ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

When connection requires refusal: AI refusal and accountable technology design in higher education

Jason Zagami

Griffith University

This non-empirical conceptual analysis examines how artificial intelligence (AI) is becoming woven into curriculum production, assessment, feedback, student support, academic integrity, accessibility and institutional decision-making. These systems are often justified through connection: faster feedback, personalised support, scalable services and more efficient learning design. Yet technical connection is not the same as educational connection. AI may leave the visible forms of university learning intact while weakening academic authorship, assessment validity, student trust, staff agency, care, equity and institutional legitimacy. The paper develops refusal-capable design: the distributed institutional capacity to decide when AI should be adopted, constrained, redesigned or refused. It integrates connected learning, responsible innovation, sociotechnical studies of educational technology, care ethics, justice-oriented educational judgment and technology non-use scholarship, using selected public examples illustratively rather than as an empirical sample. Six heuristic refusal zones organise recurring legitimacy tensions, and the Distributed Refusal-Capable University Design Cycle translates those tensions into accountable redesign. The contribution is a practical and theoretical account of refusal as a legitimate design outcome rather than simple opposition to AI.

Keywords: artificial intelligence, higher education, educational technology, refusal, backlash, learning design, conceptual analysis

Introduction

Artificial intelligence (AI) is no longer peripheral to university learning. It now appears in course production, learning management systems, writing assistants, assessment workflows, feedback tools, student support chatbots, accessibility technologies, academic integrity systems, analytics dashboards and administrative decision-making. Some systems are centrally procured. Others arrive through enterprise platforms, individual experimentation or student practice. The result is a distributed reorganisation of university learning, work and governance (Miao & Holmes, 2023; Sharples, 2023; Williamson & Eynon, 2020). This reorganisation is usually justified through connection. AI is said to connect students to faster feedback, academics to design support, learners to personalised resources, support services to students after hours, and institutions to data about engagement, risk and performance. These promises are attractive because universities face large cohorts, diverse student needs, delayed feedback, workload intensification, accessibility obligations, academic integrity concerns and demands for flexible learning (Selwyn, 2022; Tertiary Education Quality and Standards Agency, 2026; University of Queensland, 2026). Yet connection, by itself, does not make a practice educational. Technical connection refers to access, availability, automation, interoperability, data flow, platform integration and response speed. Educational connection depends on relationships and practices that support learning, belonging, disciplinary formation, academic authorship, care, agency, assessment validity, professional responsibility and public purpose (Biesta, 2015; Ito et al., 2013; Selwyn, 2022). Educational thinning is the substitution of technically adequate outputs or interactions for the academic, relational and ethical work through which those forms acquire educational meaning. It is not automation, reduced contact time or poor-quality technology. Thinning occurs when the visible form remains credible – curriculum, feedback, support, assessment or connection – while responsibility, disciplinary judgment, reciprocal relationship, valid evidence of learning or meaningful contestability is diminished. Curriculum may remain while academic authorship is diluted; feedback may remain while academic responsibility is obscured; support may remain while care is automated; assessment may remain while evidence of learning becomes uncertain; and connection may remain while trust is weakened (Bayne, 2015; Biesta, 2015; Noddings, 2013). The guiding question is: how can higher education interpret and respond to resistance to AI in ways that protect educational purpose without

treating all objection as anti-technology? Backlash is used narrowly for visible, collective or institutionally consequential contestation that attracts organisational or public response. Read through responsible innovation and voice theory, it can signal a perceived gap between technological implementation and educational purpose (Hirschman, 1970; Stilgoe et al., 2013). The analysis advances three contributions: AI backlash is theorised as a legitimacy signal, refusal-capable design is developed as an institutional response, and the Distributed Refusal-Capable University Design Cycle is proposed as a framework for accountable technology redesign.

Methodological positioning

This paper is a non-empirical conceptual analysis illustrated through purposively selected public examples. It does not estimate the prevalence of AI refusal or claim to represent the full range of institutional practice. The examples were identified within the public reporting, university guidance, regulator material and union documents cited in the paper. They were retained when they were publicly traceable, involved an articulated stakeholder concern or institutional response, and illuminated a distinct higher-education function. Selection therefore sought conceptual range across curriculum, assessment, support, data governance, workload and equity, not statistical or sectoral representativeness. The analysis proceeded through iterative concept mapping rather than qualitative coding of a dataset. Each example was read against five bodies of scholarship – connected learning, sociotechnical studies of educational technology, responsible innovation, care ethics and educational judgment – by asking four consistent questions: what educational purpose is claimed; which relationship or responsibility is at stake; what role AI performs; and what response stakeholders seek. Similar tensions were grouped and tested against the literature. The resulting six refusal zones are proposed heuristic domains: analytically derived categories for organising discussion, not empirical findings, a complete taxonomy or mutually exclusive classifications. The five-stage design cycle was then developed abductively by linking the recurring questions in the examples to responsible innovation's commitments to anticipation, reflexivity, inclusion and responsiveness. Educational purpose and human connection specify what should be sustained; AI mediation makes the intervention inspectable; distributed judgment broadens whose knowledge counts; and accountable redesign or refusal makes non-adoption, constraint and revision legitimate outcomes. The examples illustrate why these stages matter, but do not validate the sequence or establish causal effectiveness. Because the analysis used only publicly available documents and did not involve human participants, private records or interaction with identifiable individuals, institutional human-research ethics review was not required. Public examples are used to illuminate concepts, not to infer the intentions of individuals or evaluate the institutions concerned. The Distributed Refusal-Capable University Design Cycle is therefore offered as a design proposition, not a validated instrument. Its purpose is to support institutional dialogue, professional learning, technology governance and learning design decisions. Further empirical work should test the usability of the zones and cycle across disciplines, institutions and stakeholder groups, including whether they reveal overlooked harms, improve contestability and change adoption decisions.

