Comparing Generative AI Disclosure Requirements: An Exploratory Document Analysis of Ten Publisher and Journal Policy Units (October 2026)

Rayan Abdurazag Khalat[1], Abdurazag Mohammed Khalat[2]

[1]Department of Dentistry, Faculty of Dentistry, Sabratha University, Surman, Libya
[2]Department of Electrical Engineering, Faculty of Engineering, Sabratha University, Surman, Libya, a.m.klat@gmail.com

Corresponding Author: E-mail: raykh757@gmail.com

Abstract:

Publishers and journals increasingly require transparency when authors use generative artificial intelligence (GenAI), but the operational details of disclosure remain uneven. This exploratory document analysis compares ten purposively selected publisher- or journal-level policy units against twelve criteria. Coding categories were explicit (E), conditional or partial (C), not explicit in sufficiently reviewed text (N), and unverified (U) when adequate primary-source evidence was unavailable. In response to reviewer-identified discrepancies, we rechecked the relevant official policy texts and updated the matrix and evidence register together. The revised matrix contains 79 E, 24 C, 17 N, and 0 U cells out of 120. The documents suggest recurring attention to human accountability and limits on AI authorship, while disclosure detail, location, language editing, research-stage use, images, and peer-review confidentiality vary across the selected documents.

Keywords: Generative artificial intelligence; Academic writing; Scholarly publishing; AI disclosure; Research integrity.

مقارنة متطلبات الإفصاح عن استخدام الذكاء الاصطناعي التوليدي: تحليل وثائقي استكشافي لسياسات عشر وحدات من جهات النشر والمجلات العلمية (أكتوبر 2026)

ريّان عبدالرزاق خلاط[1]، عبدالرزاق محمد خلاط[2]

[1] كلية طب الأسنان، جامعة صبراتة، ليبيا، raykh757@gmail.com
[2] قسم الهندسة الكهربائية، كلية الهندسة، جامعة صبراتة، ليبيا، a.m.klat@gmail.com

الملخص:

تتزايد مطالبة دور النشر والمجلات بالشفافية عند استخدام الذكاء الاصطناعي التوليدي، لكن التفاصيل الإجرائية للإفصاح لا تزال متفاوتة. تقارن هذه الدراسة الوثائقية الاستكشافية عشر وحدات مختارة قصدياً من سياسات الناشرين والمجلات وفق اثني عشر معيارًا. استُخدمت الفئات: صريح (E)، ومشروط أو جزئي (C)، وغير مذكور في النص الرسمي الذي تمت مراجعته بما يكفي (N)، وغير متحقق منه (U) عند عدم توافر أدلة أولية كافية. وبعد ملاحظات المراجعين، أُعيد فحص النصوص الرسمية ذات الصلة وتحديث مصفوفة الترميز وسجل الأدلة معًا. تتضمن المصفوفة 79 تصنيفًا صريحًا، و24 تصنيفًا مشروطًا أو جزئيًا، و17 تصنيفًا غير مذكور في النص الذي تمت مراجعته، ولا تتضمن خانات غير متحقق منها، من أصل 120 خانة. وتشير الوثائق إلى اهتمام متكرر بمسؤولية الباحثين وحدود إسناد التأليف إلى الذكاء الاصطناعي، مع اختلاف تفاصيل الإفصاح ومكانه ومعاملة التحرير اللغوي والاستخدام في مراحل البحث والصور وسرية التحكيم. وتقترح الدراسة إطارًا أدنى للإفصاح، دون الادعاء بقياس موثوق لاتساق السياسات؛ فالعينة قصدية ولم يُجرَ ترميز مستقل ثانٍ. وتظل النتائج استكشافية ولا تمثل تقديرات لقطاع النشر بأكمله.

الكلمات المفتاحية: الذكاء الاصطناعي التوليدي، الكتابة الأكاديمية، النشر العلمي، الإفصاح عن الذكاء الاصطناعي، النزاهة البحثية

Disclosure of Artificial Intelligence (AI) Tools: ChatGPT (OpenAI; GPT-6; accessed 10 October 2026) and Claude (Anthropic; Claude Sonnet 5.5; accessed 10 October 2026) were used for language editing, internal-consistency review, source discovery, an initial retrieval and draft-coding pass of policy texts, and re-verification

of selected policy sources and references (October 2026), and revision of the manuscript text and tables. The authors reviewed the final text and take full responsibility for its content. Because independent second coding was not completed, the policy classifications remain first-pass (Section 2.6).

1. Introduction

Generative artificial intelligence (GenAI) has moved from experimental use to an increasingly visible component of scholarly work. Large language models and related systems can support literature discovery, drafting, language refinement, coding, translation, data analysis, and other research tasks. Their rapid diffusion has created a policy problem for scholarly publishing: journals and publishers must permit useful forms of assistance without weakening authorship standards, confidentiality, reproducibility, or the reliability of the scientific record. The central difficulty is not whether AI should be used at all, but how its use should be classified, documented, and governed.

Early policy responses were heterogeneous. Ganjavi et al. (2024) found substantial variation among leading publishers and journals in the scope of permitted uses, disclosure location, and the level of detail required from authors, and they identified conflicts between some journal-level and publisher-level policies. In early 2023, leading journals stated that AI tools cannot be credited as authors (Nature, 2023; Thorp, 2023). The landscape has since continued to evolve rather than converge. Lendvai and Aczél (2026) examined AI-related policy discourse across five major publishers and reported a shared baseline around non-authorship and disclosure alongside differences in scope, specificity, and normative framing. Professional bodies have also set expectations, notably the ICMJE recommendations (ICMJE, n.d.) and the COPE position on authorship and AI tools (COPE, n.d.).

