

G1 — Supervisor–Student AI-Use Agreement

A one-sitting agreement in which an advisor or supervisor and a doctoral student decide, task by task, where AI can handle the routine work and where the judgment stays human, and write the decision down.

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

About this instrument

Purpose	To translate the framework's labor-versus-judgment demarcation into a signed, revisitable agreement for a specific student and project.
Who completes it	The doctoral student and the principal advisor or supervisor, together, in one sitting. Co-advisors, co-supervisors, and advisory committee members countersign. I use <i>advisor or supervisor</i> throughout to cover the US advisor as well as the UK and Australian supervisor. The short form <i>supervisor</i> is used where the formal supervisory relationship is the main point. I use <i>doctoral student</i> throughout, whatever the local term might be, such as <i>PhD student</i> or <i>graduate student</i> in the United States and much of Europe, <i>research student</i> in the United Kingdom, or <i>higher degree by research student</i> in Australia. <i>Thesis</i> means the thesis or dissertation, and <i>examiners</i> refers to the examiners or the dissertation committee members in US usage. Finally, <i>oral examination</i> refers to the final oral examination of the thesis, such as the viva in the United Kingdom and Australia, or the thesis defense in the United States.
Time required	About 30 minutes for the initial completion, and about 10 minutes at each scheduled revisit.
Related report sections	§2.1 (Dimension 1, human-in-the-loop discipline), §2.2 (Dimension 2, responsible use in practice), Appendix G (labor-versus-judgment task taxonomy), Appendix C (AI tool classes), and §5.3 (recommendations for graduate school leadership).
Cells of the 20-cell grid this instrument touches	D1–process, D1–people, D2–process, D2–people, D3–systems, D4–people.
What it produces	A completed and signed schedule, a named log location, and a scheduled revisit date. These provide three concrete artifacts that a graduate school can audit.

Table 1. Overview and core purpose of the artificial intelligence agreement instrument.

Drafting note for [INSTITUTION NAME]. Adopt this document as-is or amend it for your local context. Two decisions must be made locally before issuing it: where the AI-use log lives and how long it is retained (Section 5), along with the escalation route for disagreements (Section 8). Everything else can be used as written.

Section 1: Why this conversation happens at the start of the doctoral program

A doctoral project is a sequence of tasks. Some tasks are routine labor where an AI tool can carry the load, any mistake is visible and easy to fix, and the speed is a genuine benefit. Other tasks require scholarly judgment, such as framing the research question, interpreting an unexpected result, deciding what counts as an outlier, or weighing conflicting claims across a literature. If AI substitutes for those judgments, the project quietly stops being research, and nobody can tell from the finished text that it happened.

An institutional policy can state that principle in general terms, but it cannot determine where the line falls on a particular project, in a particular discipline, or for a particular student. In fact, *Responsible AI in Academic Research: A Competency Framework for Research Training* found that six of the thirty-eight universities in its sample reach Class D by adding codified competencies for graduate students and supervisors, curriculum expectations, and examiner-side rules (§1.3). Setting these boundaries is a job for the two people who actually know the project.

It protects the student. Use that was permitted, logged, and verified under the governing policy and approvals should not later be reclassified solely because expectations changed, subject to continuing legal, ethical, contractual, confidentiality, funder, and publisher obligations. The student will also have a clear written answer to any question an examiner is entitled to ask.

It protects the advisor or supervisor. It records that the conversation took place, when it happened, and what was agreed upon. This provides the supervisory oversight evidence that the report's Class D institutions build into their doctoral policies (§1.3).

It protects the work. Citation fabrication was measured at 18 to 55 percent in the Walters and Wilder citation-generation benchmark (2023) and at 28.6 to 91.4 percent across a medical literature test (Chelli et al. 2024, discussed in §2.2 and §2.4). Retrieval-anchored tools reduce this error rate without removing it entirely, and they shift the failure mode from an invented reference to a misattributed real one, which is much harder to spot (§2.3).

Nothing in this agreement permits anything that the institution's AI policy, the funder's grant conditions, or the research ethics approval prohibits. If there is any conflict, those formal policies and approvals take precedence.

Evidence this section produces: a dated agreement, filed against the student's program record.

Section 2: How to complete this in about 30 minutes

I recommend reading Sections 1, 5, 6, and 8 beforehand so you can use the meeting itself to make decisions.

Minutes	What you do
0–3	Fill in Section 3. Confirm which institutional AI policy, funder condition, and ethics approval apply.
3–20	Work through Sections 4A to 4F. Check one box per row. Write a condition only where the choice needs one. Focus on recording your specific decision rather than restating general policy.
20–24	Add any discipline-specific rows in Section 4G. Agree on where the log will live and who will review it (Section 5). Both parties should read the restricted data list in Section 6 and initial it.
24–27	Set the revisit dates in Section 7 aligned with the student's program milestones.
27–30	Read Section 8 and sign Section 10.

