERA Open*, *11*. https://doi.org/10.1177/23328584251389621

Kuteeva, M., & Andersson, M. (2024). Diversity and standards in writing for publication in the age of AI—Between a rock and a hard place. *Applied Linguistics, 45*(3), 561-567. https://doi.org/10.1093/applin/amae025

Kim, E. (2025). Does publisher volume matter? A cross-sectional analysis of Scopus journal publishing patterns. *Publications*, *13*(2), 17. https://doi.org/10.3390/publications13020017

Kim, J., & Lee, B. (2023). Ai-augmented surveys: Leveraging large language models and surveys for opinion prediction. *arXiv*. https://doi.org/10.48550/arXiv.2305.09620

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

Lendvai, G. F., & Petra, A. (2026). Artificial intelligence in academic practices and policy discourses across 'Big 5'publishers. *Research Evaluation, 35*, rvag004. https://doi.org/10.1093/reseval/rvag004

Lenzner, T., & Hadler, P. (2026). Evaluating Large Language Model Performance in the Cognitive Pretesting of Questionnaires: A Human-AI Comparison. *ResearchGate*. https://doi.org/10.13140/RG.2.2.16697.33128

Liang, W., Zhang, Y., Wu, Z., Lepp, H., Ji, W., Zhao, X., Cao, H., Liu, S., He, S., Huang, Z., Yan, D., Potts, C., Manning, C., & Zou, J. (2025). Quantifying large language model usage in scientific papers. *Nature Human Behaviour*, 1-11. https://doi.org/10.1038/s41562-025-02273-8

Lin, H., & Zhang, Y. (2026). Navigating the Risks of Using Large Language Models for Text Annotation in Social Science Research. *Social Science Computer Review, 44*(3), 403-427. https://doi.org/10.1177/08944393251366243

Lindemann, N. F. (2025). Chatbots, search engines, and the sealing of knowledges. *AI & Society, 40*, 5063-5076. https://doi.org/10.1007/s00146-024-01944-w

Mühlhoff, R. (2018). Digitale Entmündigung und User Experience Design: Wie digitale Geräte uns nudgen, tracken und zur Unwissenheit erziehen. *Leviathan, 46*(4), 551–574. https://doi.org/10.5771/0340-0425-2018-4-551

Naddaf, M. (2026). Surge in fake citations uncovered by audit of 2.5 million biomedical-science papers. *Nature, 653*(8115), 658-659. https://doi.org/10.1038/d41586-026-00748-w

Nejjar, M., Zacharias, L, Stiehle, F., & Weber, I. (2025). LLMs for science: Usage for code generation and data analysis. *Journal of Software: Evolution and Process, 37*(1), Article e2723. https://doi.org/10.1002/smr.2723

Pangakis, N., & Wolken, S. (2024). Knowledge distillation in automated annotation: Supervised text classification with LLM-generated training labels. In *Proceedings of the sixth workshop on natural language processing and computational social science* (NLP+ CSS 2024) (pp. 113-131). https://doi.org/10.18653/v1/2024.nlpcss-1.9

Pereira, V., Basilio, M. P., & Santos, C. H. T. (2025). PyBibX–a Python library for bibliometric and scientometric analysis powered with artificial intelligence tools. *Data Technologies and Applications, 59*(2), 302-337. https://doi.org/10.1108/DTA-08-2023-0461

Peters, U., & Chin-Yee, B. (2025). Generalization bias in large language model summarization of scientific research. *Royal Society Open Science, 12*(4), 241776. https://doi.org/10.1098/rsos.241776

Revilla, M., Ochoa, C., Höhne, J.K., Couper, M.P. (2026). Transcribing and coding voice answers obtained in web surveys: comparing three leading automatic speech recognition tools. *Journal of Survey Statistics and Methodology.* https://doi.org/10.1093/jssam/smaf028

Rothschild, D., Marlar, J., Amaya, A., Barari, S., Buskirk, T., Cobb, C., Gennai, J., Hillygus, D., Korlakai Vinayak, R., Krupenkin, M., Lee, S., Steiger, D., & Webb, B. (2026). *Responsible AI Integration in Survey Research*. American Association for Public Opinion Research. https://aapor.org/wpcontent/uploads/2026/05/Responsible-AI-Integration-In-Survey-Research.pdf

Schwitter, N. (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), 61-65. https://doi.org/10.1177/20597991251313876

Stuhler, O., Ton, C. D., & Ollion, E. (2025). From Codebooks to Promptbooks: Extracting Information from Text with Generative Large Language Models. *Sociological Methods & Research, 54*(3), 794-848. https://doi.org/10.1177/00491241251336794

von der Heyde, L., Haensch, A.-C., & Wenz, A. (2026). Vox Populi, Vox AI? Using Large Language Models to Estimate German Vote Choice. *Social Science Computer Review, 44*(3), 549-571. https://doi.org/10.1177/08944393251337014