The need for a current comparative analysis is practical rather than merely chronological. Several publishers have revised their guidance recently (Wiley, 2025), and individual journals may add requirements of their own. The same activity, such as language editing or translation, may therefore be treated differently depending on the destination. This study updates the comparison to a fixed observation date and asks whether publisher-level and journal-level documents differ in the operational elements they require.

The aim is to describe patterns of convergence and divergence in GenAI disclosure requirements across ten purposively selected policy units and to derive a provisional minimum disclosure framework. The objectives are: (1) to code each unit against twelve criteria; (2) to summarize the distribution of first-pass codes without treating them as validated consistency estimates; (3) to classify disclosure location and level of detail where the official text supports a decision; (4) to compare publisher- and journal-level documents while acknowledging dependent units and incomplete verification; and (5) to propose a minimum disclosure framework for future evaluation.

Four research questions guide the analysis: (RQ1) What patterns of convergence and divergence appear across the ten policy units on the twelve criteria? (RQ2) What patterns exist in where disclosure must appear and how much detail it must contain? (RQ3) How do publisher-level and journal-level documents appear to differ on criteria where the evidence is sufficiently clear? (RQ4) Which disclosure elements recur in the retrieved documents and could form a provisional minimum framework?

2. Methodology

2.1 Study Design

A comparative document analysis with descriptive quantification was used to examine how scholarly publishers and selected journals regulate the use of GenAI in academic writing and publication (Bowen, 2009). Document analysis was selected because the object of the study is the wording, structure, and operational requirements of publicly available policies rather than the attitudes or behavior of individual researchers. The study is therefore non-interventional and does not involve human participants.

2.2 Sampling Strategy

The complete Research Evaluation journal-specific author-guidelines page was retrieved and reviewed on 17 October 2026. It requires disclosure in the cover letter and in Methods or Acknowledgements, and lists examples including generating content/images, writing code, processing data, and translation. The row is coded from this journal-specific page only; it does not imply that every Oxford University Press policy was reviewed.

2.3 Data Sources

Official publisher or journal policy pages were the intended primary sources; third-party summaries were not treated as evidence for coding. URLs and access dates were recorded where available, but complete archived copies and a cell-by-cell quotation-backed evidence log were not available for every unit. Reviewer comments identified several final codes that exceeded what the logged excerpts supported. Those cells have been downgraded to U or C rather than treated as resolved. This was targeted document retrieval, not a database-based systematic review; no PRISMA screening counts are claimed. ICMJE, CDC and Oxford University Press (general journal policy) guidance and peer-reviewed studies were used as contextual sources, not as coded units. Verification levels (V) are reported for each unit in Table A1; levels have been reduced where source coverage was incomplete.

2.4 Coding Framework

A codebook was written before the final comparison, giving a definition, a decision rule, and an example excerpt for each of twelve criteria: (1) prohibition of AI authorship; (2) requirement to disclose AI use; (3) identification of the AI tool; (4) description of the purpose or task; (5) model or version information; (6) explicit human verification or accountability; (7) specified disclosure location; (8) treatment of AI used in research methods, data analysis, or coding; (9) confidentiality rules for reviewers and editors; (10) treatment of AI-generated or AI-altered images; (11) prohibition of fabricated or manipulated data, results, or references; and (12) explicit consequences for non-compliance. The criteria were developed for this exploratory comparison with reference to domains examined in prior policy analyses, including Ganjavi et al. (2024). The operational questions are shown in Table 2.

Each criterion was coded E (explicit), C (conditional or partial), N (not explicit in the official text reviewed), or U (unverified because the text could not be retrieved or was incomplete). A code of E records that a requirement exists, not that two documents require the same thing; disclosure location and level of detail were therefore also classified into types (Sections 3.4 and 3.5). The purpose was not to assign a quality score to publishers, but to identify convergence, divergence, and areas of ambiguity.

2.5 Analysis

The analysis proceeded in two descriptive stages. First, each policy unit was coded against the twelve operational criteria. Second, the code distributions were summarized across the 120 possible cells, and criterion-level E/C/N distributions were calculated among codable cells only (U excluded from the denominator). These descriptive majority shares are not inter-coder reliability estimates or a validated consistency index. No independent second coding or Cohen's kappa was performed.

2.6 Coding Reliability

The first-pass coding was performed by a single coder. Independent second-coder assessment, percentage agreement, and Cohen's kappa were not available at the time of writing; inter-rater reliability was therefore not calculated. Because AI tools assisted the first retrieval and draft-coding pass (Section 2.7), the codes should not be treated as independent of that assistance. Future work should add independent double coding of at least a subset of cells, with agreement statistics reported only from actual double coding. The findings are reported as exploratory and should not be interpreted as a reliability-tested classification.

2.7 AI Tools Used in the Research Process

No generative AI output was treated as scientific evidence or as a substitute for source verification. AI assistance was limited to language editing, checking the internal consistency of the manuscript, and a first retrieval and draft-coding pass of policy texts. Policy classifications are retained as provisional wherever the official source text could not be confirmed (coded U or flagged in Table A1). This approach follows the conference template's requirement that AI use be described in the methodology and that authors retain responsibility for the accuracy and integrity of generated outputs.