Table 2. Recommended agenda and estimated time allocation for completing the agreement meeting.

If discussing a row takes more than 90 seconds, select **Negotiated per instance** and move on. That is exactly what that option is designed for.

Section 3: Parties, project, and the documents that outrank this one

Field	Entry
Institution	[INSTITUTION NAME]
School or college / graduate school	
Doctoral student	
Student's stage at signing	☐ Before the first formal progress review ☐ Year 1 ☐ Year 2 ☐ Year 3 ☐ Later
Principal advisor or supervisor	
Co-advisor(s) / co-supervisor(s) / advisory committee	
Project title	
Discipline and sub-field	
Data types in the project	☐ Public documents ☐ Published literature only ☐ Human-participant data ☐ Identifiable or re-identifiable data ☐ Special-category personal data ☐ Third-party licensed data ☐ Industry-partner confidential data ☐ Other:
Research ethics approval number and expiry	
Institutional AI policy that governs this agreement (title, version, date)	
Funder conditions that apply	
Date of this agreement	[DATE]
Next scheduled review	[DATE]

Table 3. Institutional details, participant roles, and governing policy documents for the project.

Evidence this section produces: the named policy version, ethics approval number, and review date, each of which is independently checkable.

Section 4: The task schedule

The 21 tasks below represent the full labor-versus-judgment catalog from Appendix G of the report, serving as its working demarcation for PhD-level research across disciplines. You should **add discipline-specific rows in Section 4G** as needed. For example, computational biology might add rows for structure prediction and simulation, while humanities might add rows for archive interpretation and close reading. The competency that the framework calls for is agreeing on and recording this breakdown, rather than adhering to any single fixed list.

Check one box per row.

- **AI-assisted, with disclosure:** An AI tool may produce this output, provided the student verifies it against the primary source, the underlying data, or a test before it enters the research record. On a judgment row, AI may scaffold but must not substitute for the human decision, and the conditions must name the human responsible for the final decision. This use must be logged under Section 5 and disclosed under Section 9.
- **Human-only:** No AI-generated content enters the research record for this task.
- **Negotiated per instance:** There is no standing permission for this task. The student asks before each use and records the agreed approach in the log. Use this option when AI is used to scaffold a judgment (such as generating options, listing counterarguments, or challenging a draft position) while the final decision that enters the record remains the researcher's own. Appendix G prohibits AI from *deciding*, but not from *assisting the human who decides*.

Row markers.

- ♦ Appendix G classifies this as a judgment that humans must make. Checking this box confirms that you discussed it, rather than implying that an alternative was available.
- ■ This task is strictly human-only due to external rules. Academic publishers have moved beyond simple demarcation to explicit prohibition on these tasks.
- ▲ AI is expected to scaffold the task while the human makes the final determination. If you select AI-assisted, the conditions column must name the specific safeguard you will use.

The class labels in brackets reflect the defaults from Appendix G. Your agreed choice is recorded by the checkbox you select.

4A: Finding and reading the literature

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed (tool class, verification step, who checks)
T1	Initial literature scoping: search-term iteration and reading-list assembly from broad queries (<i>Labor</i>)	☐	☐	☐	
T2	Citation accuracy verification at draft time (<i>Labor</i>)	☐	☐	☐	
T3	Literature synthesis judgment: deciding which papers are central to the argument and resolving conflicting findings (<i>Judgment</i>)	☐	☐	☐	

Table 4. Literature search and review tasks with agreed assistance modes and conditions.

The essential condition to get right here. Regardless of what is checked for T1 and T2, every citation that remains in a draft must be verified against the primary source before it is cited. An AI tool may propose a candidate citation, but the student must confirm that the source exists, that it actually supports what the student claims, and that it comes from the publisher of record. This step is not negotiable, and the citation fabrication rates documented in §2.4 show why it is so critical.

Evidence: a reference manager with resolved DOIs, along with a log entry recording the verification check and its result.

4B: Designing the study

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed
T4	Hypothesis formulation prior to data collection (<i>Judgment</i>)	☐	☐	☐	
T5	Method choice and analytic-plan design (<i>Judgment</i>)	☐	☐	☐	
T6	Inclusion and exclusion criteria for a systematic review or meta-analysis (<i>Judgment</i>)	☐	☐	☐	
T7	Active-learning screening of titles and abstracts (<i>Labor</i>)	☐	☐	☐	

Table 5. Study design tasks with permitted levels of assistance and agreed conditions.

