

B2 — Briefing for Deans and Associate Deans of Research

The five decisions that school and college leadership own regarding AI use in doctoral research, five practical questions to test whether those decisions have been made, and how your unit fits into the overall institutional assessment.

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

About this instrument

Purpose	Make and provide evidence for the school- or college-level decisions that an institutional AI policy cannot make for you.
Who reads it	Deans, Associate Deans for Research and Research Training, department chairs, and heads of school. You don't need any prior reading of <i>Responsible AI in Academic Research: A Competency Framework for Research Training</i> to use this briefing.
Time required	15 minutes to read, and around 90 minutes in a meeting with your department chairs and heads of school to work through the five questions.
Related report sections	§1.3, §2.1–§2.5, §3.2, §4.1–§4.3, §5.2–§5.4, Appendices A, C, and G.
Cells touched	D1-policy · D1-process · D2-policy · D2-process · D3-systems · D4-people · D4-process · D5-people · D5-process

Table 1. Overview of key features and the operational purpose of this instrument.

Why this reaches your desk

In the report, I reviewed 38 top-tier doctoral universities across 15 countries and jurisdictions and found that only six publish an AI policy that reaches beyond basic research integrity into AI literacy, supervision, and examination. In fact, around 40 percent of these institutions stop at the level of basic plagiarism policies.

In most institutions, the missing operational layer is the school or college. Across the major research universities in the report's United States sample (including Harvard, MIT, Michigan, Stanford, Yale, and Princeton), doctoral-research AI policy is devolved down to individual schools, departments, advisors, and supervisors, while central university policies remain thin. Similarly, while the University of Oxford has a substantive central policy on researcher accountability and disclosure, its discipline-specific expectations are left entirely to departmental guidance.

Devolution itself isn't necessarily a bad thing, but uncoordinated or undesigned devolution creates major gaps. For example, German universities also devolve policy to local faculties. At LMU Munich, the binding doctoral policy is issued directly by its medical school, and RWTH Aachen relies on guidance from its Biology and Chemistry departments. However, those local guidelines rest on a clear national standard: the legally binding sworn declaration that every doctoral student is required to sign.

The practical question for your leadership team is whether your school or college is the place where university policy gets translated into actionable disciplinary practice, or whether it's simply where central policy runs out.

The five decisions you own

Each decision has a tangible artifact (such as a document, register, or curricular plan) that proves it has been made and supplies evidence toward the named cell or cells. No artifact advances a dimension by itself. Advancement is determined only through a fresh A1 assessment across all four cells for that dimension. The framework in the report evaluates five core dimensions across four operational axes (policy, people, systems, and process) at four developmental levels: *absent*, *nascent*, *established*, and *leading*.

1. Where the labor and judgment line falls in your disciplines. Dimension 1, **human-in-the-loop discipline**, requires a clear, published demarcation between mechanical *labor* (tasks AI can handle, provided a human verifies the output) and expert *judgment* (tasks where substituting AI silently degrades academic rigor and quality). The catalogue in the report covers tasks such as literature scoping, citation checking, prose drafting, code generation, method choice, outlier adjudication, and examiner judgment (in US usage, examiners are dissertation committee members). The catalogue is brief and applicable across disciplines so that individual schools and colleges can adapt and extend it rather than simply copy it. *Artifact*: A dated school or college extension of the task catalogue that lists the specific disciplines covered and names the [COMMITTEE] that formally approved it. *Evidence toward*: D1-policy → nascent, and later D1-policy → established once the taxonomy is actively taught in your doctoral methods curriculum and applied during student advising and supervision.

2. Your school or college task-level rules for the four AI-use modes. Dimension 2, **responsible use in practice**, is operationalized across four distinct modes of engagement: **search, co-author, validator, and tutor**. While central university policies often mention these modes, they rarely explain what a researcher or graduate student in your specific field should or shouldn't do within each mode. Your local school or college rules can be more detailed and specific than central institutional policy, but they must never contradict it. *Artifact*: Published school or college guidance that cites the specific institutional policy clause it sits under and establishes a clear review schedule. *Evidence toward*: D2-policy → established.

3. How advisors and supervisors are equipped, and whether that preparation is verified. A graduate student cannot be expected to develop a competency that their advisor or supervisor does not demonstrate in their own practice. Central graduate schools typically fund and deliver supervisor training, but your school or college must decide whether completing that training is a formal condition for supervisory registration, and who is responsible for checking the register. Without that requirement, training remains voluntary and this dimension will remain stuck at *nascent*. *Artifact*: A supervisory register that includes a training completion date field, along with a documented count of completions against the full roll of active supervisors. *Evidence toward*: D4-people → established.

