

P2 — AI-Use Disclosure Standard

One institutional rule for what must be disclosed about AI use in research, what it must say, and where it goes for each kind of output, with statement templates a researcher can use unedited.

Browse the [full online instrument catalogue](#) to read or download every resource in HTML, Word, PDF, and spreadsheet formats.

About this instrument

Purpose	To give [INSTITUTION NAME] one defensible disclosure rule across theses and dissertations, journal articles, grant applications, conference papers, data and code releases, and research reports. Where this standard says <i>thesis</i> , it means the thesis or dissertation, whichever the institution calls it.
Who completes it	Nobody "completes" this form. Instead, [COMMITTEE] adopts it as an institutional standard, researchers, advisors, supervisors, and graduate students apply it, and the graduate school and research office check it at the points named in Section 4.
Time required	It takes about 45 minutes to read and mark local decisions, about 90 minutes of [COMMITTEE] time to adopt, and about three minutes for a researcher to write a disclosure once it is in force.
Related report sections	§1.2 (what "responsible" means: ethical, valid, reproducible, transparent), §2.1 (Dimension 1, human-in-the-loop discipline), §2.2 (Dimension 2, responsible use in practice, and the four AI-use modes), §2.3 (Dimension 3, tooling that promotes responsible use: auditability and reproducibility at a known model version), §2.4 (Dimension 4, AI-literate humans: competencies 1 and 2), §2.5 (Dimension 5, institutional benchmarking grid), §3.1 (funder positions), §3.3 (publisher convergence), Appendix C (the eleven AI tool classes), and Appendix G (labor-versus-judgment task taxonomy).
Cells of the 20-cell grid this instrument touches	D1-policy, D2-policy, D2-systems, D2-process, D4-people, D5-process.
Companion instruments	P1 (model institutional AI-in-research policy), where P2 serves as the disclosure clause of P1 written to stand alone. G1 (supervisor-student AI-use agreement), where G1 decides what is permitted, while P2 explains how it is recorded.
What it produces	An adopted standard, a disclosure statement on each covered output, and a retained log where AI touched an analysis. All three are inspectable by an examiner (or dissertation committee member, in US usage), editor, auditor, or funder.

Table 1. Overview of the policy instrument purpose, scope, and operational background.

Drafting note for [INSTITUTION NAME]. This standard is written to be adopted as-is. Four decisions must be made locally before it is issued: **(a)** The retention period and custodian for AI-use logs, in Clause 14. **(b)** Whether translation of a researcher's own text is disclosable and permitted, or prohibited outright, in Clause 6. Institutions in the sample split on this. **(c)** Which office receives disclosures that describe a prohibited use, in Clause 29 (which should not be the misconduct office by default). **(d)** Whether the thesis declaration is a positive statement of use, a two-option declaration, or an amendment to an existing sworn instrument, in Clause 17.

Section 1: What disclosure is for, and what it is not

Disclosure exists for two reasons.

The first is **reproducibility**. A reader has to be able to work out which parts of an output depend on a tool whose behavior is not fully repeatable and whose version will eventually be retired or withdrawn. Setting a model's temperature to zero does not make it deterministic. Floating-point non-associativity, batch-invariance failure, framework optimizations, and infrastructure-level variation mean that the exact same prompt with the exact same settings can return different output across runs. In addition, providers deprecate model versions on schedules ranging from months to years. This means an analysis reported against a general model name without a specific version number is, within a year or two, an analysis that nobody can re-run (§2.4, competency 2). If this information is not recorded, that irreproducibility remains invisible.

The second reason is **honest provenance**. A tool cannot bear responsibility for research work, so it cannot bear authorship of it. This principle represents the core of the no-AI-co-author rule on which every major publisher has converged (§1.1, §3.3). Disclosure is simply the positive form of that same claim: it states which parts a human generated, which parts a human verified, and who ultimately stands behind the result.

Disclosure is not a confession, and it is not a penalty. That sentence belongs in your institutional standard in those exact words. UK Research and Innovation states this point directly: applicants are expected to be transparent about generative-AI use, and "this information will not affect the assessment process." The way an institution frames disclosure directly influences how researchers behave. For example, the British Medical Journal moved AI disclosure into a structured submission-form field (the first publisher to do so), and recorded a 5.7 percent disclosure rate across 25,114 submissions to 49 BMJ journals in the seven months that followed. During that same period, AI assistance was becoming widely embedded in literature searching, drafting, and code development at scale (§1.1, §3.3). That gap will not close by making disclosure harder to write or more dangerous to file.

Disclosure is also not the whole of responsible use. Publishers converged on disclosure, but their standards "do not test whether the AI use was responsibly conducted" (§2.2). Responsible use means AI use that is ethical, valid, reproducible, and transparent, and no three of these properties can substitute for the fourth (§1.2). Disclosure delivers transparency and serves as a precondition for reproducibility. Verification is what makes a use valid, and a disclosure records that obligation without replacing it.

Section 2: The disclosure threshold

An undefined threshold produces two failures at once: researchers who disclose nothing because they cannot tell what counts, and researchers who disclose everything, which buries the disclosures that matter.

The rule, in one sentence.

Disclose any AI use that a competent reader would need to know about to judge the work, reproduce it, or rely on it.

Clauses 1 to 8 make that sentence operational. They rest on the labor-versus-judgment demarcation that the framework requires institutions to publish (§2.1, Appendix G). In this framework, *labor* refers to work an AI tool can handle where an error is visible and easily corrected. In contrast, *judgment* refers to work where substituting AI silently degrades the quality and integrity of the research product.

Clause 1: The three-part test. A use is disclosable if the answer to any of the following is yes.

Test	Question
1a <i>Trace</i>	Did any AI output survive into the finished artifact: text, code, a figure, a transcript, a screening decision, a reference, or a data value?
1b <i>Judgment</i>	Did an AI tool touch a task on the judgment side of the institution's labor-versus-judgment demarcation, even as a scaffold rather than a primary decision-maker?
1c <i>Reproducibility</i>	Would another researcher, given the same inputs, need to know the specific tool and its version to arrive within a stated tolerance of the same output?