Note on study design. Tasks T4 through T6 define the core intellectual stance of the research. An AI tool may help generate candidate hypotheses, list potential methodological approaches, or draft sample wording for criteria. That is appropriate scaffolding, and Negotiated per instance is usually the best choice for it. However, what the student ultimately registers in a pre-registration or protocol represents their own intellectual commitment. Task T7 represents the routine labor component of T6. An open-source active-learning screener prioritizes the queue, but the student must review every borderline decision and every exclusion near the threshold.

Evidence: the pre-registration or protocol record (such as PROSPERO, OSF, or an institutional equivalent) with its timestamp, along with the screening tool's full decision history.

4C: Data, code, and analysis

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed
T8	Transcription of interviews and focus groups, with audio retained for human verification (<i>Labor</i>)	☐	☐	☐	

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed
T9	Code generation for data-cleaning and analysis pipelines (<i>Labor</i>)	☐	☐	☐	
T10	Statistical assumption checking and outlier adjudication (<i>Judgment</i>)	☐	☐	☐	

Table 6. Data processing and analysis tasks with assistance classifications and agreed conditions.

Two conditions to get right here. For T8, the original audio must be retained and checked against the generated transcript. Automated transcription engines can invent content (occurring in 187 of 13,140 audio segments in the Koenecke et al. 2024 evaluation, discussed in Appendix C, Class 10), and invented text is never flagged by the model as uncertain. Deciding which transcription deployment is permitted to process participant audio is covered in Section 6. For T9, the student must inspect the code diff, run the tests, and sign the commit. Around 29.5 percent of generated Python snippets in the 2024 ACM TOSEM evaluation contained a classified security weakness, and hallucinated package names remain a documented vulnerability in this class of tools (Appendix C, Class 5).

Evidence: original audio files retained alongside transcripts with a verification note, and version-control history showing the student as the committer along with passing test runs.

4D: Writing

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed
T11	Drafting prose from a researcher-generated outline, converting bullet points to paragraphs or drafting sections from notes (<i>Labor</i>)	☐	☐	☐	
T12	Drafting substantive content that the researcher cannot independently produce or validate (<i>Judgment</i>)	☐	☐	☐	
T13	Language polishing on researcher-generated text, such as grammar, register, and length adjustments (<i>Labor</i>)	☐	☐	☐	
T14	Translation of researcher-generated text between languages (<i>Labor</i>)	☐	☐	☐	

Table 7. Academic writing tasks with permitted assistance levels and specified conditions.

The distinction that matters here. In task T11, the researcher already knows the substantive content, and the AI tool simply helps adjust its form. In task T12, the tool supplies content that the researcher could not have produced or cannot independently verify, which creates a major risk of fabrication. As a practical test, if a student cannot explain a sentence without looking back at the tool's output, it falls under T12. Task T13 falls within the permitted exceptions of most major academic publishers and represents the lowest-risk activity here. Students writing in an additional language should certainly not be discouraged from using tools for language polishing. Task T14 requires that the original source text be retained so that the translation remains fully auditable.

Evidence: the outline or notes that preceded the draft with an earlier timestamp, along with the retained source text for any translation.

4E: Figures and images

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed
T15	Image or figure generation for illustrative or schematic purposes, such as talks, slides, teaching materials, and cover art (Labor)	☐	☐	☐	
T16	Image or figure generation for data presentation, including graphs, microscopy, gel images, and structural diagrams (Judgment: prohibited)	☐	☐	☐	Check Human-only.

Table 8. Visual and figure generation tasks with assistance classifications and agreed conditions.

Task T16 is where general guidance becomes an outright prohibition. Nine major academic publishers (Springer Nature, Nature, Science, Elsevier, Cell Press, The Lancet, JAMA, the New England Journal of Medicine, and the British Medical Journal) enforce an explicit prohibition against using generative AI to create or alter images in submitted manuscripts. The only exception is the rare case where the AI system itself is the explicit object of study (§3.3). Generating a data figure is not simply a matter of disclosure, because altering or generating data visually is treated as data fabrication.

Evidence: figure source files that trace directly to the underlying data, along with provenance metadata on any illustrative images.

4F: Review, approval, and examination

#	Task	AI-assisted, with disclosure	Human- only	Negotiated per instance	Conditions we agreed
T17 ▲	Adversarial review of one's own argument, such as red-teaming, devil's-advocate testing, or pre-mortem analysis (<i>Judgment, with AI augmentation</i>)	☐	☐	☐	
T18 ◆	Supervisor approval of an analysis chapter (<i>Judgment: humans only</i>)	☐	☐	☐	Check Human-only.
T19 ◆	Examiner judgment at the oral examination (<i>Judgment: humans only</i>)	☐	☐	☐	Check Human-only.
T20 ◆	Co-authorship determination (<i>Judgment: humans only</i>)	☐	☐	☐	Check Human-only.
T21 ◆	Peer-review decision on a submitted manuscript (<i>Judgment: humans only</i>)	☐	☐	☐	Check Human-only.