Theoretical framing: From connection to refusal-capable design

AI adoption in higher education is often framed as capability, compliance or risk. Institutions ask how AI can be used, how students should disclose use, how staff should be trained, how assessment can be secured, and how misuse can be detected. These questions are necessary, but they do not explain the depth of concern now visible in the sector. The deeper issue is legitimacy. AI unsettles who is authorised to teach, assess, advise, support, classify, generate curriculum and make decisions about students. Connected learning theory frames learning as socially embedded, interest-driven and connected to opportunity (Ito et al., 2013). In higher education, connection is also disciplinary, institutional, professional and civic. Students connect with fields of knowledge, academic communities, peers, teachers, professional standards and public purposes. University learning is not access to information alone. It involves induction into ways of knowing, practising, questioning, creating and contributing. AI complicates connected learning because it can generate the surface forms of connection without sustaining their depth. A chatbot can simulate dialogue. A feedback generator can simulate academic response. A content generator can simulate curriculum. A dashboard can simulate care through risk identification. The problem is not simulation alone, but the institutional temptation to treat simulated connection as equivalent to educational relationship (Bayne, 2015; Sharples, 2023). Biesta's account of educational purpose clarifies the problem. Education involves qualification, socialisation and

subjectification: knowledge and skill acquisition, participation in social and disciplinary practices, and the formation of persons capable of agency and responsibility (Biesta, 2015). AI may assist parts of qualification. It is less able to sustain the relational and normative dimensions through which students become members of disciplines, professions and publics. Educational technologies are never just tools. They reorganise practice, labour, authority, visibility and accountability. Sociotechnical studies of educational technology show how platforms, data systems and digital infrastructures shape what institutions value and they can see (Selwyn, 2022; Williamson & Eynon, 2020). AI intensifies this dynamic because it participates in producing the artefacts and judgments of education. Responsible innovation emphasises anticipation, reflexivity, inclusion and responsiveness (Stilgoe et al., 2013). In university AI adoption, these commitments require institutions to anticipate consequences, examine their assumptions, include affected stakeholders and respond to concern. Refusal is one way stakeholders indicate that an innovation is not yet educationally legitimate.

Refusal-capable design extends rather than replaces responsible innovation, human-centred design and technology non-use scholarship. Responsible innovation supplies procedural commitments; human-centred approaches foreground needs and usability; and non-use scholarship shows that absence of use is socially meaningful. Refusal-capable design adds an institutional decision architecture: protected opportunities to question or decline particular uses, distributed authority across affected roles, explicit non-adoption as a legitimate design outcome, and a route for converting objections into auditable changes to governance, workload, human review and non-AI alternatives. Hirschman's distinction between exit, voice and loyalty also helps explain the institutional value of refusal (Hirschman, 1970). In universities, refusal is a form of voice. Students who petition against AI use, staff who resist an AI workflow, support teams who demand human escalation, and academics who reject detector-led misconduct processes are not necessarily rejecting the institution. They are contesting the conditions under which AI enters it. These theoretical elements converge in educational thinning. Technical connection increases access, speed, automation and scale. Educational connection requires academic responsibility, human relationship, disciplinary meaning, care, equity, valid assessment and public purpose. AI-mediated systems produce thinning when technical connection substitutes for, simulates or weakens these relations. Backlash may make thinning visible, but the two are not equivalent: thinning can occur without public contestation, and backlash can arise for reasons other than thinning (Hirschman, 1970; Star & Ruhleder, 1996; Stilgoe et al., 2013). Educational thinning can be investigated through observable indicators rather than inferred from AI use alone. Warning signs include course materials released without accountable academic approval; feedback that cannot be traced to a responsible educator; support systems without timely human escalation; assessment decisions for which no person can explain or defend the judgment; efficiency claims that omit verification and repair labour; compulsory AI use without an equivalent pathway; and affected students or staff who cannot contest the system. These indicators are prompts for inquiry, not a scoring scale: their meaning depends on educational purpose, context and consequence.

