



Fast AI, Slow Policy

Governing Generative AI in Higher Education

Insights from a national ACPHIS study

46 universities • Australia and Aotearoa New Zealand

[ACPHIS GenAI Project:
Welcome Survey](#)





Acknowledgement of Country

ACPHIS acknowledges the Traditional Custodians of the lands on which we meet, work and learn. We pay our respects to Elders past and present, and extend that respect to all Aboriginal and Torres Strait Islander peoples joining us today.

Meet the Project Team



Project Leads



Lubna



Deborah

Technical Lead



Lemai

Chief Investigators and Research Team



Meena



Rohini



Thuc



Kranthi




Jiahe



Sophia



Mauricio




Bernadette



Mimi



Erwin



Antonett



Soonja



Winyu



Amara



Ana



Shashi



Saima



Ruhi

**IS Community and Capability building
Voluntary efforts**

Working collaboratively with a large team of 21 academics - Senior, MCRs, ECR, HDR students

<https://www.acphis.org/genai-policy-review>




QR code for the project website

Project Aims

The project has two overarching aims:

1. To explore how Generative AI is currently being used across the sector, including institutional policies, teaching practices, research workflows, and ethics and governance approaches; and
2. To develop evidence-informed recommendations that can guide institutions, policymakers, and professional bodies in fostering responsible, innovative, and inclusive AI adoption.

Generative AI & the Tertiary Sector Project Milestones (2025-26)



Publications

- Alam, S. L., Hyland-Wood, B., Chen, J., Nguyen, L., Bunker, D., Sadiq, S., & Abedin, B. (In press). Generative AI and the Tertiary Sector: Current Issues, Key Opportunities and Risks for Institutions and Policymakers. Communications of the Association for Information Systems, 59, pp-pp. Retrieved from <https://aisel.aisnet.org/cais/vol59/iss1/25>
- Alam S, Nguyen L, Bunker D, Hyland-wood B, Fielt E, Marrone M, Chen J, Tsai M, Yang W, Balapumi R, Nguyen T, Hol A, Chinthammit W, Addanki K, Jha M, Yeom S, Mendoza A, Duan S, Charles S, Qutab S, Bajaj R. (2026). Governing Generative AI in Higher Education: Insights from a National ACPHIS Study. HERDSA 2026. <https://conference.herdsa.org.au/2026/program>

Presentations

- Upcoming presentation at ACPHIS SGM on 31 July 2026 from 2-3 pm
- Panel at ACIS 2025 on “Generative AI and the Tertiary Sector: Current Issues, Key Opportunities and Risks for Institutions and Policymakers. Communications of the Association for Information Systems”.

Panelists: Alam, S. L., Hyland-Wood, B., Chen, J., D., Sadiq, S., & Abedin,

Facilitator: Bunker, D.



01 — Background

Why this matters now



02 — Three Studies

Educational practice ·
Research practice ·
Institutional governance



03 — Synthesis and Implications

What institutions can do
next

Why this matters now

GenAI is evolving faster than institutions can learn, govern and adapt.

FAST AI

Rapid capability gains

- widespread experimentation
- expanding teaching and research uses
- uncertain boundaries of acceptable use

SLOW POLICY

Institutional adaptation takes time

- policy responses evolving unevenly
- capability-building still emerging
- trust, quality and accountability at stake

This project asks not only what GenAI can do, but how institutions can respond responsibly.

What do we already know from existing research?



- Existing Australian policy reviews consistently suggest that:
- Universities have responded rapidly to Generative AI.
- Institutional policies have largely converged around academic integrity and TEQSA principles.
- Assessment redesign is emerging, but discipline-specific guidance remains limited.

Atif, A., Jha, M., Richards, D., & Cibin, A. (2026). Institutional alignment with TEQSA's generative AI principles in Australian universities: A national document analysis with implications for ICT education. *Australasian Journal of Educational Technology*. <https://doi.org/10.14742/ajet.11001>

The sector challenge

Universities are balancing competing priorities rather than solving a single problem.

INTEGRITY

assessment, disclosure and academic standards

CAPABILITY

AI literacy, judgement and responsible use

RESEARCH PRACTICE

epistemic authority, authorship and accountability

PRIVACY / IP

data governance, copyright and institutional risk


INNOVATION

experimentation without loss of trust

The policy task is therefore dual: enable innovation while safeguarding trust, quality and accountability.

Three complementary studies

A public-facing pathway from educational practice to institutional governance.