Table 9. Review, approval, and examination tasks with assistance levels and agreed conditions.

Task T17 is the row most often misunderstood. Having an AI model agree with a student's argument does not validate the research. In fact, agreement with a user's prompt framing is a well-documented failure mode of language models (§2.4, competency 5). An adversarial review is only genuine when the critique comes from a different model family than the one that assisted in drafting the argument (§2.4, competency 4), and when any disagreement is treated as an invitation for human judgment rather than a final verdict. If you check AI-assisted for T17, the conditions column must name that second model family.

Tasks T18 through T21 must remain human because they provide the verification mechanism for all the other work. These are the stages where a qualified person takes professional responsibility for the research. An automated tool cannot assume that responsibility, which is why none of the eighteen major academic publishers surveyed in §3.3 permit AI to be listed as an author. Task T21 also binds the student when serving as a peer reviewer. Uploading an unpublished manuscript to a public generative AI tool violates reviewer confidentiality at sixteen of those eighteen publishers, regardless of whether AI was used to help form the evaluation.

Evidence: the advisor's or supervisor's written approval on the chapter, the log entry naming the second model family used for T17, and the journal's reviewer disclosure record.

4G: Discipline-specific rows we are adding

You can add rows here for project-specific tasks that the general catalog does not cover. The competency framework treats the general catalog as a baseline that departments, schools, and supervisory teams can adapt and extend for their specific disciplines.

#	Task	AI-assisted, with disclosure	Human-only	Negotiated per instance	Conditions we agreed
T22		☐	☐	☐	
T23		☐	☐	☐	
T24		☐	☐	☐	

Table 10. Additional discipline-specific tasks with selectable assistance modes and agreed conditions.

Evidence: completed task rows, or a signed note explaining why the general catalog was sufficient for the project.

Section 5: Disclosure and logging

Being on the labor side of the line is not a blanket permission slip. Every AI use that meets any of P2's trace, judgment, or reproducibility tests must be logged and disclosed, regardless of its task category. The absence of a surviving output does not by itself make a use non-disclosable. AI use that touches a judgment task must be logged and disclosed even when none of the tool's output survives.

What is recorded. One entry per use, capturing ten specific fields:

1. Date.
2. Task number from Section 4.
3. Tool class and the specific tool used, including its model version or date-stamped endpoint.
4. Deployment tier: institutional enterprise tenant, locally executed model, or open consumer tier.
5. What was provided to the tool, described in one line, explicitly stating whether any unpublished or personal data was included.
6. The prompt text, or a stable file path to the stored prompt.
7. What the tool returned, described in one line, along with the location where that output is saved.
8. The verification step performed and its result, expressed as a count whenever possible (for example, "14 candidate citations returned, 12 resolved to the primary source, and 2 did not exist and were removed").
9. Whether the output entered the formal research record.
10. Where the use will be formally disclosed, such as in the thesis declaration or a manuscript declaration, or, only where all three P2 tests are no, marked as not disclosable.

Where it lives. Store the log on [INSTITUTION NAME]'s research data management platform within this project's records area, maintaining one file per project. Do not store it in a personal cloud account or rely on the tool's built-in session history. Vendor histories sit outside institutional control, and the vendor can delete or alter them at any time.

Who does what. The student records each entry at the time the tool is used, rather than reconstructing entries later. The advisor or supervisor reviews the log at each milestone listed in Section 7 and initials that review. Neither party should edit an earlier entry. If a correction is necessary, add an explanatory note as a new entry.

How long it is kept. Retain the log for [RETENTION PERIOD], aligned with [INSTITUTION NAME]'s research records retention schedule.

What the log is not. This log is designed to record tool use, not to evaluate the student's thinking. It is neither a productivity metric nor a surveillance mechanism. Its real purpose is to ensure that the student can confidently answer any question about their methods during an examination, and

to ensure that any research integrity inquiry can be resolved based on clear documentary evidence rather than unassisted memory.

Drafting note. *If your institution already uses a formal thesis submission declaration, you should point field 10 to that existing form rather than creating a separate disclosure mechanism. Structured disclosure fields provide clear records, whereas vague, free-text afterthoughts do not.*

Evidence: the log file itself, stored at a documented institutional path, with contemporaneous entries and the advisor's or supervisor's initials at each scheduled milestone.

Section 6: Data and confidentiality: what never goes into a third-party tool

Pasting text or uploading files into an AI tool constitutes a disclosure to that tool's commercial operator. Both parties must initial this section to confirm that they have read and understood these restrictions.