Refusal, non-use and the duty to understand

Refusal must be distinguished from non-use. Technology non-use scholarship warns against treating non-use as absence, deficit or lag (Baumer et al., 2015; Selwyn, 2003). A student who does not use generative AI may be making a principled refusal, lacking access, unsure of rules, using AI covertly or simply seeing no need. Refusal needs careful interpretation before it is counted. Refusal implies an explicit, reasoned decision not to use, or to restrict a particular use, because it conflicts with a value, responsibility or educational purpose. Non-use is the broader descriptive condition and may reflect refusal, lack of access, uncertainty, irrelevance or covert use. Resistance is the broader set of practices that challenge, reshape or slow implementation, of which refusal is one form. Backlash is the visible, collective or institutionally consequential expression of resistance. Policy caution delays or conditions action while evidence or safeguards are developed. Institutional non-adoption is an organisational decision not to procure or implement a system and may result from refusal, caution, cost or strategic priorities. Keeping these categories distinct prevents every hesitation or absence of use from being interpreted as refusal (Baumer et al., 2015; Selwyn, 2003). A refusal-capable university must also distinguish the right to refuse from the duty to understand. The right to refuse protects students and staff who choose not to use AI in personal intellectual, creative or scholarly work, provided non-use does not harm others or breach role responsibilities. Private refusal may defend authorship, privacy, attention, epistemic independence, environmental commitments or scholarly craft (Bender et al., 2021; Crawford, 2021; Selwyn, 2003). The duty to understand applies where a role carries responsibility for teaching, assessment, research, student support, procurement, governance or system administration. A course convenor need not adopt every

AI tool, but must understand enough to decide how AI affects learning outcomes, assessment validity and student guidance. A learning designer need not become an AI advocate, but must understand enough to distinguish augmentation from academic displacement. A technology administrator need not approve every feature, but must understand enough to evaluate privacy, auditability, accessibility and support implications (Miao & Holmes, 2023; Tertiary Education Quality and Standards Agency, 2026; University of Queensland, 2026). This distinction prevents two weak responses. One is compulsory adoption, where the duty to understand is misused as pressure to use AI for productivity, efficiency or institutional signalling. The other is uninformed refusal, where refusal becomes a way of avoiding systems that already shape learning and work. Mature AI literacy includes the capacity to use, refuse, restrict, audit and govern AI in context (Miao & Holmes, 2023; Tertiary Education Quality and Standards Agency, 2026).

Ethical stance: Care, justice and public accountability

The ethical stance of refusal-capable design is stronger than caution. It begins from the claim that universities have obligations to people before they have obligations to platforms. AI adoption should be judged by whether it protects learning, care, justice and public accountability, not only by whether it increases output or reduces friction (Biesta, 2013; Noddings, 2013; Selwyn, 2022). Care is the first ethical commitment. AI can route information, generate text and identify patterns, but it cannot bear responsibility for a student, a class, a disciplinary standard or a community. Care requires attentiveness to context, capacity to notice vulnerability and willingness to remain answerable for consequences. AI may support care by reducing delay or improving navigation, but it cannot replace the human and institutional responsibility to care (Noddings, 2013; Sharples, 2023). Justice is the second commitment. AI risks and benefits are unevenly distributed. Students with money, confidence, high digital access and strong prior knowledge may benefit from AI while others face confusion, suspicion, surveillance or exclusion. Multilingual students, disabled students, neurodivergent students, casual staff and professional staff working under workload pressure may carry disproportionate burdens when AI systems are poorly governed. Refusal is unevenly available as well: secure students and staff can refuse AI more safely than those who rely on AI for accessibility, language support, rapid feedback or employment security. Ethical AI adoption must ask who gains time, who loses agency, who becomes more visible to surveillance, and who repairs system failures (Baumer et al., 2015; Bender et al., 2021; National Tertiary Education Union, 2024; Selwyn, 2003). Public accountability is the third commitment. Universities are not private productivity services. They are public knowledge institutions that certify learning, steward disciplines, protect student futures and shape civic trust. If AI generates course materials, advises students, assesses work, classifies risk or mediates support, the university must be able to defend publicly the educational rationale, governance arrangements and limits placed on the system. Hidden AI use undermines that accountability because students and staff cannot know when an educational relationship has been technologically mediated (Jasanoff & Kim, 2015; Selwyn, 2022; Williamson & Eynon, 2020). The ethical test is thus not whether AI can do something, but whether the university should allow that task to be done in that way. Some educational work should remain deliberately slow, relational or difficult because difficulty is part of formation. Some decisions should remain human because responsibility cannot be delegated to a system. Some forms of refusal should be protected because they defend the conditions under which higher education remains worthy of trust (Biesta, 2013, 2015; Noddings, 2013).