4. What issues leave your school or college for central escalation. There are four distinct types of AI-related cases, each requiring different evidentiary standards and procedural responses: undisclosed AI use, disclosed but unverified AI use, AI-fabricated content (such as fake citations, fabricated data, or synthetic images), and broader AI-assisted research misconduct. When institutions fail to separate these categories, they often end up applying standard student plagiarism rules to complex research integrity matters, or treating serious data fabrication as routine drafting issues. Your leadership team must decide which situations should be handled locally through informal conversations between students and advisors, which matters must be formally escalated to the research integrity office, and what specific thresholds trigger that

escalation. *Artifact*: A one-page escalation flowchart naming the responsible officer at each step, together with a summary log of matters routed over the past twelve months categorized by case type. *Evidence toward*: D1-process → established and D2-process → established.

5. What your school or college funds and supports locally. The main recurring commitment is curriculum time. The six core AI-era research competencies (**citation verification, model-and-parameter specification, prompt-as-fork-in-the-garden discipline, model-heterogeneity in adversarial review, sycophancy detection and human-as-verifier discipline, and structured failure-mode reporting**) need to be integrated as explicit learning outcomes in your doctoral methods curriculum, with an assigned faculty owner and a published review schedule. The primary one-off commitment is surveying the AI tools that your researchers and graduate students report using and mapping them against six key procurement criteria (**verifiable citation, data residency, uncertainty reporting, reproducibility, auditability, and open-source or local hosting options**). Doing this exercise may uncover specialized tools that central IT procurement has never evaluated. *Artifact*: A named methods curriculum lead, a published syllabus review schedule, an allocated budget line for faculty release time, and a completed software mapping submitted to the central institutional procurement committee. *Evidence toward*: D4-process → established and D3-systems → nascent.

Handling disciplinary difference honestly

The exact same research task can sit on opposite sides of the line depending on the academic field.

	Where AI is labor, with human verification	Where AI cannot substitute
Experimental life sciences (wet lab)	Drafting experimental protocols from researcher notes, transcribing lab meetings against retained audio recordings, or generating structural priors from prediction tools checked against reported confidence scores (keeping in mind that roughly one in ten high-confidence atomic positions is still materially inaccurate)	Generating or retouching any empirical data figure (including microscopy, gel images, and structural diagrams), which is the framework's one flat prohibition across all research
Humanities thesis or dissertation	Polishing style and clarity on the student's own drafted prose, converting academic registers, or assembling an initial reading list for human verification	Interpreting an archival document, or conducting a close reading of a primary text
Computational fields	Generating routine code for data cleaning and analysis pipelines, provided the researcher carefully reviews the diff, runs full test suites, and signs the commit (recognizing that a substantial minority of AI-generated code snippets contain known security vulnerabilities)	Selecting research methods and designing the analytic plan, adjudicating data outliers, or determining whether an empirical finding holds up under adversarial testing

Table 2. Appropriate artificial intelligence labor and non-substitutable tasks across academic disciplines.

While specific task boundaries vary across academic disciplines, three baseline principles must remain constant across all schools and departments:

- **The requirement to articulate standards.** The primary competency is that every school or college publishes a clear task taxonomy for its fields, not that every school shares the exact same rules. A school or college that fails to document and publish its taxonomy scores *absent* on this dimension.

- **The disclosure requirement.** Classifying an activity as acceptable mechanical labor is not an exemption from transparency. Any AI use on labor tasks still requires full disclosure according to institutional reporting guidelines.
- **The core human-only list.** Four critical academic judgments must remain strictly human in every discipline without exception: examiner evaluation at the final oral examination (such as the dissertation defense or viva), advisor approval of a student's analysis chapter, formal authorship determination, and the final peer-review recommendation on a submitted manuscript. No school, department, or committee may grant an exception to this list.

Drafting note. If a department chair or head of school argues that their discipline requires an exemption that institutional policy does not permit, that represents a formal request to change institutional policy. That request must be submitted to [COMMITTEE] for consideration, because it is not a decision that a school or college can make independently.

Five questions to ask your department chairs, heads of school and supervisors

Every answer must be backed up by tangible documentation. When you work through these questions with your chairs and school heads, focus on recording verifiable evidence rather than general impressions or informal intentions.

1. Which research tasks in this department's or school's disciplines may AI carry, and which stay human, and where is that written down?

☐ Yes ☐ Partial ☐ No (Evidence: the written document, its approval date, the approving body, and the specific disciplines covered. An informal verbal understanding scores a "No".)

2. For graduate students who submitted their theses or dissertations in the last twelve months, is there a formal record of an AI-use conversation between the student and their advisor or supervisor?

☐ Yes ☐ Partial ☐ No (Evidence: the specific form or milestone tracking field that captures this discussion, along with the count of completed records relative to total student submissions.)

3. When a graduate student presents an AI-assisted analysis chapter, what specific competencies or verification steps must they demonstrate before an advisor approves it?