Table 2. Evaluation criteria and diagnostic questions for the three-part disclosure threshold test.

If all three answers are no, no disclosure is required. This is the **no-trace exemption**, stated positively in Clause 6 so that researchers can rely on it with confidence.

Evidence this clause is operating: The published standard contains the three-part test, and a sample of ten disclosures shows researchers applying it rather than disclosing indiscriminately.

Clause 2: Trace triggers a disclosure. Except for the uses exempted under Clause 6, where AI output survives into the artifact, the use must be disclosed. The test is survival into the final record, not whether the author believes it was significant. Significance is a judgment that the reader is entitled to make.

Evidence: Disclosure statements on a sample of recent outputs name the specific sections in which AI output survives.

Clause 3: Judgment triggers a disclosure even where nothing survives. If an AI tool scaffolded a judgment task (such as generating candidate interpretations of an unexpected result, proposing analytic options, running a devil's-advocate pass, or suggesting inclusion criteria), the use is disclosed even if none of the tool's actual words appear in the finished artifact. The reader assessing that judgment is entitled to know that a tool was used in the process. This requirement is what distinguishes a research disclosure standard from a simple plagiarism rule.

Evidence: The standard explicitly names validator-mode use as disclosable, and a sample of theses using AI adversarial review carries a corresponding statement.

Clause 4: Reproducibility triggers the extended content set. Where the AI use touched the analysis rather than the writing (such as classification, coding, extraction, screening, statistical work, or code that produced a reported result), the disclosure must include the four additional fields set out in Clause 13.

Evidence: Outputs reporting AI-mediated analysis carry the model version, parameters, prompt deposit, and a re-run stability result.

Clause 5: The four AI-use modes, applied. The framework's four modes (**search, co-author, validator, tutor**) sit at different points on the disclosure threshold.

Mode	Default disclosure position	Why
<i>Search</i>	Disclose when AI-surfaced sources are cited in the artifact. Do not disclose when the tool only reordered search results that the researcher then found and read independently.	Studies have measured citation fabrication at 18 to 55 percent, along with a collective error rate above 60 percent across eight consumer AI search engines on 1,600 queries. This makes the provenance of any citation list a critical fact for readers (§2.2, §2.4, and Appendix C, Class 1).
<i>Co-author</i>	Disclose when AI-generated or AI-substantially-altered text survives into the artifact. Do not disclose spelling and grammar correction of the researcher's own sentences (Clause 6).	This is where publisher rules apply directly, and where the trace test is easiest to apply.
<i>Validator</i>	Disclose. Always.	This is inherently judgment-adjacent work, and it remains completely invisible in the finished text unless disclosed (Clause 3).
<i>Tutor</i>	Do not disclose in research outputs. Tutor use is a learning activity, not a research method, and belongs in the advising or supervision record.	Disclosing background learning generates unnecessary volume without informing the reader.

Table 3. Default disclosure positions and rationales across the four artificial intelligence use modes.

Evidence: The standard reproduces this table, and advisors and supervisors can explain these four positions without needing to consult the document.

Clause 6: What need not be disclosed (the no-trace exemption). The following uses are exempt from disclosure:

- Spelling, grammar, and punctuation correction applied to sentences the researcher wrote.
- Predictive text, autocomplete, and keyboard suggestions.
- Reference-style formatting and bibliography generation from records the researcher selected.
- Typesetting, layout, and the accessibility features of ordinary software.
- Search-engine ranking, database relevance ordering, and reference-manager deduplication.
- Plagiarism-similarity checking run by the institution as part of its own administrative process.
- Tutor-mode use for the researcher's own personal learning, where no output enters the artifact.

Publishing and academic sectors support each of these exemptions. The Association for Computing Machinery treats "basic word processing systems" as an exception to its disclosure requirement. The Institute of Electrical and Electronics Engineers places editing and grammar enhancement outside the intent of its policy. Springer Nature states that AI-assisted copy-editing does not need to be declared. *Responsible AI in Academic Research: A Competency Framework for Research Training* found that SAGE is the only publisher in its sample with an explicit assistive-versus-generative threshold, exempting tools "that make suggestions to improve or enhance your own work" (§3.3).

Local decision. Translation of a researcher's own text between languages is treated in Appendix G as labor with an audit trail and a researcher signature. However, at least one institution in the report's 38-university sample prohibits AI translation of text drafted in another language outright. [INSTITUTION NAME] must choose between disclosable and permitted or prohibited. Do not leave this unstated.

Evidence: The exemption list appears in the published standard, and a researcher survey shows the list is widely known.

Clause 7: A disclosure obligation is not a permission. Where an institutional policy, funder condition, publisher rule, or ethics approval prohibits a specific use, that prohibition governs. In those cases, a disclosure serves as evidence of a policy breach rather than a cure for it. The clearest example is AI-generated data figures, which nine major publishers prohibit as a category (§3.3, Appendix C, Class 9).

Evidence: The standard clearly states this rule, and the institution can name where a disclosed-but-prohibited use would be routed (Clause 29).

Clause 8: The standard binds everyone who produces research. This standard applies on identical terms to graduate students, research staff, advisors, supervisors, and professional staff producing research outputs. If an institution applies disclosure rules only to students, it is simply running a plagiarism policy under a research-integrity label. Across the 38 top-tier doctoral universities in 15 countries and jurisdictions surveyed for the report, around 40 percent of institutional AI policies stop at that ceiling (§1.3).

Evidence: The scope clause of the published standard explicitly names both staff and students, and at least one staff-authored output carries a disclosure.

Section 3: What a disclosure must contain

Clause 9: The five required fields. Every disclosure must include these five fields, in any order and in any format the target destination allows.

