

B3 — Briefing for Graduate School and Doctoral College Leadership

The five decisions on AI in research training that this office owns, and the documents that would prove each one has taken effect.

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

About this instrument

Purpose	Show the leader who owns doctoral research training where AI enters their remit, and what specific evidence proves those decisions have taken effect.
Who completes it	Deans of Graduate Studies, Associate Deans of Research Training, Directors of Doctoral Colleges and Graduate Academies, alongside the officers who manage doctoral enrollment, supervision, and examination.
Time required	15 minutes to read, plus about an hour with your team to answer Section 5 using actual institutional documents.
Related report sections	§1.1, §1.2, §1.3, §2.1, §2.2, §2.4, §2.5, §3.2, §4.1, §4.2, §4.3, §5.3, Appendices A, C, G.
Cells touched	Ten of twenty: D1-policy, D1-people, D1-process, D2-policy, D2-process, D4-policy, D4-people, D4-process, D5-people, D5-process. Dimension 3 belongs to procurement, but it places a ceiling on what this office can teach (§4.3).
Terminology	This briefing uses terms that carry the same meaning across different higher education systems. <i>Graduate students</i> and <i>doctoral students</i> refer to the researchers being trained, regardless of local enrollment terminology. <i>Advisors and supervisors</i> are the academic staff who guide them (these are two names for the same role). <i>Examiners</i> (or dissertation committee members in US usage) evaluate the final work. The <i>thesis or dissertation</i> is the written document submitted for the doctoral degree, and the <i>oral examination</i> (the viva or thesis defense) is the formal oral defense. Throughout this briefing, the shorter terms (<i>thesis</i> , <i>examiner</i> , and <i>oral examination</i>) carry these broader meanings.

Table 1. Core specifications, purpose, and implementation details of the doctoral training instrument.

1. The finding, and why it lands on this desk

Academic publishers converged on a strict no-AI-authorship rule within about three months. Research funders took longer, moving from June 2023 to April 2026. *Responsible AI in Academic Research: A Competency Framework for Research Training* surveyed fourteen major funding bodies and found that their definitions of "responsible AI use" mostly point back to host institution policy (§1.1, §1.2). The same report found that two in five of 38 top-tier doctoral universities across 15 countries and jurisdictions stop at treating AI merely as a potential plagiarism issue. The most common posture, shared by 17 of the 38 institutions, extends further to address research integrity and reproducibility. Only six universities currently reach the level of AI literacy and valid research practice, which includes explicit competency expectations for graduate students and their supervisors, as well as clear rules for examiners during the oral defense and in the final thesis declaration (§1.3, §3.2).

Framing AI primarily around plagiarism is a mistake. A large language model is an interactive, agentic research tool, much like a statistical software package or an automated transcription engine. For example, nobody treats SPSS output as plagiarism. The real question is whether the researcher's use of AI is *ethical, valid, reproducible, and transparent* (§1.2, Appendix C). Framing AI around plagiarism can lead institutions to rely on automated detection tools, and that detection fails in practice. In benchmark testing across seven commercial detectors, they achieved only 39.5 percent baseline accuracy, which dropped to 17.4 percent under simple adversarial techniques, alongside documented bias against non-native English speakers (Perkins et al. 2024, Appendix C, Class 11).

By contrast, framing AI around research integrity leads directly to building core competencies, which is what a graduate school is designed to do. The real risks in doctoral research are things like a literature review with citations that do not resolve to real papers, an empirical analysis where the AI model version and parameters were never recorded, or an interpretation of findings that the graduate student cannot explain and defend. Developing that scholarly judgment happens through advising and supervision, teaching it belongs in the research training curriculum, and testing it occurs during examination. All three of those core responsibilities sit directly with your office.

How to read your institutional profile. The assessment evaluates five dimensions (human-in-the-loop discipline, responsible use in practice, tooling that promotes responsible use, AI-literate humans, and the institutional benchmarking grid) across four organizational axes (policy, people, systems, and process) at four developmental levels: absent, nascent, established, and leading. An institution's level on any given dimension is determined by the lowest score among its four axes, not an average across them, and the assessment does not reduce your university to a single uninformative number (§2.5, §4.1, §4.3). Notably, every observable condition required to reach the *established* level on Dimension 4 (including embedding the six competencies as named curriculum outcomes, training advisors and supervisors, training examiners, and reviewing the curriculum on a regular schedule) is a direct operational responsibility of the graduate school (Appendix A).