Von Eschenbach, W. J. (2021). Transparency and the black box problem: Why we do not trust AI. *Philosophy & technology*, *34*(4), 1607-1622. https://doi.org/10.1007/s13347-021-00477-0

Walters, W. H., & Wilder, E. I. (2023). Fabrication and errors in the bibliographic citations generated by ChatGPT. *Scientific reports, 13*(1), 14045. https://doi.org/10.1038/s41598-023-41032-5

Wang, Z., & Gong, M. (2026). A Cross-Disciplinary Analysis of AI Policies in Academic Peer Review. *Learned Publishing*, *39*(1), e2035. https://doi.org/10.1002/leap.2035

Wuttke, A., Aßenmacher, M., Klamm, C., Lang, M. M., Würschinger, Q., & Kreuter, F. (2025). AI conversational interviewing: Transforming surveys with LLMs as adaptive interviewers. In *Proceedings of the 9th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2025* (pp. 179–204), Albuquerque (USA). https://doi.org/10.18653/v1/2025.latechclfl-1.17

Yin, S., Huang, S., Xue, P., Xu, Z., Lian, Z., Ye, C., Ma, S., Liu, M., Hu, Y., Lu, P., & Li, C. (2025). Generative artificial intelligence (GAI) usage guidelines for scholarly publishing: a cross-sectional study of medical journals. *BMC medicine*, *23*(1), 77. https://doi.org/10.1186/s12916-025-03899-1

Zhou, Y., Liu, H., Srivastava, T., Mei, H., & Tan, C. (2024). Hypothesis generation with large language models. In *Proceedings of the 1st Workshop on NLP for Science (NLP4Science)* (pp. 117-139), Miami (USA). https://doi.org/10.18653/v1/2024.nlp4science-1.10

Appendix

Table A1. General AI provisions in the 'Big 5' publisher guidelines

Categories	
Authorship	
AI can be author	0
AI cannot be author	80
Not mentioned	20
Accountability	
Lies fully with author	100
Lies partially with author	0
Lies not at all with author	0
Not mentioned	0
Extent of permission	
All use cases	20
Some use cases	80
Not at all	0
Not mentioned	0
Extent of disclosure	
All use cases	40
Some use cases	60
Not at all	0
Not mentioned	0
Place of disclosure	
Predetermined	60
Author's choice	0
Not mentioned	40
Predetermined place of disclosure: Manuscript sections (e.g., methods)	
Yes	66.7
Predetermined place of disclosure: Acknowledgments	
Yes	33.3
Predetermined place of disclosure: Cover letter	
Yes	0
Predetermined place of disclosure: Disclosure statement	
Yes	66.7
Predetermined place of disclosure: Other place	
Yes	33.3
Reference to journal- or publisher-specific AI guidelines	
Yes	0
Reference to external AI guidelines (such as COPE and ICMJE)	
Yes	0

Note. For every category (e.g., "authorship"), each AI guideline was assigned one code only (e.g., "AI can be author" or "AI cannot be author"). Shares are reported in percentages. The *Place of disclosure* category is split into five dichotomous variables (i.e., "yes" or "no"), following the notion of check-all-that-apply categories, as some publishers may list several options for the disclosure of AI use in their guidelines. For the five *Predetermined place of disclosure variables*, only "yes" percentages are reported (with "no" percentages inferred as complements to 100%), based only on guidelines specifying a predetermined place of disclosure. Due to rounding, percentages may not add up to 100%. Percentages are based on the five largest academic publishers in the sample (n = 5): Elsevier, SAGE, Springer Nature, Taylor & Francis, and Wiley.

Table A2. Provisions on specific AI use cases in the 'Big 5' publisher guidelines

	Is the use permitted?			Is disclosure required?		
	Yes	No	NM	Yes	No	NM
Text production	20	60	20	20	20	60
Language editing and grammar correction	100	0	0	60	40	0
Brainstorming and summarization	60	0	40	40	20	40
Classification and translation	100	0	0	100	0	0
Analysis code production	80	0	20	80	0	20
Literature search	60	0	40	60	0	40
Figure creation	80	0	20	80	0	20
Data generation	40	40	20	40	0	60
Peer-review assistance	60	20	20	60	0	40

Note. Every use case (e.g., text production) was coded twice: whether it is permitted and whether disclosure is required. Shares are reported in percentages. Due to rounding, percentages may not add up to 100%. NM stands for not mentioned. Percentages are based on the five largest academic publishers in the sample (n = 5): Elsevier, SAGE, Springer Nature, Taylor & Francis, and Wiley.