3. Results and Discussion

3.1 Composition of the Policy Sample

The dataset contains ten purposively selected policy units: five publisher-level policies and five journal-level or journal-group guidance documents. Nature Methods is a dependent unit because it falls within Nature Portfolio. The complete official Research Evaluation author-guidelines page was retrieved and reviewed on 17 October 2026; its row is coded from that journal-specific page only, and this does not imply that every OUP-level policy was reviewed. These constraints limit cross-unit comparison.

Table (1): Policy documents included in the comparative sample

Level	Policy unit	Reason for inclusion / principal focus
Publisher	Elsevier	Current journal policy; disclosure, authorship, methods, images, reviewers
Publisher	Springer Nature	Current risk-based AI framework across authors, reviewers and editors
Publisher	Wiley	Current author AI guidelines; disclosure, verification, rights and accountability
Publisher	Taylor & Francis	Current AI policy; permitted/prohibited uses and disclosure
Publisher	SAGE	Current AI policy; assistive, generative and prohibitive categories
Journal	Nature Methods	Journal in the Springer Nature portfolio; applies the Nature Portfolio AI policy (dependent unit)
Journal	JAMA Network	Detailed author disclosure requirements for manuscript and research AI use
Journal	PLOS ONE	Current PLOS AI disclosure and integrity requirements
Journal group	BMJ journals	Current BMJ AI use policy for authors, reviewers and editors
Journal	Research Evaluation	Journal-level unit included for verification; only criteria supported by retrieved official text are coded

3.2 Coding Framework

Table (2): Structured coding framework used in the analysis

Criterion	Operational question
AI authorship	Does the policy explicitly prohibit AI from being an author/co-author?
Disclosure	Is substantive AI use required to be disclosed?
Tool identification	Must the AI tool/platform be named?
Purpose/task	Must the author describe why/how the tool was used?
Model/version	Is model/version information explicitly required?
Human accountability	Does the policy place responsibility and verification on human authors?
Disclosure location	Does the policy specify where disclosure must appear?
Research-stage AI	Does the policy distinguish manuscript preparation from AI used in methods/data analysis?
Reviewer confidentiality	Are unpublished manuscripts protected from public AI systems?
Images	Does the policy regulate AI-generated or AI-altered images?
Data integrity	Are fabrication/manipulation of data, results or references prohibited?
Enforcement	Does the policy state consequences for inappropriate or undisclosed use?

The revised matrix contains 120 policy–criterion cells: 79 explicit (E, 65.8%), 24 conditional or partial (C, 20.0%), 17 not explicit in sufficiently reviewed official text (N, 14.2%), and no unverified (U) cells. These descriptive counts from Table A1 (visualized in Figure 1) are not a policy-quality score or a validated estimate of consistency. Percentages may not total 100 because of rounding. U cells are not interpreted as evidence that a requirement is absent.

Criterion-level descriptive distributions (E/C/N; U excluded) were: C1, 9/1/0 (n=10; E=90.0%); C2, 9/1/0 (n=10; E=90.0%); C3, 6/1/3 (n=10; E=60.0%); C4, 5/3/2 (n=10; E=50.0%); C5, 2/3/5 (n=10; no category exceeded 50%); C6, 9/0/1 (n=10; E=90.0%); C7, 9/1/0 (n=10; E=90.0%); C8, 8/2/0 (n=10; E=80.0%); C9, 7/2/1 (n=10; E=70.0%); C10, 7/3/0 (n=10; E=70.0%); C11, 5/4/1 (n=10; E=50.0%); and C12, 3/3/4 (n=10; no category exceeded 50%). These exploratory proportions describe only this purposive sample and are not a validated measure of policy consistency.

3.3 Areas of Convergence

The reviewed documents commonly address AI authorship and human accountability, but the pattern should be read cautiously because policy-unit dependence and source scope limit inference. The matrix distinguishes a criterion not explicitly addressed in the reviewed official text (N) from a substantive requirement that is present (E or C).

The retrieved policies generally emphasize transparency for substantive AI use, but they differ in the threshold for disclosure and in whether low-risk language editing is exempt. Springer Nature's current risk-based framework treats disclosure as enhancing transparency for assistive use, while Nature Portfolio's Nature Methods editorial (Nature Methods, 2026) explicitly says that use of generative AI to write or edit for readability must be declared in the Methods section. These are related but not identical statements, so they should be coded separately and interpreted in context.

Human accountability is a recurring principle in the official policy text that was adequately reviewed. Wiley's official guidance explicitly states that authors remain fully accountable and must personally review and verify AI-generated content. The matrix records only whether each criterion is explicit in the particular policy text reviewed; it does not generalize a requirement to every unit without evidence.

3.4 Variation in Disclosure Location

Disclosure location varies across the official sources retrieved. The Nature Methods editorial (Nature Methods, 2026) states that generative AI used to write sections or edit for readability must be declared in Methods. Oxford University Press's general journal policy (Oxford University Press, 2026; a contextual source, not a coded unit) says AI use should be disclosed in both the cover letter and Acknowledgements, with Methods disclosure also relevant when AI is integral to the research methodology; journal-specific instructions may vary. For other publishers, locations are reported only where the relevant official policy page was checked. The revised matrix records N where a disclosure location is not specified in the text reviewed and C where the guidance is conditional or partial.

The JAMA Network's updated 2026 guidance, PLOS ONE, and BMJ specify disclosure locations in their official documents (Flanagin et al., 2026; PLOS ONE, n.d.; BMJ, n.d.). Research Evaluation's official author guidelines state that AI use should be disclosed in the cover letter and in the Methods or Acknowledgements section (Research Evaluation, n.d.). Nature Methods applies the Nature Portfolio framework, and its 2026 editorial explicitly states that generative AI used to write manuscript sections or edit for readability must be declared in Methods. The Research Evaluation coding is based on its journal-specific author-guidelines page; criteria not addressed in that text are coded N rather than inferred.