Never upload or paste the following items, regardless of what has been checked in Section 4:

- Identifiable or re-identifiable data about human participants, including raw interview audio and transcripts that have not been fully de-identified.
- Special-category personal data, such as information regarding health, genetics, biometrics, racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, sex life, or sexual orientation (under GDPR Article 9(1)).
- Any participant data where the ethics approval and participant consent forms do not explicitly permit disclosure to third-party commercial processors. If the consent form is silent on this point, the answer is an immediate no, rather than an assumption that it is probably fine.
- Data governed by a license, a data-transfer agreement, or custodian terms that prohibit onward disclosure.
- Commercial-in-confidence or proprietary material provided by an industry partner.
- Unpublished third-party work shared in confidence, such as manuscripts received for peer review or grant applications received for assessment.
- Material under formal embargo, patent-pending disclosures, or other intellectual property restrictions.
- Personal data about identifiable colleagues, students, or co-authors.

The deployment tier rule. Subject to the absolute prohibitions above, other unpublished research data or writing may be entered only into an enterprise-tenanted platform or a locally executed model where training on user inputs is disabled and data residency can be explicitly identified. Open consumer tiers typically enable model training on user inputs by default, so they must be reserved strictly for public material. This is fundamentally an institutional procurement safeguard rather than an individual choice at the keyboard. The clear demonstration of this risk was the 2023 Samsung incident, where three confidential data leaks occurred within twenty days, prompting a complete company-wide ban on consumer generative AI tools (§2.3).

The three-question test before pasting any data:

1. Would I be comfortable having this text held indefinitely by a third party outside institutional control?
2. Do the ethics approval and participant consent forms explicitly permit disclosure to a third-party processor?
3. Can I clearly identify the deployment tier, verify that training on input is turned off, and confirm the data residency?

If the answer to any of these three questions is no, do not paste the text. Seek clarification first.

Escalation. Questions concerning data protection and confidentiality should not be decided informally between the advisor and the student. They must be referred to [DATA PROTECTION OFFICER / RESEARCH ETHICS OFFICE], whose determination is binding on both parties.

Student initials: __ **Advisor or supervisor initials:** ____ Date: [DATE]

Evidence: an initialed section, and the named deployment tier recorded in log field 4 for every entry involving unpublished material.

Section 7: When we revisit this agreement

AI tools evolve quickly, model versions change, and research projects develop over time. An agreement that is filed away and never reopened will become completely outdated by year three.

Milestone triggers: set the review dates now.

Milestone	Date	What is checked	Output
First formal progress review		The full task schedule and a review of the log	Reconfirmed or amended agreement
Mid-program review		The full task schedule and a review of the log	Reconfirmed or amended agreement
Completion seminar		Sections 4F, 5, and 9, along with a practice run of the "describe and defend" questions	Readiness note
Thesis submission		Section 9 in full, ensuring the formal declaration is complete	Signed declaration
Oral examination		The log is made available if the examiners request it	Log provided on request
Examiner-report adjudication		Any AI-related examiner comments	Response record
First publication from the thesis		The target publisher's specific disclosure rules	Manuscript declaration

Table 11. Review milestones, scheduled dates, evaluation checkpoints, and expected outputs.

Event triggers: either party can request a review within [10] working days whenever any of the following events occur:

- The student wants to use a new AI tool class that is not covered in the current schedule.
- The underlying model version for an existing tool undergoes a major update.
- The project begins using a new data type, such as collecting human participant data for the first time, receiving industry partner data, or acquiring a licensed dataset.
- The project's research ethics approval is amended, renewed, or expires.
- A new advisor or supervisor joins the supervisory team, an advisor leaves, or the principal advisor changes.
- [INSTITUTION NAME] updates its institutional AI policy, or the funding body or target journal updates its requirements.
- **A procedural failure occurs.** For example, if a fabricated citation is discovered in a draft, unverified AI output is submitted to an advisor, or a confidentiality rule from Section 6 is breached. Reopening the agreement after a failure is designed as a constructive correction rather than a disciplinary action, and the steps in Section 8 guide how to resolve it.

Evidence: dated revisit entries on this agreement initialed by both parties, along with a record of completed reviews matched against scheduled milestones.

Section 8: What happens when we disagree

Differences of opinion are a normal part of research supervision. This section provides a clear, step-by-step process so that disagreements do not create unnecessary risks for either the student or the supervisor.