#	Field	What it means	Why a reader needs it
1	Tool class	Which of the eleven AI tool classes in Appendix C was used: search, co-author, validator, tutor, coder, retrieval-augmented assistant, literature-synthesis, specialized research assistant, image generation, voice and transcription, or AI detection.	It identifies the specific failure mode to look for. For example, a transcription tool tends to fail by inventing content, while a retrieval-augmented assistant tends to misattribute a real source.
2	Model and version	The named model and its version or date-stamped endpoint. Where the provider does not expose a version number, state that clearly and give the date of access.	Model-version drift makes a disclosure without a version number unreproducible within a year.
3	Date or date range of use	When the tool was used, to the specific day or across a date range.	It pins the disclosure to the specific model state that existed at that time.
4	What it was used for	The task, described in the researcher's own words, tied directly to the part of the artifact affected.	This is the field the reader actually relies on to evaluate the work.

#	Field	What it means	Why a reader needs it
5	What verification was performed	The specific check that was performed, who performed it, and its result.	Verification is the actual obligation, and the disclosure is simply its record.

Table 4. Definitions and reader rationales for the five required disclosure fields.

Fields 1 to 4 match what the more operationally developed institutions and publishers already ask for. The University of Helsinki requires "the name of the tool and the version used, the date of use, how the tool was used, and for what step or purpose." TU Delft requires name, version, model, and source. Taylor & Francis requires tool name, version, purpose, and manner of use. Oxford University Press requires the location of each use, the tool and version, the manner of use, the justification, and the steps taken to validate the outputs. The Public Library of Science requires the tools used, how the authors evaluated the validity of the outputs, and which aspects of the article were affected (§3.3).

Field 5 is the one that most publisher templates only imply. It is mandatory here because Dimension 1 makes human verification an institutional commitment, and the first competency under Dimension 4 (citation verification) makes it the researcher's direct responsibility (§2.1, §2.4). A disclosure that describes what a tool did without stating what a human checked tells only half the story.

Evidence: A sample of ten disclosures scored against the five fields, and the institution can state the completion rate per field.

Clause 10: The minimal form. Where the target destination permits a single sentence, the five fields compress into one (Section 5). The minimal form is not an inferior option. It is the appropriate form for AI use that leaves a trace in the text but has no analytical consequence.

Clause 11: The full form. Where the AI use touched more than one part of the artifact, or where the destination requires a dedicated section, the disclosure is written for each part separately rather than as a single global sentence (see the bad example in Section 8).

Clause 12: Name a human for each verification. In multi-author work, field 5 must name which specific author performed the verification. While collective responsibility is the default for any shared output, the verification check itself was performed by an individual.

Evidence: Multi-author disclosures name an individual verifier for each task.

Clause 13: The reproducibility set (four additional fields), required where Clause 4 applies.

#	Field	Requirement
6	Prompt or prompt set	Provided verbatim in an appendix, or deposited in a named repository with a resolvable identifier. Where a prompt contains personal or confidential data, deposit it under controlled access and disclose the location, never the private content itself (Clause 30).
7	Parameters	Temperature, random seed where the provider exposes one, and any retrieval corpus or system configuration. Where a parameter is not exposed, state that explicitly.
8	Re-run stability check	The number of re-runs performed, and the observed agreement between them. Report the result even when agreement is poor, especially when it is poor.
9	Deposit location	Where the prompts, outputs, logs, and code live, and how long they will remain accessible there.

Table 5. Requirements for the four additional fields in the reproducibility disclosure set.

Field 8 is not just decorative paperwork. AI models are rarely deterministic at the raw output level, and they are rarely fully stable across even five re-runs at identical settings. Prompt choice is a significant researcher degree of freedom. A systematic evaluation of 13 million model-generated labels across 18 models found that prompt and model choices alone produced incorrect conclusions in 31 to 50 percent of cases, even when using highly capable models (§2.4, competency 3). A disclosure of an AI-mediated analysis that omits a stability check leaves out that analysis's largest known source of uncertainty.

Clause 14: Retention. The institution retains the material behind field 9 for [RETENTION PERIOD], under the custody of [CUSTODIAN]. This requirement should be established as a clear rule in the standard rather than decided case by case. Auditability through session logs is one of the six tool properties that Dimension 3 includes in the procurement test (§2.3). Having a clear retention rule turns that technical property into a practical institutional capability.

Local decision. Set [RETENTION PERIOD] to match the institution's existing research-data retention schedule rather than inventing a new one. Align [CUSTODIAN] with the research-data-management function, not with the academic integrity office.

Evidence: The records schedule names AI-use logs with a designated retention period and a custodian.

Section 4: Where the disclosure goes, by artifact type

Clause 15: The destination's rule on location governs, while this standard's rule on content serves as the baseline. Where a publisher, funder, or examination regulation specifies where the disclosure must sit, that placement rule governs. If the destination asks for less content than Clause 9, the researcher supplies all five fields anyway. If it asks for more, the researcher supplies the additional details.

Clause 16: The map.

Artifact	Where the disclosure goes	Who checks it, and when
Doctoral thesis	The thesis declaration page, plus a statement in the methods or an appendix wherever AI touched the analysis.	Advisor or supervisor at submission, examiners during examination, and the graduate school at final submission.
Journal article	The destination journal's mandated location (see Clause 19).	Corresponding author before submission, and editor at triage.
Grant application	The funder's specified field or section. Where the funder specifies none, the disclosure goes in the institution's internal pre-award record, not in the application text.	Research office at internal sign-off.
Conference paper	The venue's policy where it has one. Otherwise, use the journal-article form.	Corresponding author, as well as the program chair if the venue requests it.
Dataset or code release	The README, plus a machine-readable field in the metadata record, plus the prompt set and model version deposited with the artifact.	Data steward at deposit.

Artifact	Where the disclosure goes	Who checks it, and when
Research report or commissioned study	A named disclosure section, at the front or immediately before the references.	Report owner at sign-off, and the commissioning body as specified by contract.
Preprint	A statement in the methods or an equivalent section. No preprint server in the report's sample mandates a specific section name, and none permits AI as an author.	Corresponding author.
Ethics application	The application's own AI field where one exists. Otherwise, use a paragraph in the methods description.	Ethics committee.
Peer-review report	Follow the journal's reviewer policy, which is more restrictive than its author policy at most publishers. Never upload a manuscript under review to a tool without the publisher's explicit permission.	Handling editor.
Examiner report	Not applicable. AI may not participate as an examiner, and examination materials remain outside the AI-tool workflow under P1 Clause 11 and Appendix G.	Graduate school.