The Distributed Refusal-Capable University Design Cycle

The Distributed Refusal-Capable University Design Cycle provides a way to examine AI-mediated systems before, during and after implementation. It operationalises the paper's theoretical synthesis by treating refusal as a legitimate outcome of responsible design. As shown in Figure 1, public signals and the six heuristic zones identify where educational responsibility may be strained; the five stages then structure deliberation from educational purpose to accountable redesign or refusal. Distributed stakeholders participate throughout, and observed effects return the decision to review. The first stage is educational purpose. The issue is not what the tool can do, but what educational purpose its use would serve. Universities need to ask what learning, assessment, support or institutional purpose is being addressed, and whether AI is appropriate for that purpose. The second stage is human connection. The issue is which relationships must be protected. AI may support interaction, but it can also displace relationship. The design process should identify whether the relevant connection is student-academic, student-peer, student-discipline, student-support staff, academic-

learning designer, student-community or staff-institution. The third stage is AI mediation. The issue is what AI is actually doing. Is it generating, summarising, marking, recommending, classifying, predicting, detecting, translating, advising or monitoring? What data does it process? Who can inspect or contest the output? What assumptions are built into the system? The fourth stage is distributed judgment. The issue is who sees benefit, risk or harm. Students, academics, learning designers, technology administrators, professional staff, privacy officers, accessibility teams, unions and leaders may see different consequences. Their concerns should be treated as design evidence. The final stage is accountable redesign or refusal. AI may be adopted, constrained, piloted, redesigned, disclosed, made optional, subject to human review, or refused. The decision should record its rationale, responsible owner, review date, communication plan, workload implications and escalation pathways. The cycle is not a linear checklist: implementation effects, new evidence and further refusal signals return the decision to educational purpose.

From diagnostic refusal to accountable design

Refusal becomes useful when universities treat it diagnostically. The question is not how to eliminate refusal, but what refusal reveals about the conditions under which AI is being introduced. Where students refuse AI-generated course materials, they may be defending the expectation that university teaching carries academic presence. Where academics refuse AI feedback systems, they may be defending the relation between assessment and disciplinary judgment. Where professional staff resist automated triage, they may be identifying risks that are invisible in platform demonstrations but visible in student support work (Bayne, 2015; Down, 2025; Noddings, 2013). Diagnostic refusal is important because individual opt-out is often impossible. Many AI systems are embedded in infrastructure that students and staff do not directly choose: spam filters, search ranking, learning analytics, plagiarism tools, productivity suites, helpdesk systems and accessibility platforms. In these cases, refusal shifts from individual non-use to institutional transparency, procurement scrutiny, audit, appeal, human review and contestability (Star & Ruhleder, 1996; Williamson & Eynon, 2020). The diagnostic value of refusal also extends beyond immediate teaching concerns. AI systems carry environmental costs, labour dependencies and extractive data practices that may be hidden from users. Stakeholders who refuse AI for environmental, labour, intellectual property or cultural reasons are not necessarily avoiding innovation. They are making visible the conditions through which educational convenience is produced (Bender et al., 2021; Crawford, 2021; Gray & Suri, 2019). Accountable design begins when these signals are translated into institutional action. Refusal can lead to non-AI pathways, clearer assessment rules, secure tools, academic approval processes, workload recognition, human escalation, stronger procurement standards, transparent disclosure or a decision not to implement a system. Educational technology staff are often central to this translation because they can turn concern into workflow, governance and design. They do not decide alone, but they help make refusal practically actionable.

Six refusal zones in higher education

The six refusal zones are proposed heuristic domains in which educational thinning and legitimacy tensions can be examined. They are neither exhaustive nor mutually exclusive. The public examples below illustrate the kinds of concerns that make each domain analytically useful; they do not constitute evidence that the domains are prevalent or universal. The relationship is sequential but iterative: examples make a concern visible; the zones locate the educational responsibility at stake; and the design cycle converts that concern into questions, distributed judgment and a documented decision. Figure 1 summarises this logic.