1. **Attempt to resolve the issue directly and document the result.** Most disagreements involve a specific task or use case and can be resolved in a single conversation. Once resolved, record the agreed outcome in the project log under the relevant task number.
2. **If the issue is not resolved within [10] working days, adopt the more restrictive position in the interim.** The student should continue their research on that conservative basis so that they are never expected to work under unresolved ethical or policy ambiguity.
3. **Seek input from a co-advisor, co-supervisor, or the advisory committee chair.** They will review the question and provide their written perspective within [10] working days.
4. **Escalate to [GRADUATE SCHOOL / DEAN OF GRADUATE STUDIES].** The graduate school will provide a formal determination, which will be recorded as an official amendment to this agreement.
5. **Data protection and confidentiality issues bypass this process entirely.** Any questions regarding participant privacy or sensitive data must be directed immediately to [DATA PROTECTION OFFICER / RESEARCH ETHICS OFFICE] as outlined in Section 6.

Four standing protections that bind both parties:

- **No retrospective reclassification.** Policy changes do not retrospectively reclassify conduct permitted when undertaken and properly logged, subject to continuing legal, ethical, contractual, confidentiality, funder, and publisher obligations. Policy amendments apply from the date they are formally signed.
- **Disagreements do not affect academic progression.** Discussing where boundaries should be set is an ordinary part of academic training, and these discussions must not negatively influence a student's formal progress review. A deliberate failure to *disclose* tool use is a separate issue that falls under [INSTITUTION NAME]'s research integrity policies rather than this supervisory agreement.
- **The student may seek independent advice.** The student is entirely free to consult [GRADUATE SCHOOL / RESEARCH INTEGRITY ADVISER / STUDENT ADVOCATE] regarding this agreement without needing prior permission from their supervisor or having to notify them.
- **Limits on what an advisor or supervisor can require.** An advisor or supervisor cannot require a student to use an AI tool if the student raises a reasoned objection to it. Furthermore, a supervisor cannot require any sensitive material listed in Section 6 to be entered into a consumer-tier tool, nor can they penalize a student solely because a later agreement revision restricts tool use that was previously approved and logged, subject to the obligations above.

Evidence: log entries documenting any disagreement and its resolution, along with any written determinations from step 3 or 4 filed alongside this agreement.

Section 9: What changes at thesis submission and examination

During the doctoral program, this document functions as an active working agreement between the student and their supervisor. Once the thesis is submitted, it becomes part of the formal institutional record.

At submission. The student completes [INSTITUTION NAME]'s formal thesis declaration and describes their AI use in the preliminary pages of the dissertation. This statement should be drawn directly from the log described in Section 5 rather than reconstructed from memory. The supervisor's countersignature confirms that this agreement was established, that it was reviewed at the required milestones, and that the supervisor has inspected the log. This signature confirms supervisory oversight rather than verifying every single prompt or output.

At the oral examination. The doctoral student must be fully prepared to describe and defend any AI use, as well as the substantive content of the thesis itself, during the final defense. This reflects the standard established by the University of Toronto School of Graduate Studies, which represents the clearest practical model for Dimension 1 at the point of examination (§5.3). The student should practice answering these questions during their completion seminar, well before the defense takes place. If the examiners ask to review the project log, it will be provided to them.

Examiner responsibilities. The final academic evaluation must remain the independent judgment of the human examiners (task T19). An AI system must never participate as an examiner or draft any portion of an examiner's formal report. Examiners must not upload an unpublished thesis or examination material into any generative AI tool, including an enterprise-tenanted or locally executed system. A dissertation under examination is confidential scholarly work under Section 6, and examiners are bound by the same confidentiality standards that govern peer reviewers.

At publication. When preparing journal submissions, disclosure transitions from the university's internal template to the specific requirements of the target publisher, which are often stricter and more detailed. Four consistent principles emerged from the report's survey of eighteen academic publishers in §3.3. First, AI cannot be listed as a co-author under any circumstances. Second, seventeen of the eighteen publishers require explicit disclosure of generative AI assistance upon submission. Third, all major biomedical and multidisciplinary publishers strictly prohibit the generative creation or visual manipulation of data figures. Finally, decisions about co-authorship on any resulting papers must be made exclusively by the human contributors (task T20).

After submission. This completed agreement will be retained in the student's official academic file for [RETENTION PERIOD]. It is not a public document and will not be shared beyond the supervisory team, the graduate school, the examination committee if requested, or an authorized research integrity panel.

Evidence: the completed thesis declaration, the statement in the dissertation's preliminary pages, the archived project log, and the supervisor's countersignature.

Section 10: Completion check and signatures

Completion checklist. Review these items before signing. If any item is marked partial or no, treat it as the immediate next step to resolve before finalizing the agreement.