Several patterns therefore coexist in the sources retrieved: a separate declaration statement, Methods or Acknowledgements at the author's choice, an Acknowledgements section, a dedicated Methods section, and an acknowledgement in the contributor section. Criterion C7 in Table A1 records only whether a location is specified; a unit-by-unit mapping of these patterns requires the cell-level evidence log that is not yet available. The documents agree on the need for transparency but not on where it must appear.

3.5 Variation in the Level of Disclosure Detail

The level of detail required in a disclosure differs among policies. Wiley's official guidance requires disclosure of AI use and recommends documenting the tool name and version among other details; because the version detail is recommended rather than an unequivocal mandatory field for every use, Wiley C5 is coded C. The Nature Methods editorial requires declaration in Methods for generative AI used to write manuscript sections or edit for readability, but does not specify a required tool name or separate purpose/task field. PLOS ONE requires the names of tools used but does not explicitly require model/version information. BMJ requires the technology name and a description of why/how it was used, but does not explicitly require model/version information. The complete Research Evaluation author-guidelines page specifies disclosure in the cover letter and in Methods or Acknowledgements; it lists code writing and data processing among examples of AI use requiring disclosure, but does not explicitly require model/version information.

Model/version information should be distinguished from naming the AI technology. The revised matrix codes model/version as explicit for Taylor & Francis and JAMA Network, conditional/partial for Elsevier, Springer Nature, and Wiley, and not explicit in the sufficiently reviewed text for SAGE, Nature Methods, PLOS ONE, and BMJ. Research Evaluation C5 is N because the complete journal-specific author-guidelines page reviewed does not explicitly require model/version information. These codes describe the specific official texts reviewed, not all possible publisher-level requirements.

3.6 Treatment of Language Editing versus Substantive Content Generation

Policies differ in their treatment of language editing. Oxford University Press's general journal policy (Oxford University Press, 2026) exempts basic spelling and grammar correction but typically expects disclosure for substantial editing, rewriting, translation, or data-related uses. Nature Portfolio's 2026 editorial states that AI used to write sections or edit manuscripts for readability must be declared in Methods. These statements illustrate variation in scope; the instructions of individual journals may differ.

More substantive rewriting, translation, literature synthesis, and code generation may trigger disclosure, but the exact threshold is publisher- and journal-specific. Wiley's current AI guidance (Wiley, n.d.; Wiley, 2025) describes disclosure for AI used in research methodology, including study design, data collection, literature review, and data analysis.

3.7 AI Use in Research Methods and Data Analysis

Research-stage AI use is addressed differently across the retrieved documents. Wiley's guidance, as announced in 2025 (Wiley, 2025), covers disclosure of AI use in study design, data collection, literature review, and data analysis. Oxford University Press's general journal policy (Oxford University Press, 2026) also distinguishes AI use integral to methodology from routine language editing. These are general publisher-level statements, and journal-specific instructions may differ.

This distinction is important because a simple writing disclosure is insufficient when AI contributes to the production or interpretation of research evidence. In such cases, readers need methodological information that permits evaluation and, where feasible, reproducibility.

3.8 Images, Data Integrity, and Prohibited Uses

Rules for images vary more than the headline principle suggests. Elsevier distinguishes explanatory images (permitted, with disclosure in the caption), data visualizations derived from underlying data (permitted when reproducible methods are reported), and primary research images (not permitted), and it bars general-purpose generative image tools for graphical abstracts (Elsevier, 2026). Springer Nature permits AI-generated or AI-assisted visual content derived from verifiable inputs and does not permit content without verifiable source material (Springer Nature, 2026). Taylor & Francis permits process diagrams and concept illustrations with a figure-legend statement naming the tool and version, and does not permit research or clinical images (Taylor & Francis, n.d.).

The retrieved policies prohibit fabrication or misrepresentation of research data and results in several explicit cases. This criterion is not generalized to every unit: the final text distinguishes directly documented prohibitions from partial evidence and cases where the reviewed policy text does not explicitly address the criterion.

3.9 Peer Review and Confidentiality

The retrieved Nature Portfolio policy explicitly warns against sharing manuscripts, peer-review reports, or sensitive data with unsecured or public AI systems. Wiley's announcement of its 2025 guidelines states that uploading unpublished manuscripts to AI tools is prohibited (Wiley, 2025), and the cell is coded explicit. Research Evaluation's own author guidelines do not address AI in peer review, and that cell is coded N.

AI use in peer review is addressed directly in several official policies. Taylor & Francis prohibits uploading unpublished manuscripts to AI tools and prohibits using AI to generate review reports, while BMJ (n.d.) and PLOS (n.d.) explicitly protect confidential submissions and restrict public-tool uploads. SAGE also prohibits AI-generated review reports and confidential manuscript upload. Coverage is not identical across the sample, and

the coding records explicit or conditional language only where it appears in the reviewed policy text. The evidence is qualitative and is not used to estimate an industry-wide prevalence.

Publisher-level and journal-level patterns are reported descriptively only. Although no cells remain coded U after the targeted source checks, source coverage still differs across policy units and several N codes mean only that a requirement was not explicit in the official text reviewed. Therefore, no between-group comparison is interpreted as a stable difference. The supplementary file records the revised matrix and source-verification limitations; future comparisons should be recalculated only after broader and consistently documented source coverage.