2. The five decisions this office owns

2.1 Doctoral milestones: where AI use enters the record. The *process* axis of the framework in the report identifies key milestones across a doctoral candidature: the first progress review, mid-program review, completion seminar, thesis submission, oral examination, examiner report adjudication, and post-thesis publication and dissemination (§2.5). You need to decide at which of these milestones AI use is formally recorded and reviewed, and who conducts that review. Relying on a single disclosure question on the final thesis submission form surfaces three or four years of student practice far too late to correct any problems. Reaching the *established* level on Dimension 1 requires that the boundary between automated assistance and scholarly judgment is taught in core methods courses, reinforced through ongoing advising and supervision, and referenced in examiner training (Appendix A).

2.2 Supervisor development: the binding constraint. Graduate students cannot acquire research competencies that their advisors and supervisors do not understand or model. For this reason, professional development for advisors and supervisors is the primary bottleneck on achieving institutional maturity on Dimension 4 at scale. This training represents a recurring annual operational cost as AI model versions, tool classes, and regulatory standards evolve, and the report places that cost directly in the graduate school budget (§5.3). You need to decide whether completing this training is a mandatory condition of supervisor registration and re-registration, or merely a voluntary offering. Only making it mandatory allows an institution to answer the question "how many of our advisors and supervisors are trained?" with a concrete, verifiable number.

2.3 The training curriculum: six competencies as named outcomes. You need to decide whether the six competencies in Section 4 are formally established as *named learning outcomes*, overseen by a designated curriculum lead with a published review cycle, or left as an optional workshop that only self-selected students attend. Only embedding them as formal curriculum outcomes satisfies the criteria for the *established* level on Dimension 4 (Appendix A). These competencies are designed to build upon and reinforce existing standards for scientific reproducibility and methodological validity, rather than replace them (§2.4).

2.4 Examination policy: both sides of the oral examination. On the student side, the doctoral candidate "must be able to describe and defend any use of generative AI, as well as the contents of the thesis during their final oral examination" (following the University of Toronto School of Graduate Studies formulation, §5.3). On the examiner side, examiners must not upload any portion of a student's thesis into a generative AI tool, and must not use external AI detection software when assessing the work (following King's College London, *Generative AI: Guidance for doctoral students, supervisors and examiners*, §3.2). Examiner judgment during the oral defense is strictly a human responsibility, because the oral examination serves as the ultimate verification mechanism for the entire research training pipeline (Appendix G).

2.5 The supervisor-student relationship: documented or devolved. An *absent* rating on Dimension 1 is defined as leaving the boundary between automated labor and human judgment up to individual advisors and supervisors without institutional guidance (Appendix A). In practice, that default means two doctoral students in the very same program may face completely different rules, and the student ultimately bears the consequences during their examination. You need to decide whether the agreement on AI use between a student and their supervisory team is formally documented, dated, and revisited across the degree, or left to be reconstructed from memory when questions arise later.

3. The supervisor-student AI-use agreement (G1)

The **G1** template provides a practical mechanism to record which specific research tasks AI tools may assist with and which must remain strictly human tasks for a given doctoral project. It is completed collaboratively by the graduate student and their advisors or supervisors, drawing directly on the report's taxonomy of research labor versus human judgment (Appendix G). This agreement is neither a bureaucratic permission slip nor a disciplinary misconduct form. Rather, it is a transparent record that protects both the student and the supervisory team by making their expectations explicit and defensible. I recommend integrating this agreement across four standard milestone touchpoints.

Touchpoint	Why here	What changes
Program start (within [NUMBER] weeks of enrollment)	Before habits form	First completion, signed by the student and each advisor or supervisor, filed in the doctoral student record
First formal progress review (confirmation or upgrade)	Methods now specified	Revised against the actual research design and methodology. The review committee or panel confirms that the agreement was formally reviewed, rather than mandating specific research choices.
Mid-program review	An analysis chapter exists	Revised to account for any new tool classes or AI model versions used during data analysis. Any AI-assisted analyses must include their corresponding model and parameter records.