Table 6. Disclosure placement locations and review responsibilities across different research artifact types.

Evidence: The institution can produce, for each row, the form or template that carries the disclosure and the office that checks it.

Clause 17: The thesis declaration is the anchor instrument. Where an institution already has a sworn or signed declaration at thesis submission, AI disclosure should be added directly to it rather than created as a separate instrument. Germany's *eidesstattliche Versicherung* tradition is the clearest working example. The affidavit long predated generative AI, and disclosure was added directly to it with established legal standing. Institutions that require doctoral students to state "when, where, why, and to what extent" AI was used inside that existing affidavit have the strongest instrument in the report's global sample, at no additional administrative cost (§1.3, §3.2).

Local decision. Choose one of three forms: **(i)** an amendment to an existing sworn declaration (the strongest approach where such traditions exist), **(ii)** a two-option declaration where the doctoral student selects either no AI use or acknowledged AI use, with the acknowledgment providing the five required fields (currently in use at one Class D institution in the report's sample), or **(iii)** a positive statement of use appended to the thesis preliminary pages and prompted in the online submission form (used at institutions in Australia and Canada). Do not create a separate standalone form when an existing instrument can easily carry the clause.

Evidence: The thesis declaration text, and the completion rate at final submission for the most recent [DATE] cohort.

Clause 18: The doctoral student must be able to defend the disclosure orally. Where an institution examines by oral examination (the viva or thesis defense), the doctoral student is expected to describe and defend any AI use, and examiners may ask about it. This is already the rule at institutions in the report's sample, and it is what gives practical force to the written

disclosure. Examiners are not expected to act as forensic detectors of AI use, but they are fully entitled to ask candidates to explain and defend whatever was declared.

Evidence: The examiner briefing note includes the question, and examiner reports from the most recent cycle show it being asked.

Section 5: Statement templates

Twelve templates follow: six artifact types, each presented in a **minimal** one-sentence form and a **full** form. Both comply with this standard, and [INSTITUTION NAME] can choose which format serves as its baseline requirement (see the drafting note at the end of this section).

These templates were written for this pack to satisfy the content requirements publishers and institutions state. They are not copied or adapted from any publisher's sample text. If a target venue mandates its own specific wording, researchers should use that wording and simply add whichever of the five required fields in this standard it leaves out.

Square brackets indicate fill-in fields. Delete any sentence that does not apply, but do not delete field 5 (verification).

5.1 Doctoral thesis

Minimal

Declaration of AI use. *In preparing this thesis I used [MODEL AND VERSION], a [TOOL CLASS] tool, between [DATE] and [DATE], to [TASK, e.g. convert my own bullet-point notes into draft paragraphs in Chapters 3 and 5]. I checked every resulting passage against my own notes and sources before it entered the thesis. Every claim, interpretation, and conclusion in this thesis is mine, and I am responsible for all of them.*

Full

Declaration of AI use. This declaration covers all use of artificial-intelligence tools in the research and writing of this thesis, in accordance with the [INSTITUTION NAME] AI-Use Disclosure Standard.

Literature searching. I used [MODEL AND VERSION], a retrieval-augmented research assistant, between [DATE] and [DATE], to scope the literature in [FIELD] and to assemble a candidate reading list. I resolved every citation that survived into Chapter 2 to its primary source and read each source in full. [NUMBER] candidate references returned by the tool did not resolve to a real source and were discarded.

Drafting. I used [MODEL AND VERSION], a large-language-model co-author tool, between [DATE] and [DATE], to convert my own outlines and notes into draft prose in Chapters [N] and [N], and to shorten passages I had already written. No section of this thesis contains substantive content I could not have produced and cannot independently defend. I revised all AI-assisted text.

Analysis. I used [MODEL AND VERSION] to [ANALYTICAL TASK, e.g. apply a coding frame to 412 interview excerpts] in Chapter [N], at temperature [VALUE] and seed [VALUE or "not exposed by the provider"]. The prompt set is reproduced at Appendix [X] and deposited at [REPOSITORY, IDENTIFIER]. I re-ran the task [NUMBER] times, and agreement between runs was [RESULT]. I hand-coded a [PERCENTAGE] random sample against the model output, where agreement was [RESULT]. All disagreements were resolved by me and are documented at [LOCATION].

Adversarial review. I used [MODEL AND VERSION], from a different model family than the tools above, to run a devil's-advocate pass on the argument in Chapter [N]. I treated its output as a source of questions, not as validation. [NUMBER] of its objections led to changes in the chapter, and those changes are described at [LOCATION].

Verification and responsibility. [NAME] performed the verification described above. The research questions, the design, the interpretation of results, and the conclusions are my own work. No AI tool is an author of this thesis, and no AI tool made any research judgment recorded in it.

Signed: [NAME], [DATE].

5.2 Journal article

Minimal

Declaration of AI use. The authors used [MODEL AND VERSION], a [TOOL CLASS] tool, on [DATE], to [TASK], in the [SECTION] section. [AUTHOR INITIALS] reviewed and verified all resulting content. The authors are responsible for the content of this article. No AI tool is an author.

Full

Declaration of AI use. In preparing this manuscript the authors used the tools listed below. [AUTHOR INITIALS] verified all AI-assisted content before submission, and the authors take responsibility for the whole of the article, including its references.

1. [MODEL AND VERSION] ([TOOL CLASS]), used on [DATE] to [TASK] in [SECTION].
Verification: [CHECK PERFORMED, BY WHOM, RESULT].
2. [MODEL AND VERSION] ([TOOL CLASS]), used between [DATE] and [DATE] to [TASK] in [SECTION]. Verification: [CHECK PERFORMED, BY WHOM, RESULT].