Table (3): Main convergence and divergence across publisher-level and journal-level policies

Policy domain	Publisher-level pattern	Journal-level pattern	Main observation
AI authorship	Explicitly prohibited in most reviewed units; PLOS ONE is conditional/partial in this matrix	Consistent with the retrieved documents, but not established for every unit	Strongest area of convergence
Disclosure	Substantive use is addressed in several policies; thresholds and exemptions differ	Disclosure is commonly addressed, but the level of detail and exact requirements vary	Common principle; inconsistent implementation
Tool/purpose details	Requirements vary; some units distinguish assistive from generative use	JAMA Network and BMJ guidance provide detail; other units differ in specificity and disclosure scope	The same activity may be exempt, conditional, or reportable depending on the policy
Language editing	Treatment of routine language correction differs; Springer Nature's manuscript-preparation page asks authors to describe even low-risk use, whereas some policies exempt basic correction	Nature Methods (Nature Portfolio) requires a Methods-section declaration even for editing for readability; coverage of other journal-level units is not fully verified	The disclosure threshold for language editing is not uniform across policies
Research methods/data	Research-stage use is explicit in several policies; remaining units differ in scope or do not state specific requirements	Methods-level reporting appears in some guidance; coverage is not fully verified	Methodological AI may require more than a writing declaration, depending on the policy
Images	Image controls are explicit in some policies; coverage varies	Insufficient verified coverage for a broad journal-level comparison	High-risk area because provenance and primary-data integrity matter
Peer review	Confidentiality is explicit in some retrieved policies; other cells remain C or U	Some policies explicitly restrict sharing manuscripts with public AI systems	The sample is too incomplete to quantify the pattern reliably
Enforcement	Explicit at SAGE, JAMA Network, and BMJ; general or conditional elsewhere; not stated on the Elsevier page read	Conditional or not retrieved	Sanctions are rarely specified; N indicates that the reviewed AI policy text did not state a specific sanction

3.10 Permitted and Prohibited Uses

Several documents use a risk continuum rather than a binary permitted/prohibited model, for example Springer Nature's three tiers, SAGE's assistive, generative, and prohibitive categories, and Elsevier's image categories. Basic language correction is commonly treated as low risk, although not uniformly (Section 3.6). Substantive rewriting, translation, literature synthesis, coding, data analysis, and image generation are more likely to require disclosure and human verification. Fabrication of data, results, citations, or primary research images is not permitted in the documents where the point was coded. The "Recommended disclosure level" column of Table 4 is the authors' proposal and is not a requirement of any single policy.

Table (4): Risk-sensitive treatment of common AI uses

AI use	Typical policy treatment	Recommended disclosure level
Spelling, grammar, punctuation	Often exempt when purely assistive	Usually not required
Substantive rewriting/restructuring	Generally permitted only with human oversight	Disclose tool and purpose
Translation of scientific text	Generally permitted with verification	Disclose when substantive
Literature discovery/synthesis	Permitted by several policies with verification	Disclose when it affects content or analysis
Code generation/checking	Permitted with validation; research use belongs in Methods	Detailed disclosure
Data analysis/interpretation	Permitted conditionally; human validation required	Methods-level disclosure
AI-generated explanatory figures	Conditional and tool-dependent	Disclose in manuscript/caption as required
AI-generated primary research evidence	Generally prohibited	Not an acceptable substitute for research
Fabricated data/references/results	Prohibited	Not remedied by disclosure
Uploading confidential manuscripts for peer review	Generally prohibited	Not permitted without a secure, authorized environment

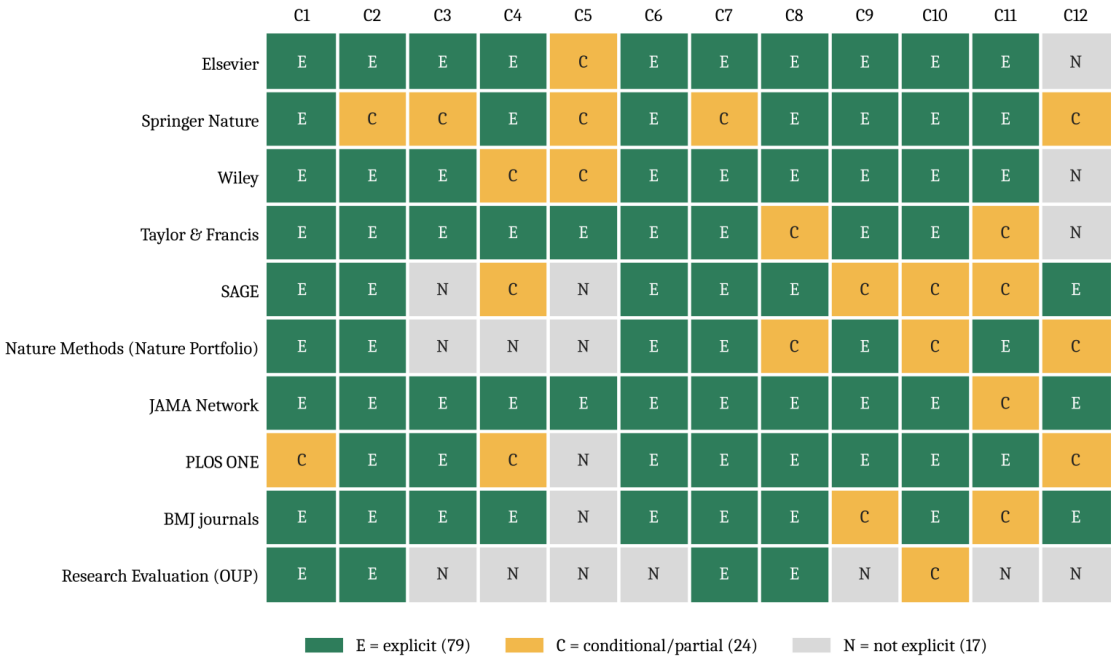


Figure (1): Revised policy coding matrix (12 criteria × 10 policy units). E = explicit, C = conditional/partial, N = not explicit in sufficiently reviewed text