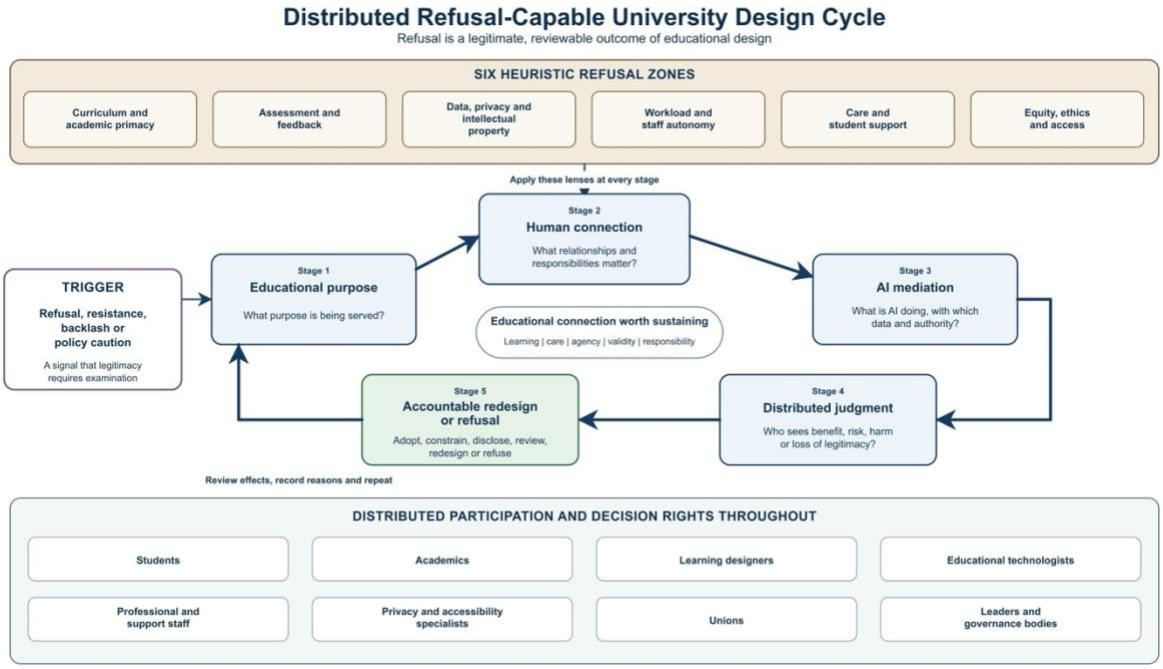


Figure 1. Distributed Refusal-Capable University Design Cycle. Selected public signals are interpreted through six heuristic refusal zones and processed through five recurring design questions. Stakeholders participate throughout, and implementation effects return the decision to review

Curriculum and academic primacy

AI-generated course materials are among the clearest examples of educational thinning. Slides, scripts, readings, quizzes, activities and weekly modules can be generated quickly. Their fluency can conceal genericity, error and weak disciplinary connection. Curriculum may appear complete while academic authorship has been displaced, a concern made visible in public backlash against AI-generated course delivery (Bayne, 2015; Down, 2025). Curriculum is more than a collection of topics, activities and assessments. It is an academic construction of disciplinary knowledge, intellectual progression, professional standards, epistemic judgment, ethical formation and public purpose. Decisions about what is taught, why it matters, how difficulty is sequenced, what evidence counts and how students are inducted into a field require academic responsibility (Biesta, 2015; Ito et al., 2013). Learning designers should not be positioned as substitutes for academics. Their work supports academic curriculum: strengthening alignment, improving accessibility, designing interaction, enhancing student experience and translating academic intent into coherent learning environments. Learning design extends academic judgment. It should not replace it. A refusal-capable university should set a clear boundary: academics retain responsibility for curriculum intent, disciplinary accuracy, intellectual sequencing, assessment standards and final approval of course content. AI-generated content should not be released to students unless an academic has verified its accuracy, relevance, disciplinary integrity, ethical appropriateness and alignment. Refusal protects academic legitimacy, student trust and learning designer professionalism (Bayne, 2015; University of Queensland, 2026).

Assessment and feedback

Assessment is a major refusal zone because AI destabilises the relationship between student work, student learning and academic judgment. If students can outsource written work, if staff can use AI to generate feedback, and if institutions use AI to detect AI use, assessment validity cannot be assumed (Furze et al., 2024; Perkins et al., 2024; Tertiary Education Quality and Standards Agency, 2026). AI-assisted assessment can be constructive. It can support formative practice, help students critique AI outputs, generate feedback prompts, support self-evaluation and encourage transparency through AI-use declarations. Frameworks such as the AI Assessment Scale help move the debate away from blanket prohibition by distinguishing levels of permitted AI involvement in assessment (Furze et al., 2024; Perkins et al., 2024). Refusal is necessary when AI weakens the

validity or fairness of assessment. AI-generated feedback should not be disguised as academic response. AI systems should not make final grading decisions without accountable human judgment. AI detection should not be used as sole evidence of misconduct. Detector scores may appear objective, but they can produce false accusations, damage trust and create inequitable suspicion (Tertiary Education Quality and Standards Agency, 2026; University of Melbourne, 2025). Assessment redesign should focus less on catching AI and more on evidencing learning. Possible responses include oral defence, in-class work, staged drafting, reflective process evidence, authentic performance, supervised components, peer critique and transparent AI-use declarations. AI may support assessment and feedback, but it must not replace academic responsibility for judging learning (Furze et al., 2024; Tertiary Education Quality and Standards Agency, 2026; University of Queensland, 2026).

Data, privacy and intellectual property

Data governance is one of the most immediate sites of AI refusal. Staff should refuse to upload student work, private information, unpublished research, confidential materials, Indigenous knowledge, colleague resources or institutional intellectual property into open tools without approval. Such refusal is not a barrier to innovation. It is an expression of institutional care (University of Melbourne, 2025). AI systems often operate through unclear data practices. User-submitted material may be retained, processed, used for product improvement, or governed by terms that staff and students do not understand. Universities cannot ask students to trust digital systems while treating their intellectual and personal materials as disposable prompt input (Bender et al., 2021; Crawford, 2021; University of Melbourne, 2025). A refusal-capable university should provide secure AI environments, clear data classification rules, prompt examples, procurement standards, staff training and embedded warnings. Refusal should be built into workflow design, not left as individual vigilance. Educational convenience does not justify unsafe data exposure (Miao & Holmes, 2023; University of Melbourne, 2025).