Where AI tools contributed to the analysis, the model version, parameters, prompt set, and re-run stability results are reported in [SECTION / SUPPLEMENT] and deposited at [REPOSITORY, IDENTIFIER].

No AI tool is listed as an author. No AI tool generated or altered any figure, image, or data value in this article. No part of this manuscript was uploaded to a tool without an enterprise agreement disabling training on input.

5.3 Grant application

Minimal

AI-use statement. [MODEL AND VERSION], a [TOOL CLASS] tool, was used on [DATE] to [TASK, e.g. shorten the applicant-written case for support to the word limit]. The applicant wrote and is responsible for the content of this application.

Full

AI-use statement. The applicant used the following tools in preparing this application: [MODEL AND VERSION] ([TOOL CLASS]) on [DATE] to [TASK], and [MODEL AND VERSION] ([TOOL CLASS]) between [DATE] and [DATE] to [TASK]. Every claim, citation, figure, and costing in this application was written or checked by the applicant. No section of this application was substantially developed by an AI tool. [Where required by the funder: the sources surfaced with AI assistance and cited in this application are marked in the reference list.]

Drafting note. Check the funder's rule at the point of application, every time (Section 7). Where the funder requires nothing, keep the statement in the internal pre-award record and out of the application text. Unrequested text in a page-limited application is an unnecessary cost with no benefit.

5.4 Conference paper

Minimal

AI-use disclosure. [MODEL AND VERSION], a [TOOL CLASS] tool, was used on [DATE] to [TASK] in [SECTION]. [AUTHOR INITIALS] verified the output. The authors are responsible for the content of this paper and no AI tool is an author.

Full

AI-use disclosure. The authors used [MODEL AND VERSION] ([TOOL CLASS]) between [DATE] and [DATE] for the following tasks, identified by section: [SECTION]: [TASK], and [SECTION]: [TASK]. Verification: [CHECK, BY WHOM, RESULT] for each. Where AI tools were used to generate or transform code, the generated code is marked in the repository at [IDENTIFIER] and was reviewed and tested by [NAME] before inclusion. Where AI tools were used in the analysis, model version, parameters, prompts, and re-run stability results are reported at [LOCATION]. No AI tool is an author of this paper. Any AI use in reviewing other submissions to this venue is governed by the venue's reviewer policy and is not covered by this statement.

5.5 Dataset or code release

Minimal

AI use in this release. [MODEL AND VERSION], a [TOOL CLASS] tool, was used on [DATE] to [TASK, e.g. generate the cleaning script at /src/clean.py]. All generated code was reviewed and tested by [NAME] before release. Generated segments are marked in the source.

Full

AI use in this release.

Scope. This statement covers AI use in producing the data, the code, and the documentation in this release.

Tools. [MODEL AND VERSION] ([TOOL CLASS]), [DATE] to [DATE], used to [TASK]. [MODEL AND VERSION] ([TOOL CLASS]), [DATE] to [DATE], used to [TASK].

Where the AI output is. AI-generated or AI-modified code is marked in the source with `# AI-assisted: [MODEL AND VERSION], [DATE]` at the head of each affected block. AI-derived data values, where any exist, are flagged in the column [COLUMN NAME] of [FILE]. No data value in this release was generated by an AI tool other than as recorded there.

Reproducing the AI steps. Prompts: [PATH]. Parameters: temperature [VALUE], seed [VALUE or "not exposed"], retrieval corpus [DESCRIPTION or "none"]. Re-run stability: [NUMBER] runs, agreement [RESULT]. Users should expect approximate rather than exact reproduction of these steps: the model is not fully deterministic at these settings, and [MODEL AND VERSION] may be withdrawn by its provider.

Verification. [NAME] reviewed all generated code, ran the test suite at [PATH], and confirmed [RESULT]. Machine-readable field: `ai_use: true` in [METADATA FILE].

5.6 Research report

Minimal

AI use in this report. [MODEL AND VERSION], a [TOOL CLASS] tool, was used on [DATE] to [TASK]. [NAME] verified the output. [ORGANIZATION] is responsible for every factual claim, calculation, and recommendation in this report.

Full

AI use in this report. This report was researched and written by [NAMES]. AI tools were used as set out below, and the authors are responsible for every factual claim, source interpretation, calculation, caveat, and recommendation.

Tool class	Model and version	Dates	Used for	Verified by, and how
[CLASS]	[MODEL AND VERSION]	[DATE] to [DATE]	[TASK]	[NAME], [CHECK], [RESULT]
[CLASS]	[MODEL AND VERSION]	[DATE] to [DATE]	[TASK]	[NAME], [CHECK], [RESULT]

Every source cited in this report was retrieved and read by a named author. No figure, chart, or image in this report was generated by an AI tool. Where AI tools contributed to analysis, the prompts, parameters, and stability results are at [LOCATION]. Where this report was prepared for [COMMISSIONING BODY], this statement satisfies clause [N] of the commissioning agreement.

Drafting note on setting the institutional floor. Set the **minimal** form as the floor where AI touched only labor tasks, and the **full** form as mandatory where AI touched analysis or judgment. If you mandate the full form for every minor use, you run directly into the noise problem described in Section 9. In short, a standard that takes too long to complete is a standard that nobody will follow.

Section 6: How this interacts with publisher policies

Clause 19: The destination's location rule governs, and the researcher checks it at submission. Publishers converged on substance and diverged on placement. For instance, Elsevier requires a section entitled "Declaration of Generative AI and AI-assisted technologies in the writing process" immediately above the references, while Cell Press requires the equivalent statement after the declaration of interests. Nature and Springer Nature place it in the Methods. Science requires it in three places: the cover letter, the acknowledgments, and the methods or supplementary materials. JAMA requires the Acknowledgment or Methods plus an answer to a submission-form question. The Association for Computing Machinery and the Institute of Electrical and Electronics Engineers place it in the Acknowledgments, with the Institute additionally requiring section-by-section identification. The Public Library of Science requires a dedicated section in the Methods. Wiley requires a cover-letter statement plus a dedicated appendix describing the AI-generated content. The British Medical Journal requires a contributor-section acknowledgment plus a fuller methods description where the use was in the research itself, and moved the question into a structured submission-form field (§3.3).