3.11 Enforcement and Consequences

Consequences of non-disclosure differ across the sources retrieved. SAGE states that editors may reject submissions where inappropriate AI use was not disclosed. Springer Nature, PLOS ONE, JAMA Network, and BMJ describe possible editorial action or restrictions in specified circumstances. The Elsevier journal policy page reviewed did not state sanctions; its separate publishing-ethics policy was not coded here. For other units, N means the reviewed AI policy text did not state a specific sanction, while C records conditional or general enforcement language. These observations describe the pages reviewed, not the full enforcement systems of each publisher.

3.12 Discussion

The documents reviewed suggest recurring attention to a common ethical baseline: AI systems are not treated as authors, and human researchers retain responsibility for published work. This interpretation is consistent with

the broad direction of the 2024 bibliometric analysis by Ganjavi et al. (2024). Targeted official-source checks resolved several reviewer-identified discrepancies, but nine cells remain U and no independent second coding was performed. Accordingly, the study identifies exploratory patterns in the selected documents rather than establishing uniformity across scholarly publishing.

The more consequential problem lies below this common baseline. Authors are not simply asking whether disclosure is required; they need to know what counts as disclosure, where it belongs, what information must be reported, and whether the same use is permitted by the target journal. The present comparison shows that these questions do not have uniform answers. A statement such as "AI use must be disclosed" therefore has limited practical value unless accompanied by an operational definition and a minimum set of reporting elements.

This finding supports a distinction between policy principle and policy implementation. At the principle level, the sampled organizations are relatively aligned. At the implementation level, they vary in disclosure location, terminology, technical detail, exemptions for language editing, treatment of research-stage AI, and enforcement. Variation also appears inside a single publisher: Springer Nature's manuscript-preparation page asks authors to describe even low-risk use, whereas the Nature Portfolio page presents disclosure of such use as enhancing trust. The result is a fragmented compliance environment in which authors may comply with one policy and still fail to meet another because the two systems use different definitions of "disclosure," different locations, or different thresholds for what counts as substantive use.

A second issue is the inconsistent boundary between language assistance and intellectual contribution. Exempting routine grammar and spelling correction is reasonable because these tools do not ordinarily alter the substantive scientific content. However, the boundary becomes less clear when AI substantially restructures paragraphs, synthesizes literature, proposes arguments, translates scientific claims, generates code, or recommends analytical approaches. Nature Portfolio's current risk-based framework is particularly useful because it distinguishes low-risk assistive uses from evaluative or interpretive uses and from prohibited uses that replace scholarly judgment (Springer Nature, 2026). A similar risk-based structure can be observed in SAGE's distinction between assistive, generative, and prohibitive uses (SAGE, n.d.).

Third, disclosure and documentation should not be treated as identical concepts. A short statement that "ChatGPT was used for editing" may satisfy a journal's disclosure requirement, but it may not be sufficient to reconstruct a research process when AI was used for analysis, coding, evidence synthesis, or interpretation. For methodological uses, the appropriate standard should therefore be closer to reproducible reporting: tool identity, version or model, task, input and output role, validation procedures, and the effect on the study. Guidance from Elsevier, Taylor & Francis, and JAMA Network moves in this direction. CDC (2026) recommends identifying the content affected, the action taken, the AI tool, the purpose, and the extent of human oversight; for analytic or methodological uses, it also recommends sufficient details on prompts, settings, inputs, and validation steps to support reproducibility, subject to security requirements. A contrasting position argues against mandatory full disclosure of prompts and outputs for ordinary writing assistance, while supporting targeted methods-level detail where scientific validity or reproducibility requires it (Earp et al., 2026).

Fourth, the policy landscape is changing quickly. Elsevier's journal policy page carries an update date of June 2026, Wiley announced new guidelines in October 2025, and BMJ states that it reviews its AI policy every six months until further notice (BMJ, n.d.). Policy audits therefore become outdated quickly, and future studies should record the exact access date, policy version when available, URL, and an archived copy of every document analyzed. The present study fixes the observation date at 7 October 2026. Complete archived copies were not retained for every policy unit.

The comparison also has implications for authors from multilingual and resource-constrained research environments. AI-assisted language editing can reduce barriers to communicating scientific work in English, and several publishers explicitly permit or recognize such assistance. At the same time, unclear disclosure thresholds may place a disproportionate compliance burden on researchers who use language tools extensively. A proportionate policy should therefore distinguish between language polishing that preserves the author's reasoning and uses that generate or transform substantive scholarly content. This is preferable to treating all AI assistance as equivalent.