Together, these studies show how capability, agency and governance evolve together in responsible AI adoption.

STUDY 1

Understanding GenAI in higher education

An umbrella review through a Functional–Relational–Responsible lens

STUDY 1

Method and review architecture

20 literature reviews

3 Integrated evidence perspectives
Functional-Relational-Responsible

1
Collection



2
Screening



3
Extraction



4
Quality appraisal



5
Coding & synthesis

Objective

Develop and refine a socio-technical framework that explains not only what GenAI enables, but how it reshapes educational relationships and responsible practice.

Three interdependent dimensions of GenAI in higher education

FUNCTIONAL

What GenAI enables

- augmentation
- automation
- productivity

RELATIONAL

How people engage

- dialogue
- collaboration
- agency

RESPONSIBLE

How use is governed

- ethics
- transparency
- accountability

The dimensions are co-constituted: capability shapes interaction, interaction reshapes roles, and governance establishes legitimate boundaries.

GenAI augments academic work across roles

The dominant pattern is augmentation, not replacement.

STUDENTS

- intelligent tutor
- study and revision assistant
- writing and feedback support

TEACHERS

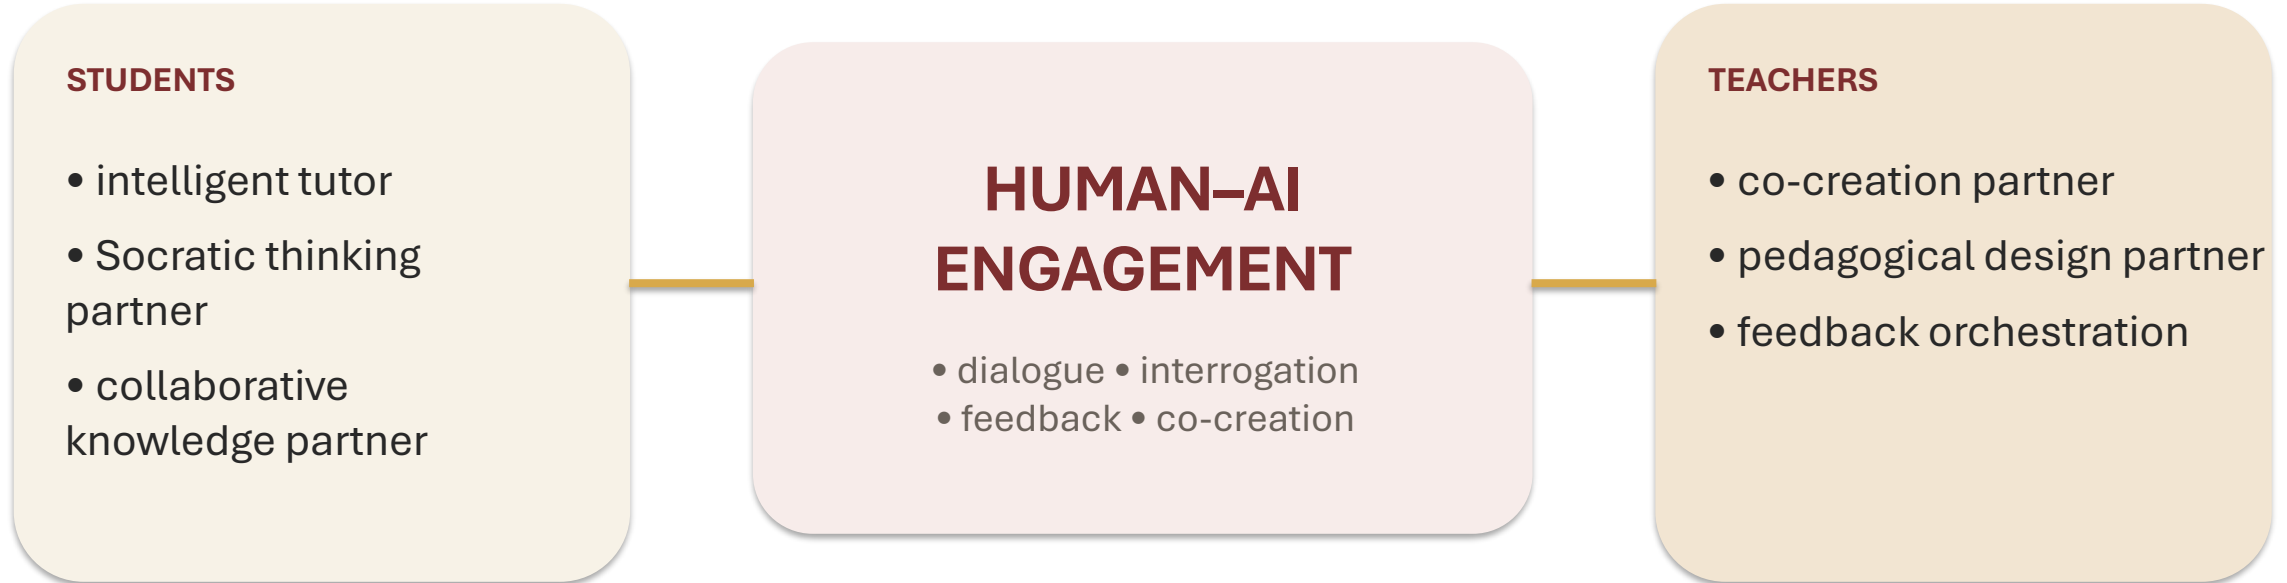
- content creation partner
- assessment and feedback support
- learning-design partner