Workload and staff autonomy

AI adoption is often justified through efficiency, but efficiency claims can conceal labour transfer. Staff may be required to redesign assessment, check AI outputs, moderate generated feedback, answer student questions, revise policies, manage appeals, evaluate accessibility, disclose AI use and learn new systems. These tasks require time. If workload is not recognised, refusal is predictable and justified (National Tertiary Education Union, 2024; Selwyn, 2022). AI also raises questions of autonomy. Academics may resist AI if it appears to reduce their authority over curriculum, assessment or feedback. Professional staff may resist AI if it automates complex support work without recognising their expertise. Learning designers may resist AI workflows that turn them into content production technicians rather than educational partners. Unions may resist AI where it increases workload, threatens intellectual property or displaces staff expertise (National Tertiary Education Union, 2024; Selwyn, 2022). A refusal-capable implementation should include workload modelling, consultation, professional learning, role clarity and recognition of quality assurance labour. Staff should not be performance-managed for refusing unsafe or educationally unjustified AI use. Nor should AI be used to remove academics from curriculum or reduce professional staff to monitors of automated processes (National Tertiary Education Union, 2024; Selwyn, 2022).

Care and student support

Student support is another area where technical connection and educational connection may diverge. Chatbots and automated support systems can help students navigate deadlines, policies, forms, learning resources and routine questions. They may improve access outside business hours and reduce friction. These are real benefits, but they are not the same as care (Noddings, 2013; Sharples, 2023). Student support often involves vulnerability. Students may seek help because of distress, disability, financial hardship, exclusion, harassment, academic failure, illness, progression risk or personal crisis. These situations require interpretation, care, escalation and responsibility. A chatbot may provide information, but it cannot hold final responsibility for care. Care ethics emphasises attentiveness, responsiveness and relational responsibility (Noddings, 2013). Universities should not confuse automated availability with care. AI may support care by improving navigation and reducing delays, but it should not replace human support where judgment, advocacy or empathy is required. Practical refusal here means clear escalation pathways, visible human contact and human review for consequential support decisions.

Equity, ethics and access

AI adoption creates equity and ethical challenges that may produce legitimate refusal. AI can support accessibility through captioning, translation, simplification and alternative formats. It can also create new exclusions. Students without access to paid tools may be disadvantaged. Multilingual students may be disproportionately suspected by detection systems. Disabled and neurodivergent students may experience automated systems as inaccessible, punitive or inflexible. Where AI use is optional, equivalent non-AI pathways help ensure that refusal remains practically available rather than merely symbolic (Tertiary Education Quality and Standards Agency, 2026). Students may also object to compulsory AI use on ethical, environmental, cultural or professional grounds. The UNSW generative AI arts controversy illustrates the point. Student concern centred not only on a teaching method, but on intellectual property, environmental cost, artistic identity and whether students should be required to use technologies they considered ethically problematic (Bender et al., 2021; Cassidy, 2025; Crawford, 2021; Gray & Suri, 2019). Universities should be especially careful where AI is used with Indigenous knowledges, community-based curriculum or culturally sensitive materials. Human and community authority must remain central. AI cannot be treated as a neutral summariser of knowledge that is relational, situated or governed by cultural protocols. AI use should not be compulsory, exclusionary or extractive where ethical, cultural, environmental or accessibility concerns cannot be addressed (Crawford, 2021; Tertiary Education Quality and Standards Agency, 2026).

Institutional conditions for refusal capability

Refusal-capability requires transparency, contestability, role clarity, workload recognition and meaningful alternatives. Transparency means that students and staff should know when AI materially shapes course content, feedback, assessment, support, analytics or administrative decisions; hidden AI creates avoidable mistrust because stakeholders cannot judge whether a system is supporting education or quietly substituting for it (Jasanoff & Kim, 2015; University of Queensland, 2026). Contestability means that AI-mediated decisions affecting learning, progression, assessment, support, workload or reputation should be open to question through appeal pathways, human review and documentation of how AI was used (Star & Ruhleder, 1996; Tertiary Education Quality and Standards Agency, 2026). Role clarity means recognising that student refusal, academic refusal, professional staff refusal and institutional refusal carry different responsibilities and consequences (Baumer et al., 2015; Selwyn, 2003). Workload recognition means asking who will check generated materials, moderate feedback, explain rules, redesign assessment, audit tools and support appeals, rather than treating AI as automatic efficiency (National Tertiary Education Union, 2024; Selwyn, 2022). Meaningful alternatives mean that, where AI use is not essential to a learning outcome or institutional function, students and staff should have viable non-AI pathways; where AI is essential, its educational, ethical and technical justification should be explicit (Tertiary Education Quality and Standards Agency, 2026; University of Queensland, 2026).