Based on the comparative findings, the study proposes a five-part minimum disclosure framework, derived from elements that recur in the documents coded (Elsevier, Taylor & Francis, JAMA Network, BMJ journals) and in CDC (2026) guidance. The framework is a proposal and has not been validated empirically:

1. **Tool identification:** name the AI system and, where relevant, model/version.
2. **Purpose and task:** state why and for what task the tool was used.
3. **Location or affected component:** identify the manuscript section, figure, code, analysis, or research stage affected.
4. **Human verification:** state how outputs were checked, validated, edited, or rejected.
5. **Responsibility and scope:** confirm that the authors retain responsibility for accuracy, originality, attribution, and research integrity.

The framework should be applied proportionately. Routine spelling and punctuation checks may remain exempt. Substantive rewriting, translation that can affect scientific meaning, literature synthesis, code generation, data analysis, and image generation should normally trigger disclosure. Uses that generate or alter primary research evidence, fabricate data or citations, or replace accountable scholarly judgment should be prohibited rather than merely disclosed. The five elements draw on the matrix as follows: tool identification on C3 and C5, purpose and task on C4, location or affected component on C7 and C8, human verification on C6, and responsibility and scope on C6 and C11. Model/version information (C5) is explicit in only two of ten units, so its inclusion is a recommendation rather than a description of prevailing practice.

This approach also addresses a weakness identified in earlier policy research: the absence of a standardized guideline-development process and the persistence of conflicting or heterogeneous journal-level rules (Ganjavi et al., 2024). Standardization should not mean forcing every journal to adopt identical wording. Instead, publishers and journals could agree on a common taxonomy and minimum disclosure fields while retaining discipline-specific restrictions. Such harmonization would reduce ambiguity without eliminating legitimate differences between research fields.

The revised matrix and supplementary evidence register were updated together for the reviewer-identified discrepancies. Wiley C3 and C6 are coded E based on the official Wiley guidance; Wiley C5 is C because the tool name/version details are recommended disclosure elements rather than an unequivocal universal requirement. Nature Methods C2 and C7 are E because the editorial says qualifying use must be declared in Methods; C3 and C4 are N within the scope of that editorial because it does not explicitly specify tool naming or purpose/task fields. PLOS ONE C3 is E and C5 is N based on the full official policy page reviewed. BMJ C3 and C12 are E and C5 is N based on the official policy page reviewed. Research Evaluation is marked A after review of the complete official journal-specific author-guidelines page. C2 and C7 are E; C8 is E because code generation and data processing are named as AI-use examples requiring disclosure; C10 is C because AI-generated images are mentioned but no comprehensive image rule is stated in this page. C3, C4, C5, C6, C9, C11 and C12 are N within the scope of this page because no explicit requirement addressing those criteria appears in the reviewed text. No U cells remain after reviewing the complete journal-specific author-guidelines page, and the counts should not be interpreted as complete verification of all policy criteria.

4. Conclusion

In the documents reviewed, the comparative analysis suggests recurring principles: generative AI cannot assume authorship, human researchers remain accountable for published work, and substantive AI use should be transparent. The evidence is not yet sufficient to conclude that these principles are uniformly implemented across all ten units. Differences remain in disclosure location, required detail, treatment of language editing, research-stage AI, reviewer confidentiality, and enforcement.

The study therefore recommends moving from generic disclosure statements toward a harmonized, risk-proportionate framework. A common minimum disclosure should identify the tool, purpose, affected content or research stage, human verification process, and continuing author responsibility. Routine language correction can remain outside mandatory disclosure, whereas AI use that influences evidence, analysis, interpretation, substantive content, code, or research images should be disclosed in a more detailed and reproducible manner.

Fabrication, falsification, undisclosed substantive use, and delegation of scholarly judgment to AI should remain prohibited.

Such a framework would not eliminate journal-specific policies. Instead, it would create a common compliance floor across publishers and journals, reducing avoidable ambiguity for authors while strengthening transparency, reproducibility, and trust in the scholarly record.

Ethical Approval: Not applicable. The study analyzed publicly available publisher and journal policy documents and did not involve human participants, personal data, intervention, or identifiable research records.

Scientific Integrity: The authors affirm that the manuscript is their original work, that sources are cited, and that the use of AI tools is disclosed in the statement above. The revised matrix contains no U cells after the complete official Research Evaluation author-guidelines page was retrieved and reviewed. This does not mean every possible related policy was reviewed; N records only that a criterion was not explicit in the sufficiently reviewed policy-unit text. The study remains a single-coder exploratory document analysis; independent second coding and inter-rater reliability statistics were not completed. The evidence register distinguishes direct official text from limited excerpts and does not treat N as proof of absence from all related policies.

Conflict of Interest: The authors declare no conflicts of interest.

Funding: No external funding was received for this study.

Data Availability: The revised coding matrix (twelve criteria for ten policy units, with verification levels) is provided in Appendix A and in the supplementary file. URLs and source-verification notes are supplied in the supplementary material. A complete cell-level quotation-backed evidence log and archived copies were not available for every decision at the time of writing; the matrix reports only the classifications supported by the official text reviewed.

References

BMJ. (n.d.). AI use. [Online]. Available at: https://authors.bmj.com/policies/ai-use/ [Accessed: 17 October 2026].

Bowen, G.A. (2009). "*Document analysis as a qualitative research method*", Qualitative Research Journal, 9 (2), pp. 27–40.

Centers for Disease Control and Prevention (CDC). (2026). Considerations for disclosing generative AI use in scientific work. [Online]. Available at: https://www.cdc.gov/ai/resources/considerations-for-generative-ai-use-in-scientific-work.html [Accessed: 7 October 2026].

Committee on Publication Ethics (COPE). (n.d.). Authorship and AI tools. [Online]. Available at: https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools [Accessed: 7 October 2026].