The instruction to researchers is simple: **read the destination's author instructions at submission, put the disclosure where they say, and include all five required fields whatever they ask for.**

Evidence: The institution's submission-support guidance carries the location table, and the research office can produce it.

Clause 20: Where the publisher's threshold is lower than this standard, this standard governs, and where it is higher, the publisher governs. SAGE exempts assistive tools that improve the author's own work, Springer Nature exempts AI-assisted copy-editing, and the Association for Computing Machinery treats basic word processing systems as an exception. All three are consistent with Clause 6. Where a publisher requires disclosure of something Clause 6 exempts, the publisher's rule applies to that submission.

Clause 21: Three publisher rules are not disclosure rules at all, and disclosure does not satisfy them.

- **AI cannot be an author.** This rule is universal across the eighteen publishers and journal families in the report's sample, as well as across all three preprint servers. It is a fundamental rule about accountability, and no disclosure statement can cure a breach of it.
- **AI-generated data figures are prohibited at nine major publishers,** outside a narrow exception where the AI system itself is the object of study. Disclosing an AI-generated microscopy image does not make it publishable.
- **Reviewers may not upload manuscripts to public AI tools.** Sixteen of the eighteen publishers state this directly, framed as confidentiality rather than integrity, which is why it holds regardless of which national research-integrity code applies. The Public Library of Science goes further and requires a reviewer who used AI to disclose that use to the authors, not only to the editor.

Evidence: The institution's reviewer guidance states the upload prohibition, and ethics-of-review training records show it is taught.

A note on sector speed, for [COMMITTEE]. Every major publisher issued a substantively identical no-AI-co-author position within about three months of ChatGPT-3.5's broad public release, compared to the years comparable cross-publisher convergences have historically taken (§1.1, §3.3). Academic publishing can converge quickly when it decides to, which is why a university can comfortably adopt a disclosure standard now rather than waiting for one to be handed down from outside.

Section 7: How this interacts with funder requirements

Clause 22: Check the funder's rule at the point of application, every time. Across the fourteen national research funders surveyed for the report (spanning eleven countries plus the European Research Council as a supranational entity), the applicant-side position runs from prohibition-adjacent to explicit non-requirement, across policies dated September 2023 to April 2026 (§3.1). Any general rule a researcher tries to carry in their head will inevitably be wrong for some funder.

Funder position	Funders in the sample	What the researcher does
Disclosure or acknowledgment expected	UK Research and Innovation and its component councils, the German Research Foundation (which alone extends the duty to published results), the Netherlands Organization for Scientific Research (in the reference list), the Australian Research Council and the National Health and Medical Research Council, the Canadian tri-agency (conditional acknowledgment), and the European Research Council	Use the Section 5.3 statement in the funder's specified field. Where the expectation is stated in principle rather than as a form field, put it in the section the funder names for methods or approach.

Funder position	Funders in the sample	What the researcher does
Disclosure encouraged	The US National Science Foundation	Use the minimal statement in the project description.
A substantive originality bar rather than a disclosure rule	The US National Institutes of Health: applications substantially developed by AI "are not considered the original ideas of applicants" and will not be considered	Disclosure is not the operative question. The applicant must be able to show the intellectual content is theirs.
No disclosure requirement	The Swiss National Science Foundation, the Swedish Research Council ("you do not need to state whether you have used AI"), the Japan Society for the Promotion of Science, and Singapore's Agency for Science, Technology and Research	Keep the statement in the internal pre-award record. Do not add it to the application.
No funder-specific policy located in the report's sample	The Research Council of Norway	Treat the position as unstated: keep the statement in the internal pre-award record and check the call conditions before each submission.

Table 7. Sample funding agency disclosure positions and corresponding compliance actions for researchers.

Evidence: The research office maintains a funder-by-funder disclosure table with a review date, and it can produce the current version on request.

Clause 23: Almost no funder regulates AI use *inside* the funded research, so the institution must.

Funders regulate the front and back of the grant lifecycle (application and assessment) and delegate the middle to host-institution research-integrity policy. The German Research Foundation is the single clear exception in the report's sample, as its disclosure duty extends to making results publicly available rather than stopping at the grant application. This regulatory gap is exactly why this standard must operate at the institutional level rather than relying on funding bodies to set the rules.

Evidence: The institution's AI policy explicitly claims the funded-research surface, and the funder-delegation language is cited directly within it.

Clause 24: The assessment side is nearly universal and nearly absolute. Reviewers are prohibited from uploading proposals to external generative-AI tools at thirteen of the fourteen funders surveyed, framed as confidentiality and intellectual property rather than integrity. Staff who assess for funders, sit on panels, or review for journals must be informed of this separately from the author-side standard, because it binds them in an evaluative role that this standard does not otherwise cover.

Evidence: The panel-member and peer-review briefing carries the prohibition, and the institution can name the last date it was issued.

Section 8: A good disclosure and three bad examples

A bad disclosure: The empty one

"AI was used in the preparation of this manuscript."

Why it fails. It includes none of the five required fields, so a skeptical auditor cannot confirm or refute a single thing in it. There is no failure mode to look for, nothing to re-run, no model state to pin it to, no affected section to inspect, and no obligation recorded as discharged. This is what happens when a rule asks only whether AI was used, and it is the exact format that teaches editors to stop reading disclosure statements altogether.

A bad disclosure: The flood

"The authors used AI assistance throughout the preparation of this manuscript, including grammar checking in Microsoft Word, spelling correction, predictive text, reference formatting in [manager], search-result ranking in [database], autocomplete in the code editor, transcription of two author meetings, language polishing of the abstract, introduction, methods, results and discussion, image cropping..." (continues for 340 words)

Why it fails. Almost everything in it is exempt under Clause 6, and the one item that is not (language polishing that produced surviving text) is buried where no reader will find it. Disclosing too much is not a safe, harmless error. It simply dumps the sorting burden onto the reader, who knows far less about what happened than the researcher does.