Check	Answer
Every row in Sections 4A to 4F has exactly one checkbox selected.	☐ Yes ☐ Partial ☐ No
Every row marked <i>AI-assisted</i> or <i>Negotiated per instance</i> includes a clear written condition.	☐ Yes ☐ Partial ☐ No
Section 4G has been reviewed, and it either contains discipline-specific rows or a brief note explaining why none were needed.	☐ Yes ☐ Partial ☐ No
The log location in Section 5 specifies a secure institutional file path rather than a personal account.	☐ Yes ☐ Partial ☐ No

Check	Answer
Section 6 has been initialed by both the student and the supervisor.	☐ Yes ☐ Partial ☐ No
Section 7 includes specific calendar dates for each scheduled program milestone.	☐ Yes ☐ Partial ☐ No

Table 12. Final completion checklist items and recorded compliance status for the agreement.

Signatures.

Role	Name	Signature	Date
Doctoral student			
Principal advisor or supervisor			
Co-advisor or co-supervisor			
Advisory committee chair (where applicable)			

Table 13. Participant names, assigned roles, signatures, and execution dates authorizing the agreement.

Next review date: [DATE] · **Version of this agreement:** __ · **Filed at:** ____

Appendix: Worked example

Illustrative only. The individuals, research project, and institution are hypothetical, included so that adopters can see what a fully completed agreement looks like.

Header. Rivermark University, School of Health Sciences. Student: Alex Ferreira, second year. Advisors and supervisors: Professor Dana Whitmore (principal), Dr. Sam Ige (qualitative methods). Project: *Air-quality exposure and pediatric respiratory admissions*. Data: linked administrative health data under a custodian agreement, along with 24 clinician interviews. Ethics approval ETH-2026-0418, expiring June 30, 2029. Institutional AI policy v2.1. Agreed on September 14, 2026, with the next review scheduled for March 2, 2027.

Completed schedule.

#	Task	Selected	Conditions recorded
T1	Literature scoping	AI-assisted	Retrieval-anchored assistant on enterprise tenant. Nothing enters the reference manager without being resolved to a primary source.
T2	Citation verification	AI-assisted	The tool proposes matches, and Alex resolves each candidate to the publisher of record. Results are logged as a count.
T3	Literature synthesis judgment	Human-only	Deciding which cohort studies are central to the argument represents the core contribution of the thesis.
T4	Hypothesis formulation	Human-only	Fixed in the OSF pre-registration before the linked dataset is released.
T5	Method choice and analytic plan	Negotiated per instance	An AI model may list candidate approaches to exposure misclassification, but Dana and Alex make the final choice. Each session is logged.
T6	Inclusion / exclusion criteria	Human-only	Written by Alex, approved by Dana, and formally registered prior to data extraction.

#	Task	Selected	Conditions recorded
T7	Active-learning screening	AI-assisted	Locally run open-source screener. Alex reviews all borderline decisions and every exclusion within 5 percent of the decision boundary. The screening history is exported.
T8	Interview transcription	AI-assisted	Locally executed only so that audio never leaves the institutional environment. The audio is retained, and every transcript is checked against the recording. Two sections were corrected for invented clause endings.
T9	Analysis code	AI-assisted	Enterprise tenant with schema descriptions only, ensuring no participant data is pasted. Alex inspects the diff, runs the tests, and signs the commit.
T10	Assumption checking and outlier adjudication	Human-only	Alex adjudicates, and Dana countersigns outlier decisions before they are applied.
T11	Drafting from Alex's outline	AI-assisted	The outline is dated and stored before drafting begins. Disclosed in the thesis declaration.
T12	Drafting content Alex cannot validate	Human-only	Any sentence that Alex cannot explain without looking at the tool's output is removed.
T13	Language polishing	AI-assisted	Portuguese is Alex's first language. Polishing is permitted on Alex's own text, with a general AI-use disclosure noted at publication.
T14	Translation	Negotiated per instance	Permitted only for preparing a Portuguese-language summary for the partner health service. The original source text is retained.
T15	Illustrative figures	AI-assisted	Conference presentation schematics only. Image provenance metadata is retained.
T16	Data figures	Human-only	Human-only. All cohort figures must be generated directly from the analysis pipeline.
T17	Adversarial review	AI-assisted	Critiques must come from a different model family than the one used for T11. Alex logs what the model raised and how it was addressed. Agreement by the model is not treated as validation.
T18	Supervisor approval of an analysis chapter	Human-only	Human-only. This requires Dana's professional judgment.
T19	Examiner judgment	Human-only	Human-only. This remains the exclusive responsibility of the examiners.
T20	Co-authorship determination	Human-only	Human-only. Alex, Dana, Sam, and the health service collaborator make all co-authorship decisions.
T21	Peer-review decisions	Human-only	Human-only. Alex reviews for two journals, and unpublished manuscripts are never uploaded to any tool.
T22	<i>Added:</i> record-linkage rule construction	Human-only	The data custodian agreement strictly prohibits disclosing linkage keys. Rules are written directly by Alex against the custodian's technical specification.