Touchpoint	Why here	What changes
Thesis submission	Feeds the thesis declaration	The final version is attached to the formal thesis declaration. Examiners receive this final disclosure statement, rather than the earlier working drafts.

Table 2. Milestone touchpoints, rationales, and documentation updates for the supervisor-student AI agreement.

Trigger rules between milestones. The agreement should be revisited whenever there is a change of advisor or supervisor, a substantive change in research methodology, or a shift into a new mode of AI use. The report identifies four distinct modes of AI interaction: *search*, *co-author*, *validator*, and *tutor* (§2.2).

Drafting note. You will need to decide locally who countersigns this document and where the official record is maintained, whether that is within the central student records system, recorded in supervisory panel minutes, or filed in the graduate school office. The critical requirement is that the record remains easily retrievable at the time of the oral examination three or four years later by administrators or examiners who were not present during the original discussions.

4. The training program, built from the six competencies

Competency (§2.4)	What the graduate student must be able to do	Where it is taught	Evidence it landed
Citation verification	Resolve every citation in a chapter to the primary source and produce the check record	Methods core, before the first progress review	A completed verification record is filed with progress review materials, and examiners may spot-check citations during the oral examination.
Model-and-parameter specification	State model version, temperature, seed where exposed, the prompt, and a re-run stability test for any AI-mediated analysis	Analysis module	The methods chapter contains the full specification, which the advisor or supervisor signs off on during chapter approval.
Prompt-as-fork-in-the-garden discipline	Commit prompts and the analysis plan in advance and report prompt sensitivity	Research-design module, with pre-registration	Prompts are documented in the pre-registration record, and the thesis includes an appendix reporting prompt sensitivity.
Model-heterogeneity in adversarial review	Route a critique through a model from a different frontier-model family than the one that produced the argument	Adversarial-review workshop, mid-program	The review record names both model families and what changed as a result
Sycophancy detection and human-as-verifier discipline	Show one instance where the tool agreed with their framing, and what they did about it	Same workshop	A worked instance presented at the completion seminar

Competency (§2.4)	What the graduate student must be able to do	Where it is taught	Evidence it landed
Structured failure-mode reporting	Report protocol, threat model, instrumentation and failure-discovery rate together	Writing-for-publication module	The thesis chapter or submitted manuscript carries the report

Table 3. Graduate student competencies, learning objectives, instructional settings, and required verification evidence.

Advisors, supervisors, and examiners should be trained across these same six competencies, though at different levels of depth. Advisors and supervisors need sufficient mastery to model these practices and guide their students, while examiners need to understand them to verify compliance during assessment. Initiatives like KU Leuven's "Responsible and Open research learning path for PhD researchers" and the University of Helsinki's MOOC sequence represent the closest current examples of curriculum integration. However, the report found that as of 2026, no surveyed institution has yet implemented all six competencies as a fully integrated, mandatory training program across its doctoral curriculum (§4.2).

5. Five questions to ask your office

Answer each question using actual documentation rather than informal recollection. If the supporting document cannot be produced within an hour, record the item as unresolved or not demonstrated and set a deadline to locate it. Record "No" only when documentary evidence establishes that the requirement is not met.

Q1. Which research tasks in our doctoral program may AI assist with, and which must remain strictly human? Where is that written, and who approved it?

☐ Yes ☐ Partial ☐ No. *Required evidence: the graduate school AI policy, the specific clause defining task-level boundaries, and the approving [COMMITTEE] along with the date of approval.*

Q2. At which milestones in the doctoral program does AI use enter the formal record, and who reviews it?

☐ Yes ☐ Partial ☐ No. *Required evidence: official milestone forms and supervisory panel templates showing where AI use is captured, along with one sample record from a currently enrolled student.*

Q3. How many currently registered advisors and supervisors completed AI-in-research training in the last twelve months, and where is that recorded?

☐ Yes ☐ Partial ☐ No. *Required evidence: the official supervisor register showing completion dates, the training curriculum outline, and the institutional rule linking training completion to supervisor registration.*

Q4. Which of the six competencies are named learning outcomes in our doctoral training program, and in which specific module are they taught?