Educational technology staff as mediators of accountable redesign

Because refusal is distributed, no single group can manage it alone. Educational technology staff, learning designers and technology administrators are nonetheless important to how refusal is handled. They are positioned between stakeholders and systems. They understand platforms, workflows, constraints, risks and implementation realities. They also see where institutional ambitions collide with educational practice (Selwyn, 2022; Williamson & Eynon, 2020). Their role is not to persuade everyone to accept AI. It is to help universities make AI decisions more educationally accountable. This includes making AI visible where it mediates feedback, content, support, assessment or decisions. It also involves building refusal pathways so that staff and students can raise concerns, pause implementation, request review and propose alternatives (Selwyn, 2022; Stilgoe et al., 2013). Educational technology staff can translate objections into design requirements. A complaint about AI-generated content can become a requirement for academic approval. A concern about privacy can become a secure workflow. A concern about workload can become pilot resourcing. A concern about automated support can become escalation design. Without this mediating work, refusal is more likely to become grievance, reputational damage or stalled implementation. With it, refusal becomes design intelligence. This mediation should be recognised as educational work. It requires technical knowledge, pedagogical judgment, ethical sensitivity and institutional courage. The educational technology professional of

the AI era is not simply an implementer. They are a steward of educational connection, working at the difficult boundary between institutional ambition and educational consequence (Bayne, 2015; Selwyn, 2022).

Counterarguments and limits

A theory of refusal must not romanticise refusal. Not all refusal is principled. Some may be reactionary, poorly informed, protectionist or exclusionary. Refusal can protect academic standards, but it can also preserve outdated practices. Refusal can defend student trust, but it can also deny students access to tools they need in professional life. Refusal can slow harmful adoption, but it can also produce institutional paralysis (Baumer et al., 2015; Selwyn, 2003). For this reason, refusal must be accountable. It should be linked to reasons, evidence, stakeholders and educational purpose. A refusal-capable university is not one where any stakeholder can block AI indefinitely. It is one where objections are heard, examined, documented and translated into design decisions. Some refusals will lead to rejection. Others will lead to safer adoption, conditional use, opt-out pathways, human review or further consultation (Miao & Holmes, 2023; Stilgoe et al., 2013). There is also a risk that institutions use refusal language to avoid responsibility. A university might refuse AI in assessment without investing in redesign. It might refuse staff use of AI without providing support. It might refuse student use while quietly adopting AI for administrative efficiency. Such positions are not refusal-capable design. They are avoidance. The challenge is to distinguish principled refusal from defensive obstruction (Stilgoe et al., 2013; Tertiary Education Quality and Standards Agency, 2026).

Discussion

The broader significance of refusal is that it reframes AI adoption as a question of university legitimacy. Universities are not just deciding which tools to use. They are deciding what kind of educational institutions they will become. Will curriculum remain academically authored? Will assessment remain valid? Will feedback remain accountable? Will support remain caring? Will student data remain protected? Will staff expertise remain central? Will AI adoption be participatory, or imposed through infrastructure? These questions place AI governance within public educational purpose, not merely technical implementation (Biesta, 2015; Jasanoff & Kim, 2015; Selwyn, 2022). Backlash should not be feared as an obstacle to progress. It should be read as evidence that the university's AI imaginary is being contested. Jasanoff and Kim's work on sociotechnical imaginaries reminds us that technological futures are not simply built; they are collectively imagined, authorised and resisted. Refusal is one way educational communities contest futures they do not trust. There is a positive path forward. AI can support more inclusive, responsive and creative university learning when it is governed through educational purpose. It can assist academics without replacing academic primacy. It can support learning designers without turning them into AI content technicians. It can help students practise, explore and critique without outsourcing formation. It can improve student support without automating care. It can help institutions see patterns without reducing students to data (Miao & Holmes, 2023; Sharples, 2023; Tertiary Education Quality and Standards Agency, 2026). The condition is refusal-capability. A university that cannot refuse AI cannot govern AI. It can only adapt to vendor systems, student practices and institutional pressure. A refusal-capable university retains agency. It can adopt AI where educationally justified, constrain it where risk is present, and reject it where it undermines learning, care, equity or public purpose (Stilgoe et al., 2013; Williamson & Eynon, 2020).

Conclusion

AI is changing the connective structure of higher education. It is reshaping how courses are produced, how students receive support, how assessment is secured, how feedback is generated, how staff work is organised, how data is used, and how universities imagine efficiency and responsiveness. These changes create opportunities, but they also create new forms of educational thinning (Selwyn, 2022; Sharples, 2023; Williamson & Eynon, 2020). Refusal provides a way to defend educational substance. It is not opposition to technology. It is a form of distributed university judgment through which students, academics, professional staff, learning designers, educational technologists, unions and leaders contest AI uses that weaken trust, academic primacy, assessment validity, equity, care or legitimacy (Baumer et al., 2015; Hirschman, 1970; Stilgoe et al., 2013). The future university will need AI. It will also need the ethical capacity to refuse AI. Without that capacity, connection becomes automation, curriculum becomes content, feedback becomes text,

support becomes routing, and education becomes a platform service. With that capacity, universities can build more accountable AI futures: technologically capable, educationally serious, publicly defensible and still human where humanity matters most.