Earp, B.D., Schüklenk, U., Savulescu, J. et al. (2026). "*Against mandatory prompt disclosure in AI-assisted scholarship*", AI & Society.

Elsevier. (2026). Generative AI policies for journals. [Online]. Available at: https://www.elsevier.com/en-gb/about/policies-and-standards/generative-ai-policies-for-journals [Accessed: 10 October 2026].

Flanagin, A., Perlis, R.H. and Bibbins-Domingo, K. (2026). "*Updated guidance for author use of AI in medical publication*", JAMA, 336 (11), pp. 926–928.

Flanagin, A., Pirracchio, R., Khera, R., Berkwits, M., Hswen, Y. and Bibbins-Domingo, K. (2024). "*Reporting use of AI in research and scholarly publication—JAMA Network guidance*", JAMA, 331 (13), pp. 1096–1098.

Ganjavi, C., Eppler, M.B., Pekcan, A., Biedermann, B., Abreu, A., Collins, G.S., Gill, I.S. and 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, e077192.

International Committee of Medical Journal Editors (ICMJE). (n.d.). Use of artificial intelligence in publishing. [Online]. Available at: https://icmje.org/recommendations/browse/artificial-intelligence/ [Accessed: 7 October 2026].

Lendvai, G.F. and Aczél, P. (2026). "*Artificial intelligence in academic practices and policy discourses across 'Big 5' publishers*", Research Evaluation, 35, rvag004.

Nature Methods. (2026). "*Using AI responsibly in scientific publishing*", Nature Methods, 23, p. 271.

Nature. (2023). "*Tools such as ChatGPT threaten transparent science; here are our ground rules for their use*", Nature, 613 (7945), p. 612.

Oxford University Press. (2026). Policy on AI use & disclosure for OUP journals authors. [Online]. Available at: https://academic.oup.com/pages/for-authors/journals/preparing-and-submitting-your-manuscript/ai-principles-for-oxford-journals/policy-on-ai-use-and-disclosure-for-oxford-journals-authors [Accessed: 10 October 2026].

PLOS ONE. (n.d.). Ethical publishing practice: Artificial intelligence tools and technologies. [Online]. Available at: https://journals.plos.org/plosone/s/ethical-publishing-practice [Accessed: 17 October 2026].

PLOS. (n.d.). Research integrity and publication ethics: Use of AI. [Online]. Available at: https://plos.org/research-integrity-and-ethics/ [Accessed: 7 October 2026].

Research Evaluation. (n.d.). Author guidelines. Oxford Academic. [Online]. Available at: https://academic.oup.com/rev/pages/author-guidelines [Accessed: 17 October 2026].

SAGE. (n.d.). Artificial intelligence policy. [Online]. Available at: https://www.sagepub.com/journals/publication-ethics-policies/artificial-intelligence-policy [Accessed: 10 October 2026].

Springer Nature. (2026). AI use in manuscript preparation. [Online]. Available at: https://www.springernature.com/gp/policies/editorial-policies/ai-manuscript-preparation [Accessed: 10 October 2026].

Taylor & Francis. (n.d.). AI policy. [Online]. Available at: https://taylorandfrancis.com/our-policies/ai-policy/ [Accessed: 10 October 2026].

Thorp, H.H. (2023). "*ChatGPT is fun, but not an author*", Science, 379 (6630), p. 313.

Wiley. (2025). New Wiley guidelines give researchers clear path forward in responsible AI use [Press release, 28 October]. [Online]. Available at: https://newsroom.wiley.com/press-releases/press-release-details/2025/New-Wiley-Guidelines-Give-Researchers-Clear-Path-Forward-in-Responsible-AI-Use/default.aspx [Accessed: 12 October 2026].

Wiley. (n.d.). Using AI tools in your research. [Online]. Available at: https://www.wiley.com/en-ae/publishing/article/ai-guidelines [Accessed: 17 October 2026].

Appendix A. Policy Coding Key

Coding key: E = explicit; C = conditional or partial; N = not explicit in sufficiently reviewed official text; U = unverified because directly relevant primary-source evidence is incomplete. V denotes source-verification status: A = official policy page or editorial retrieved and read; B = relevant official excerpts reviewed but the complete text was not available; C = secondary summary only; U = no sufficiently verified journal-specific primary text. U must not be interpreted as absence of a requirement. C1–C12 follow the criteria in Table 2. The revised matrix contains 79 E, 24 C, 17 N, and 0 U cells and remains a single-coder exploratory analysis; no second-coder reliability statistic is reported.

Table (A1): Revised policy coding matrix (12 criteria × 10 policy units; observation date 7 October 2026; source-check revision 13 October 2026)

Policy unit	V	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12
Elsevier	A	E	E	E	E	C	E	E	E	E	E	E	N
Springer Nature	A	E	C	C	E	C	E	C	E	E	E	E	C
Wiley	A	E	E	E	C	C	E	E	E	E	E	E	N
Taylor & Francis	A	E	E	E	E	E	E	E	C	E	E	C	N
SAGE	A	E	E	N	C	N	E	E	E	C	C	C	E
Nature Methods (Nature Portfolio)	A	E	E	N	N	N	E	E	C	E	C	E	C
JAMA Network	A	E	E	E	E	E	E	E	E	E	E	C	E
PLOS ONE	A	C	E	E	C	N	E	E	E	E	E	E	C
BMJ journals	A	E	E	E	E	N	E	E	E	C	E	C	E
Research Evaluation (OUP)	A	E	E	N	N	N	N	E	E	N	C	N	N