☐ Yes ☐ Partial ☐ No. *Required evidence: the graduate program handbook or official syllabus listing competencies by module, the designated curriculum owner, and the published schedule for curriculum review.*

Q5. What written instruction do our examiners receive regarding AI, covering both the student's permitted use and the examiners' own conduct?

☐ Yes ☐ Partial ☐ No. *Required evidence: the examiner appointment pack or briefing guidelines, the policy clause requiring students to defend AI use during the oral examination, and the explicit rules protecting thesis confidentiality and prohibiting the use of unapproved AI detection tools.*

If your office can answer "Yes" to all five questions with verifiable documents attached, your institution is operating at the *established* level across most of the areas within your remit. On the other hand, answering "No" across all five is not a sign of institutional failure. Rather, it simply reflects the current state of higher education, placing you alongside the 32 of 38 universities in the report's international sample whose published policies do not yet reach Class D scope (§1.3, §3.2).

6. What to commission next

To put these recommendations into practice, I suggest commissioning several specific instruments from the Instats AI Readiness Pack:

First, adopt **G1** to document supervisory expectations across the four milestone touchpoints described in Section 3. Second, run **A3**, the supervisor survey, to examine self-reported current supervisory practices and size your professional development needs before allocating training budgets. Third, deploy **A2**, the doctoral student survey, to uncover any gaps between what faculty believe they have agreed to and what students report doing. Fourth, conduct **A1**, the full institutional self-assessment, when [INSTITUTION NAME] is ready to evaluate all twenty cells of the maturity grid with [COMMITTEE] by [DATE]. Then complete **A4**, the gap-analysis worksheet and action ladder, to turn the A1 results and survey evidence into concrete next steps.

Your operational plan should also incorporate relevant regulatory deadlines where applicable to your jurisdiction. For institutions operating within or serving the European Union market, the EU AI Act Article 27 Fundamental Rights Impact Assessment obligation takes effect on **2 August 2026**. For institutions subject to Australian regulation, the Australian Privacy Act 1988 automated decision-making transparency requirement comes into force on **10 December 2026** (§2.5). Neither regulation applies universally, but you should designate a responsible officer to oversee whichever compliance requirement applies to your institution.

At present, no university in the report's research sample publicly reports its scoring against this framework, largely because these benchmarking standards have only recently been established (§4.2). A graduate school that takes the initiative to benchmark its practices and systematically address these gaps will be stepping onto a developmental rung on that ladder that none of the sampled institutions have yet achieved.

Crosswalk: briefing item to report section

Briefing item	Report section
§1 Publisher and funder convergence, with funders pointing back to host institutions	§1.1, §1.2
§1 The 38-university baseline and class distribution	§1.3, §3.2
§1 Plagiarism as a flawed framing, and standards for ethical, valid, reproducible, and transparent use	§1.2, Appendix C
§1 Detection tool accuracy and bias	Appendix C, Class 11
§1 Framework axes, maturity levels, and avoiding a single score	§2.5, §4.1, §4.3
§1 Dimension 4 <i>established</i> conditions	§2.4, Appendix A, Dimension 4
§2.1 Doctoral milestones and the process axis	§2.5, Appendix A, Dimension 1
§2.2 Supervisor development as the primary constraint	§2.4, §4.2, §5.3
§2.3 Curriculum as named outcomes	§2.4, §5.3, Appendix A, Dimension 4
§2.4 Examination policy for students and examiners	§2.1, §3.2, §5.3, Appendix G
§2.5 The <i>absent</i> definition on Dimension 1	§2.1, Appendix A, Dimension 1
§3 G1 touchpoints and filing	§2.1, §2.5, Appendix G
§3 Trigger rules and the four AI-use modes	§2.2
§4 All six competencies	§2.4
§4 Curriculum integration examples	§4.2
§5 Five institutional questions (task demarcation, milestone records, supervisor training, named outcomes, examiner instructions)	§2.1, §2.4, §2.5, §5.3
§6 The two regulatory deadlines and their jurisdictions	§2.5, Appendix A, Dimension 5
§6 Dimension 5 as a forward-looking standard	§4.2

Table 4. Mapping of specific briefing items to their corresponding sections within the report.

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.