References

Baumer, E. P. S., Ames, M. G., Burrell, J., Brubaker, J. R., & Dourish, P. (2015). Why study technology non-use? *First Monday, 20*(11). https://doi.org/10.5210/fm.v20i11.6310

Bayne, S. (2015). Teacherbot: Interventions in automated teaching. *Teaching in Higher Education, 20*(4), 455-467. https://doi.org/10.1080/13562517.2015.1020783

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610-623). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Biesta, G. (2013). *The beautiful risk of education*. Paradigm Publishers.

Biesta, G. (2015). What is education for? On good education, teacher judgment, and educational professionalism. *European Journal of Education, 50*(1), 75-87. https://doi.org/10.1111/ejed.12109

Cassidy, C. (2025, August 29). Professor defends new UNSW course on AI art against student backlash despite "layer cake of nightmares." *The Guardian*. https://www.theguardian.com/australia-news/2025/aug/29/university-nsw-generative-ai-art-course-students-push-to-abolish-rejected

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press. https://doi.org/10.12987/9780300252392

Down, A. (2025, November 20). "We could have asked ChatGPT": Students fight back over course taught by AI. *The Guardian*. https://www.theguardian.com/education/2025/nov/20/university-of-staffordshire-course-taught-in-large-part-by-ai-artificial-intelligence

Furze, L., Perkins, M., Roe, J., & MacVaugh, J. (2024). The AI Assessment Scale (AIAS) in action: A pilot implementation of GenAI-supported assessment. *Australasian Journal of Educational Technology, 40*(4), 38-55. https://doi.org/10.14742/ajet.9434

Gray, M. L., & Suri, S. (2019). *Ghost work: How to stop Silicon Valley from building a new global underclass*. Houghton Mifflin Harcourt.

Hirschman, A. O. (1970). *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states.* Harvard University Press.

Ito, M., Gutierrez, K., Livingstone, S., Penuel, B., Rhodes, J., Salen, K., Schor, J., Sefton-Green, J., & Watkins, S. C. (2013). *Connected learning: An agenda for research and design*. Digital Media and Learning Research Hub. https://researchonline.lse.ac.uk/id/eprint/48114/

Jasanoff, S., & Kim, S.-H. (Eds.). (2015). *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*. University of Chicago Press.

Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

National Tertiary Education Union. (2024). *Discussion paper on the impacts of artificial intelligence technologies in the tertiary education sector.* https://researchers.westernsydney.edu.au/en/publications/discussion-paper-on-the-impacts-of-artificial-intelligence-techno

Noddings, N. (2013). *Caring: A relational approach to ethics and moral education* (2nd ed.). University of California Press.

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale: A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice, 21*(6). https://doi.org/10.53761/q3azde36

Selwyn, N. (2003). Apart from technology: Understanding people's non-use of information and communication technologies in everyday life. *Technology in Society, 25*(1), 99-116. https://doi.org/10.1016/S0160-791X(02)00062-3

Selwyn, N. (2022). *Education and technology: Key issues and debates*. (3rd ed.). Bloomsbury Academic.

Sharples, M. (2023). Towards social generative AI for education: Theory, practices and ethics. *Learning: Research and Practice, 9*(2), 159-167. https://doi.org/10.1080/23735082.2023.2261131

Star, S. L., & Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research, 7*(1), 111-134. https://doi.org/10.1287/isre.7.1.111

Stilgoe, J., Owen, R., & Macnaghten, P. (2013). Developing a framework for responsible innovation. *Research Policy, 42*(9), 1568-1580. https://doi.org/10.1016/j.respol.2013.05.008

Tertiary Education Quality and Standards Agency. (2026). *Gen AI: Academic integrity and assessment reform*. https://www.teqsa.gov.au/guides-resources/higher-education-good-practice-hub/gen-ai-knowledge-hub/gen-ai-academic-integrity-and-assessment-reform

University of Melbourne. (2025). Using AI in teaching, learning and assessment. *Teaching and Learning Innovation*. https://www.unimelb.edu.au/tli/news/articles/Using-AI-in-teaching%2C-learning-and-assessment

University of Queensland. (2026). UQ's rules for using AI in assessment. *Institute for Teaching and Learning Innovation*. https://itali.uq.edu.au/teaching-guidance/generative-ai-teaching-learning-and-assessment/rules-using-ai-assessment

Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology, 45*(3), 223-235. https://doi.org/10.1080/17439884.2020.1798995

Zagami, J. (2026). When connection requires refusal: AI refusal and accountable technology design in higher education. In Cochrane, T., Narayan, V., Poole, T., Lai, P., Lake, P., Moore, O., & Li, L. (Eds.), *The Future is Connection: Collaboration, Creativity and Compassion in a Digital Society*. Proceedings ASCILITE 2026. Brisbane (pp. 419–429). https://doi.org/10.65106/apubs.2026.3132

Note: All published papers are refereed, having undergone a double-blind peer-review process.

The author(s) assign a Creative Commons by attribution license enabling others to distribute, remix, tweak, and build upon their work, even commercially, as long as credit is given to the author(s) for the original creation.

© Zagami, J. 2026