RESEARCHERS

- literature discovery
- summarisation and editing
- research productivity support

Human expertise shifts from **performing routine tasks** to **designing, evaluating and validating** AI-assisted work.

The important shift is from tool use to human–AI engagement



Agency shifts from producing answers to questioning, evaluating, refining and validating AI outputs.

Ethical concerns are sociotechnical—not isolated technical failures

GENAI-CENTRIC

- bias & epistemic distortion
- digital inequality
- privacy & governance
- opacity & explainability

STUDENT-CENTRIC

- academic integrity
- learner agency
- psychosocial risks
- authorship & misuse

TEACHER-CENTRIC

- professional identity
- emotional & cognitive labour
- accountability & liability

The literature foregrounds student risks; teacher-centred consequences remain comparatively underexamined.

Each dimension carries a distinct logic—and a distinct tension

FUNCTIONAL

Primary logic

Task–technology fit supports structured academic work.

Emerging tension

As complexity rises, fit weakens and oversight becomes more important.

RELATIONAL

Primary logic

Dialogue, feedback and co-creation reshape learning.

Emerging tension

Agency shifts towards questioning and validating AI outputs.

RESPONSIBLE

Primary logic

Governance establishes accountability and legitimacy.

Emerging tension

Greater autonomy amplifies ethical and institutional risk.

STUDY 1 SYNTHESIS

From capability to educational practice

Functional explains what GenAI enables

Relational shows how engagement reshapes agency

Responsible sets the conditions for legitimate use

Together, the three dimensions explain not only what GenAI does, but how it reshapes educational practice, relationships and governance.

STUDY 2

Calibrating Human–AI co-agency in research

Where can GenAI assist—and where must human epistemic and ethical authority remain non-delegable?

STUDY 2

Method and analytical objective

REVIEW PROCESS


- Systematic literature review and PRISMA-guided screening
- Mapping of GenAI use across the research lifecycle
- Assessment of epistemic stakes
- Classification of Human–AI configurations
- Interpretive synthesis across activities and scholarly positions

OBJECTIVE

Examine how GenAI is positioned across the research lifecycle—and identify where human epistemic authority is treated as non-delegable.

The unit of analysis is not the tool alone, but the relationship between activity, epistemic stakes and allocated authority.

Where should Human–AI boundaries be drawn?



Epistemic stakes increase as activities move closer to knowledge creation, interpretation and final judgement.

The literature agrees that humans retain responsibility for knowledge claims, but differs on how far substantive work may be delegated.

Activity → epistemic stakes → Human–AI configuration

1 Research activity

Where does GenAI enter the process?

2 Epistemic stakes


How strongly does the activity shape claims, interpretation or judgement?

3 Human–AI configuration

How should agency, oversight and decision authority be allocated?

Human–AI co-agency should be calibrated to epistemic stakes—not to technical capability alone.

Four epistemic zones of Human–AI co-agency



The emerging pattern suggested four epistemic zones.

From capability to agency calibration

Epistemic agency

Who creates, interprets and validates knowledge?

Ethical agency

Who remains responsible and authorised to decide?

Governance implication

Calibrate AI use to epistemic and ethical stakes—not technical capability alone.

The closer an activity moves towards interpretation, justification and final scholarly judgement, the stronger the case for retained human authority.