A bad disclosure: The one that discloses a prohibition

"Figures 2 and 3 were generated using [image model] from descriptions of our experimental results."

Why it fails. Nine major publishers prohibit AI-generated data figures as a category. The statement is complete, honest, and well-written, but it is still fatal to the submission because disclosure is not permission (Clause 7). The only correct action here was not to generate the figures with AI in the first place.

A good disclosure

Declaration of AI use. *We used [model], version [dated endpoint], a retrieval-augmented research assistant, between 4 and 19 February [YEAR], to scope the literature on [TOPIC] and build a candidate reference list. Every reference cited in Section 2 was retrieved and read by SK at the publisher of record. Eleven candidate references returned by the tool did not resolve to a real source and were discarded.*

We used [model], version [dated endpoint], a large-language-model co-author tool, on 3 and 4 March [YEAR], to compress the Discussion from 1,900 to 1,200 words. The compressed text was written from our own draft. MR checked every retained sentence against the original for changed meaning, and restored four passages the tool had altered.

We used [model], version [dated endpoint], to apply our published coding frame to 412 free-text survey responses (Section 4.3), at temperature 0, seed not exposed by the provider. The prompt set is at Supplement S2 and deposited at [repository, DOI]. We ran the classification five times: four responses out of 412 changed class between runs, and all four are listed in Supplement S3. SK independently hand-coded a 20 percent random sample, where agreement with the model was 94 percent. All 12 disagreements were adjudicated by SK and MR together, and were resolved in favor of the human hand coding.

No AI tool is an author. No AI tool generated or altered any figure or data value. No unpublished participant data was entered into any tool outside our institutional enterprise tenancy with training on input disabled.

Why it works. All five fields appear for each use, and the reproducibility set appears for the use that touched the analysis. The verification is specific, attributed to a named person, and reports exact numbers rather than vague assurances, including numbers that do not flatter the tool. Eleven unresolvable references and four unstable classifications are exactly the facts a reader needs to know, and exactly the facts an anxious researcher might be tempted to omit. The whole statement is under 300 words, took only a few minutes to write because logs were kept during the research, and allows an auditor to check every claim against the deposit.

Section 9: What not to require

An overly broad rule does not produce better information. It produces noise, non-compliance, and usually both. Each of the following points binds the institution, not the researcher.

Clause 25: Do not require disclosure of no-trace uses. Clause 6 exists specifically for this reason. A rule that captures spellchecking captures everything, and a rule that captures everything quickly gets ignored.

Clause 26: Do not require prompt logs for every use. Require them only where the AI output entered the analysis (Clause 4). The Canadian tri-agency revised its guidance with the stated aim of increasing clarity and lowering researchers' administrative burden while acknowledging the increased complexity of monitoring and enforcing detailed AI disclosures. This was a major funder conceding publicly that strict disclosure rules are much harder to operationalize than they appear at the design stage.

Clause 27: Do not require anything the institution cannot check. An unverifiable requirement teaches researchers that the standard is merely decorative, and that cynicism quickly generalizes

to the requirements that *are* checked. Before adopting any clause, [COMMITTEE] should clearly identify the artifact that proves compliance and the office responsible for reviewing it. If neither exists, the clause is not ready to be adopted.

Clause 28: Do not use AI-detection tools as evidence, and do not require researchers to run them. Seven detectors evaluated together showed 39.5 percent baseline accuracy falling to 17.4 percent under simple adversarial techniques, with documented bias against authors who are not native English speakers (Appendix C, Class 11). At least one Class D institution in the report's sample explicitly prohibits examiners from running a thesis through AI-detection software. Detection output should, at most, trigger a constructive conversation. It can never serve as proof of misconduct, because it treats surface authorship signals as a proxy for research validity, which they are not.

Clause 29: Do not route disclosures to the misconduct office by default. A disclosure is a routine research record. It belongs with the advisor, supervisor, editor, research office, or data steward, depending on the artifact. Only a disclosure describing an explicitly prohibited use should go anywhere else, and [INSTITUTION NAME] must define that escalation route clearly so that researchers can see the normal process is not punitive. Where the institution can state, as UK Research and Innovation does for grant applications, that disclosure will not affect assessment, it should say so plainly in the standard.

Clause 30: Do not require the contents of prompts that contain personal or confidential data. Require a pointer to a controlled-access deposit instead. A policy that pushes participant data into a public appendix in the name of transparency creates an immediate privacy breach in the name of research integrity. Data minimization applies to disclosure just as strictly as it applies to the research itself.

And one thing not to abolish. A negative attestation (a statement confirming that no AI was used) works well when it is attached to an instrument that already exists and carries legal weight, such as a doctoral affidavit or thesis declaration. It does not work as a new standalone form attached to every routine output, where it generates paperwork volume while proving nothing.

Section 10: Adoption checklist

For [COMMITTEE] at the point of adoption, and for annual review thereafter. Answers are *Yes*, *Partial*, or *No*, and each carries the artifact that would prove it.

A1. Has [INSTITUTION NAME] published a disclosure threshold that states both what must be disclosed and what need not be?

☐ Yes ☐ Partial ☐ No

Evidence: Name the document and clause, provide the URL, and state both the last review date and the next scheduled review.

A2. Does the published standard require all five content fields: tool class, model and version, date, use, verification?

☐ Yes ☐ Partial ☐ No

Evidence: Quote the clause. A standard that omits the verification field scores a "No."

A3. Does the standard require the reproducibility set (prompts, parameters, re-run stability, deposit) where AI touched an analysis?

☐ Yes ☐ Partial ☐ No

Evidence: Quote the clause, and produce one research output from the last twelve months that includes it.

A4. Does the thesis declaration or equivalent submission instrument carry an AI clause?

☐ Yes ☐ Partial ☐ No

Evidence: The current declaration text, and the completion rate at final submission for the most recent cohort.