☐ Yes ☐ Partial ☐ No (Evidence: written supervisory expectations and criteria, along with at least one milestone record demonstrating how they were applied in practice. An excellent example of this standard comes from the University of Toronto School of Graduate Studies: the doctoral candidate must be fully prepared to describe and defend any use of generative AI, alongside the intellectual contents of the thesis, during their final oral defense.)

4. Which AI tools are researchers and graduate students in your unit currently using on unpublished research data, and which of those tools have formally passed the institutional procurement and security review?

☐ Yes ☐ Partial ☐ No (Evidence: the department or school tool inventory, matching approval entries in the central procurement register, and the resulting list of active tools that have never gone through institutional review. That unvetted list is often your most informative finding.)

5. What formal procedure happens in your unit when a citation in a researcher's or graduate student's work cannot be verified or resolved to a real primary source?

☐ Yes ☐ Partial ☐ No (Evidence: the documented escalation pathway, the specific officers named at each stage of review, and a log of matters raised over the past twelve months along with their recorded outcomes.)

Your school or college in the institutional assessment

The comprehensive assessment for [INSTITUTION NAME] evaluates 20 distinct cells across the institution, representing five core dimensions crossed with four operational axes (policy, people, systems, and process). Your school or college completes these same 20 cells for your local unit. Four fundamental scoring rules govern how you interpret your results:

- **A dimension's overall level is determined by the lowest score among its four axes.** The report requires all four axes to operate for the *established* level. The pack applies that rule by setting the dimension at the level of its lowest-scoring axis. It calls that limiting axis the *binding axis*. For instance, if your school has an excellent written policy (*established*) but lacks any operational workflow (*nascent*), the dimension as a whole remains at *nascent*.

- **Maturity levels progress cumulatively like rungs on a ladder.** Each level builds directly on the previous one. An institution or school cannot claim a *leading* status on any dimension if it has not yet satisfied all the baseline requirements for *established*. In other words, you have to secure each foundational rung before you can step up to the next level.
- **There is no single summary score for a school or college, nor should there ever be.** Reducing your school to a single average score obscures where the actual strengths and vulnerabilities lie. Instead, report your level across each of the five dimensions, identify the binding axis holding each dimension back, and pinpoint the primary dimension that constrains your school overall. That detailed diagnostic profile shows you exactly where policy and resources need to be directed.
- **Assessment coverage is critical.** If your team answers fewer than 80 percent of the diagnostic indicators, the assessment result must be marked as *Provisional* until the remaining evidence is gathered and the information gaps are resolved.

To get started immediately, assign a specific faculty or administrative lead to oversee your school's assessment profile, and establish a firm date to present the completed findings to [COMMITTEE]. Taking those two simple steps supplies evidence toward D5–people and D5–process at the school level. Dimension 5 advances only if a fresh A1 rescore confirms that all four D5 cells qualify at the target level.

Comparing your local profile to the central university profile is often the most valuable diagnostic exercise you can conduct. If your school has reached *established* on human-in-the-loop discipline while the central institution remains at *nascent*, you should share your local artifacts and guidance across the wider university. Conversely, if your school sits at *absent* on preparing AI-literate researchers while the central institution is already at *established*, your unit clearly needs targeted development, and this assessment provides the exact evidence you need to request institutional funding.

One important note of caution: the report's four-class description of institutional policy scope is a descriptive categorization of policy breadth, which is completely distinct from the maturity grid. You should never attempt to infer one classification from the other.

Crosswalk: briefing item to report section

Item	Report section
Why this reaches your desk (baseline, class distribution, and devolution)	§1.3, §3.2
Decision 1: labor and judgment line	§2.1, Appendix G, §5.2
Decision 2: task-level rules for the four AI-use modes	§2.2, Appendix A (D2), §5.2
Decision 3: supervisor equipping and its verification	§2.4, §4.2 (D4), §5.3
Decision 4: escalation and the four AI case types	§2.1, §2.2, §5.4
Decision 5: curriculum ownership, review schedule, and tool mapping	§2.3, §2.4, §5.2
Disciplinary profiles (wet lab, humanities, and computational fields)	Appendix G, Appendix C (Classes 5, 8, 9, 10)
The three invariants (articulation, disclosure, and human-only list)	Appendix G (closing notes), §2.2
Question 1: written task demarcation	§2.1, Appendix A (D1)

Item	Report section
Question 2: recorded candidate and supervisor AI conversation	§4.2 (D1), §5.3
Question 3: demonstration before chapter approval	§2.4, §5.3, Appendix G
Question 4: tools in use versus tools gated	§2.3, §5.2
Question 5: unresolvable citation escalation	§2.4, §5.4
Assessment: this pack's binding-axis rule and avoiding single scores	§2.5, §4.1, §4.3, Appendix A
Assessment: divergence diagnostic and separating scope from maturity	§1.3, §4.3, §5.2

Table 3. Mapping of specific briefing items to their corresponding report sections.

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.