STUDY 3

Generative AI policy across Australian and New Zealand universities

Control at the ceiling, agency in motion



46 universities

38 Australia · 8 New Zealand

807 webpages

Publicly available institutional evidence

15 indicators

Five per analytical dimension

 $\kappa = 0.848$

Final consensus dataset

1

**Collection**

2

**Screening**

3

**Extraction**

4

**Quality assurance**

5

Coding & analysis

Quality assurance was embedded throughout the review—not added at the end.

Three complementary dimensions of GenAI governance

CONTROL

How universities regulate use

- Rules and restrictions
- Disclosure and accountability
- Governance and oversight

AGENCY

How universities enable use

- Capability building
- Critical guidance
- Innovation and participation

DEFICIT FRAMING

How policies position users

- Users framed as risks
- Misconduct and surveillance
- Restrictive default assumptions

Together, the dimensions show how policy establishes boundaries, builds capability and signals perceived risk.

Strong control, emerging agency



4.54 / 5

CONTROL

Strong and widespread

3.63 / 5

AGENCY

Present but uneven

2.33 / 5


DEFICIT FRAMING

Moderate, not dominant

Control structures are firmly established, while agency is still finding room to take flight.

The pattern is not blanket prohibition: regulation is mature, while capability-building remains uneven.

Patterns beyond the averages



The averages describe the landscape.

The heatmap reveals how policy elements are configured across institutions.

Heatmap reveals the maturity of policy discourse from publicly available sources

A computational method for measuring how universities govern generative AI



What 1,162 provisions across 35 universities reveal about how the sector governs GenAI.

807

pages collected
46 universities, AU + NZ

human-coded →
policy? · relevant?

126

governance documents
35 institutions · 29 AU · 6 NZ

segment →
provisions · sentences

1,162

AI-relevant provisions
the unit of analysis

Findings: GenAI provisions concentrate in dedicated documents, and governance language follows what is governed, not who governs.

Next step: human validation of the content measures (see back-up).

GenAI governance concentrates in guidance, not enforceable policy


LLMs classified each document's genre and scope, logging evidence and a confidence rating per label (human validation next); provision counts are deterministic keyword matches.

Genre \ Scope	Dedicated	Embedded
Guideline	16	2
Webpage	18	9
Policy	21	39
Procedure	3	18
All genres	58	68

WHERE THE 1,162 PROVISIONS ARE

Dedicated  **974 provisions**

58 docs

Embedded  **187 provisions**

68 docs


84% of the 1,162 provisions are in the 58 dedicated documents.

Implication: the sector governs GenAI through soft regulation: expectations concentrate in guidance, and little is yet codified in enforceable policy, consistent with symbolic rather than substantive adoption.

Governance language follows what is governed, not who or where



Do policies governing the same topic use the same language, whoever wrote them? (58 documents with at least three provisions; 5,000 label shuffles.)



*Within- minus between-group median cosine similarity · *** $p=0.0002$ · * $p<0.05$, uncorrected · n.s. = not significant*

Finding. University and country predict nothing about policy language: both effects are indistinguishable from chance, and both labels come from metadata, not from the text.

SYNTHESIS AND IMPLICATIONS

What the three studies tell us together

Cross-study synthesis

Study 1 Educational practice

GenAI is most valuable when it augments human capability, reshapes participation and remains subject to responsible oversight.

Study 2 Research practice

Human–AI collaboration should be calibrated to epistemic stakes, with stronger human authority as work moves closer to interpretation, justification and final judgement.

Study 3 Institutional governance

Universities have established strong control mechanisms, but capability-building, participation and enabling guidance remain uneven.

**Capability, agency and governance must develop together.
Strengthening one without the others may produce incomplete institutional responses.**

From rules to responsible institutional capability

- 1. Build capability, not only compliance**

Policies need to be supported by practical guidance, AI literacy, exemplars and opportunities to develop informed judgement.

- 2. Calibrate governance to context and risk**

Permitted use, oversight and human responsibility should reflect the educational or epistemic stakes of the activity.

- 3. Preserve meaningful human agency**

Students, educators and researchers should remain active participants in questioning, evaluating and validating AI-assisted work.

- 4. Enable accountable experimentation**

Institutions need safe mechanisms for piloting, evaluating and refining GenAI practice rather than relying only on prohibition or static rules.

- 5. Treat policy as an ongoing learning process**

GenAI governance requires review, participation and adaptation as technologies, practices and evidence evolve.

Fast AI requires more than fast rules.