#	Task	Selected	Conditions recorded
T23	<i>Added</i> : qualitative coding-frame development	Negotiated per instance	Requires Sam's approval each time. An AI model may suggest potential code labels from de-identified interview excerpts, but the coding frame is decided by Alex and Sam.

Table 14. Sample completed task schedule showing selected assistance modes and recorded conditions.

Log excerpt (three of 46 entries).

Date	Task	Tool / version	Tier	Verification result	Entered record?
2026-10-03	T2	Retrieval-anchored assistant, v4.2	Enterprise tenant, EU residency	31 candidate citations returned, 27 resolved to the primary source, 3 misattributed and corrected, and 1 did not exist and was removed	Yes
2026-11-19	T8	Local transcription model, medium weights	Local execution	Interview 11 checked against audio recording, and two invented clause endings were corrected at 14:22 and 31:08	Yes
2027-01-27	T17	Model family B, v2.0	Enterprise tenant	Two threats to validity raised. Seasonal confounding was valid and led to an analysis change, while the second point was a misreading and was rejected with an explanatory note	Yes

Table 15. Sample artificial intelligence usage log entries detailing tools, tiers, and verification outcomes.

One recorded disagreement. On November 17, 2026, Alex proposed using AI for a first-pass coding of interview transcripts. Sam considered this an interpretive task equivalent to a T3 judgment in qualitative research. Because they could not resolve the question during the meeting, the more restrictive human-only position applied for eleven days. On November 28, they resolved the issue by adding task T23 as *Negotiated per instance* with an explicit de-identification requirement. This resolution was logged, and no further escalation was necessary.

Signatures. Alex Ferreira and Dana Whitmore signed on September 14, 2026, and Sam Ige signed on September 16, 2026. Filed at: RDM/HS/ETH-2026-0418/governance/ai-use-agreement-v1.

Crosswalk to the Responsible AI framework

Section of this instrument	Report section
About this instrument (cells touched)	§2.5 (the four axes), Appendix A (D1 and D2 rows)
Section 1: Why this conversation happens	§2.1 (Dimension 1), §1.3 (institutional baseline), §2.4 (fabrication-rate evidence), and §5.3 (graduate-school recommendations)
Section 2: How to complete it	§2.5 (process axis)
Section 3: Parties and scope	§2.5 (people axis)

Section of this instrument	Report section
Section 4A: Finding and reading the literature (T1–T3)	Appendix G, rows 1–3, §2.2 (<i>search</i> mode), and Appendix C, Classes 1 and 6
Section 4B: Designing the study (T4–T7)	Appendix G, rows 9, 11, 12, and 13, §2.1, and Appendix C, Class 7
Section 4C: Data, code, and analysis (T8–T10)	Appendix G, rows 14, 8, and 10, and Appendix C, Classes 5 and 10
Section 4D: Writing (T11–T14)	Appendix G, rows 4–7, §2.2 (<i>co-author</i> mode), and Appendix C, Class 2
Section 4E: Figures and images (T15–T16)	Appendix G, rows 15–16, §3.3, and Appendix C, Class 9
Section 4F: Review, approval, and examination (T17–T21)	Appendix G, rows 17–21, §2.2 (<i>validator</i> mode), §2.4 (competencies 4 and 5), and §3.3
Section 4G: Discipline-specific rows	Appendix G, closing note on local extension
Section 5: Disclosure and logging	§2.2 (disclosure), §2.3 (auditability and reproducibility at a known model version), §2.4 (competency 2), and §3.3
Section 6: Data and confidentiality	§2.3 (data residency and training-on-input governance), §3.3 (reviewer confidentiality), and Appendix C, Classes 2 and 10
Section 7: When we revisit this	§2.5 (process axis: the progress milestones), and Appendix A (D1 <i>established</i> : known cadence)
Section 8: What happens when we disagree	§2.1, and §5.4 (adjudication framework for AI-related cases)
Section 9: Submission and examination	§2.1, §3.3 (publisher convergence), and §5.3 (examiner-side rules)
Section 10: Completion check and signatures	Appendix A (D1 and D2 <i>established</i> criteria)
Appendix: Worked example	Appendix G, applied

Table 16. Mapping between sections of the agreement instrument and corresponding report sections.

Note on prior art and independence

Supervisor-student AI agreements are an emerging genre. The University of New Hampshire published a *Collaboration Agreement for Use of Generative AI in a Research Project Template* in April 2026 under a CC BY-NC-SA 4.0 license, which is incompatible with this pack's CC BY 4.0 license. This instrument was therefore authored independently, developed directly from the labor-versus-judgment task taxonomy in Appendix G of the companion report. No text, structure, or task list is derived from the New Hampshire document, which is noted here as prior art rather than as a source.

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 \(CC BY 4.0\)](#). You may adapt them for institutional use with attribution.