A5. Does the standard bind staff as well as students, on the same terms?

☐ Yes ☐ Partial ☐ No

Evidence: The scope clause, plus at least one staff-authored output carrying a disclosure.

A6. Is there a named owner and a stated retention period for AI-use logs?

☐ Yes ☐ Partial ☐ No

Evidence: The records-retention schedule entry naming AI-use logs, the period, and the custodian.

A7. Does the institution maintain a current funder-by-funder disclosure table for the funders its researchers apply to?

☐ Yes ☐ Partial ☐ No

Evidence: The table, with its review date. A table older than twelve months scores Partial.

A8. Is the reviewer-side prohibition on uploading manuscripts and proposals to public AI tools issued separately to staff who review and assess?

☐ Yes ☐ Partial ☐ No

Evidence: The briefing document and the date of its last issue.

A9. Are examiners briefed that they may ask about a declared AI use and may not run theses through detection software?

☐ Yes ☐ Partial ☐ No

Evidence: The examiner briefing note, along with examiner reports from the last cycle demonstrating that the question was asked.

A10. Does the institution measure its own disclosure completion rate, and report it to [COMMITTEE]?

☐ Yes ☐ Partial ☐ No

Evidence: The most recent report, with the denominator stated. This is the item that separates established from leading on D2-process.

Crosswalk to the Responsible AI framework

Section or clause of this instrument	Report section
About this instrument (cells touched)	§2.5 (the four axes), and Appendix A (D1 and D2 rows)
Section 1: What disclosure is for	§1.1 (the use-versus-disclosure gap and the BMJ rate), §1.2 (ethical, valid, reproducible, transparent), §2.2 (publishers converged at disclosure but do not test responsible conduct), and §2.4 (competency 2, non-determinism and model-version drift)
Clauses 1 to 4: The three-part test and its triggers	§2.1 (Dimension 1, labor versus judgment), and Appendix G (task taxonomy)
Clause 5: The four AI-use modes applied	§2.2 (search, co-author, validator, tutor), §2.4 (competencies 1 and 4), and Appendix C, Classes 1 to 4 and 6
Clause 6: The no-trace exemption	§3.3 (SAGE assistive threshold, Springer Nature copy-editing carve-out, ACM and IEEE exceptions), and Appendix G, rows 6 and 7
Clause 7: Disclosure is not permission	§3.3 (image-side category prohibitions), and Appendix C, Class 9
Clause 8: Scope covers staff and students	§1.3 (the plagiarism-policy ceiling, and Class A and B distribution)
Clauses 9 to 12: The five required fields	§2.2 (disclosure operationalization), §2.4 (competencies 1 and 2), §3.2 (institutional field requirements), and §3.3 (publisher content requirements)
Clause 13: The reproducibility set	§2.3 (reproducibility at a known model version and auditability), and §2.4 (competencies 2 and 3)
Clause 14: Retention	§2.3 (auditability via session logs), and §2.5 (systems axis)
Clauses 15 to 18: Where the disclosure goes	§2.5 (process axis: doctoral milestones), §3.2 (thesis-declaration and affidavit instruments), §3.3 (publisher locations), and §4.2 (Dimension 1 at <i>leading</i>)
Section 5: Statement templates	§2.2 (disclosure per the publisher convention applying to the destination), §4.2 (Dimension 2 at <i>leading</i>), and Appendix A (D2 <i>established</i> : institutional disclosure template aligned with publisher policy convention)
Clauses 19 to 21: Publisher interaction	§1.1 and §3.3 (the ten-week convergence, three convergences, two divergences, and disclosure locations)
Clauses 22 to 24: Funder interaction	§3.1 (the four-part funder shape, the NIH-to-Swedish variance, the DFG output-side duty, and the reviewer-upload prohibition)
Section 8: Worked good and bad disclosures	§2.2, §2.4 (competency 6, structured failure-mode reporting), and §3.3
Clauses 25 to 30: What not to require	§3.1 (the tri-agency burden concession), §2.4 (competency 1), Appendix C, Class 11 (AI-detection failure rates), and §2.3 (data-residency and privacy governance)
Section 10: Adoption checklist	Appendix A (D2 and D5 rows: <i>established</i> and <i>leading</i> criteria, including named owners per axis), and §4.2 (Dimension 5 at <i>leading</i>)

Section or clause of this instrument	Report section
Scope note	§1.2 (the four properties of responsible use), §2.4 (Dimension 4), and §4.1 (dimensional interdependence)

Table 8. Mapping between policy instrument sections and corresponding sections of the report.

Scope note

This standard specifically governs disclosure. It does not decide which AI uses are permitted in the first place. That broader decision is the work of your institution's overall AI-in-research policy (P1) and, for an individual student project, the supervisor-student agreement (G1). In the same way, this standard does not evaluate whether a disclosed use was technically competent. That capability belongs to Dimension 4 (AI-literate humans). A university that adopts this standard on its own will establish honest provenance across its research outputs, but it will still need human-in-the-loop discipline, responsible use in practice, tooling that supports responsible use, AI-literate researchers, and an institutional benchmarking grid to ensure the work behind those outputs is genuinely sound.

Requiring clear disclosure, done plainly and without punitive threats, is the most straightforward step in the framework and the one most institutions will likely adopt first. It is well worth adopting on its own, provided it is not mistaken for the entire job.

instats

Michael J. Zyphur, PhD · Professor and Director, Instats · instats.org · support@instats.org

Cite the pack. Zyphur, M. J. (2026). *The Institutional AI Readiness Pack: Self-Assessment and Implementation Tools for Responsible AI in Academic Research*. Instats Policy Series. <https://doi.org/10.61700/bv2nulyhht>

Companion report. Zyphur, M. J. (2026). *Responsible AI in Academic Research: A Competency Framework for Research Training*. Instats Policy Series. <https://doi.org/10.61700/t31oy23grr>

License. The pack and its instruments are licensed under [Creative Commons Attribution 4.0 International \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/). You may adapt them for institutional use with attribution.