Responsible institutional adoption emerges when:

capability enables
informed use

agency preserves
human judgement and
responsibility

governance provides
legitimate boundaries,
accountability and
learning mechanisms

**Institutions will not govern GenAI responsibly through policy alone.
They must also become capable of learning responsibly.**

SYNTHESIS

Governance must do more than contain risk

Control

builds the fences

Agency

grows the wings

Deficit framing

negotiates risk and
trust

But if only the fences remain... who will learn to fly?

References

- An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22(1), 10.
- Atif, A., Jha, M., Richards, D., & Cibir, A. (2025). *Navigating the AI assessment horizon: A foresight-driven framework for ICT education*. In *Proceedings of the 36th Australasian Conference on Information Systems (ACIS 2025)*. AIS Electronic Library (AISeL). <https://aisel.aisnet.org/acis2025/>
- Atif, A., Jha, M., Richards, D., & Cibir, A. (2026). Institutional alignment with TEQSA's generative AI principles in Australian universities: A national document analysis with implications for ICT education. *Australasian Journal of Educational Technology*. <https://doi.org/10.14742/ajet.11001>
- Chan, C. K. Y. (2023). "A comprehensive AI policy education framework for university teaching and learning." *International Journal of Educational Technology in Higher Education* 20(1).
- Dai, Y., Lai, S., Lim, C. P., & Liu, A. (2025). University policies on generative AI in Asia: promising practices, gaps, and future directions. *Journal of Asian Public Policy*, 18(2), 260-281. <https://doi.org/10.1080/17516234.2024.2379070>
- Davison, R. M., Laumer, S., Tarafdar, M., & Wong, L. H. (2023). Pickled eggs: Generative AI as research assistant or co-author?. *Information Systems Journal*, 33(5). doi. 10.1111/isj.12455
- Gregor, Shirley (2024) "Responsible Artificial Intelligence and Journal Publishing," *Journal of the Association for Information Systems*, 25(1), 48-60. <https://aisel.aisnet.org/jais/vol25/iss1/11>
- Jiang, Y., Xie, L., & Cao, X. (2025). Exploring the effectiveness of institutional policies and regulations for generative AI usage in higher education. *Higher Education Quarterly*, 79(4), e70054.
- Kane, Gerald C. and Tarafdar, Monideepa (2026) "Writing an "A" Paper in the Age of GenAI: A Guide to Responsible Augmentation," *Journal of the Association for Information Systems*, 27(4), 831-845. <https://aisel.aisnet.org/jais/vol27/iss4/8>
- Kankanhalli, Atreyi (2024) "Peer Review in the Age of Generative AI," *Journal of the Association for Information Systems*, 25(1), 76-84. <https://aisel.aisnet.org/jais/vol25/iss1/9>
- Ngwenyama, Ojelanki and Rowe, Frantz (2024) "Should We Collaborate with AI to Conduct Literature Reviews? Changing Epistemic Values in a Flattening World," *Journal of the Association for Information Systems*, 25(1), 122-136. <https://aisel.aisnet.org/jais/vol25/iss1/5>
- Pan, S. L., Nishant, R., Tuunanen, T., & Nah, F. F. H. (2023). Literature review in the generative AI era: How to make a compelling contribution. <https://doi.org/10.1016/j.jsis.2023.101788>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S.,...Moher, D. (2021). The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Paré, G., Trudel, M.-C., Jaana, M., & Kitsiou, S. (2015). Synthesizing information systems knowledge: A typology of literature reviews. *Inf. Manag.*, 52, 183-199.
- Smith, S. M., Tate, M., Freeman, K., Walsh, A., Ballsun-Stanton, B., & Lane, M. (2025). A university framework for the responsible use of generative AI in research. *Journal of higher education policy and management*, 1-20.



THANK YOU

Acknowledgement

Meet the Project Team



Project Leads



Lubna



Deborah

Technical Lead



Lemai

Chief Investigators and Research Team



Meena



Rohini



Thuc



Kranthi



Jiahe



Sophia



Mauricio



Bernadette



Mimi



Erwin



Antonett



Soonja



Winyu



Amara



Ana



Shashi



Saima



Ruhi


IS Community and Capability building
Voluntary efforts
Working collaboratively with a large team of 21 academics - Senior, MCRs, ECR, HDR students

Thank You for listening!

Any Questions?

We'd like to hear from you





ABOUT

ACTIVITIES

AWARDS

ACPHIS GenAI Policy Review Project

Governing Generative AI in the Tertiary Sector: Insights from Australia and New Zealand
2025/HE001075

