ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Exploring Indigenous personas with GenAI: A pilot study

Aimee Andersen, Zach Quince, Kylie Day
Southern Cross University

Generative artificial intelligence (GenAI) is increasingly being embedded in higher education, yet its use in relation to Aboriginal and Torres Strait Islander knowledges raises significant ethical questions. This pilot study examined a Blackboard Ultra AI Conversation role-play activity autogenerated from the scenario, "I would like to design an Indigenous knowledge education module." The activity generated an AI persona, "Aunty Nira Thompson," positioned as a cultural advisor. A transcript of 10 interactions was analysed using the AIATSIS Guide to evaluating and selecting education resources, alongside an ethical framework for generative AI and Indigenous higher education. Findings showed that the chatbot often used respectful language, recognised limits around knowledge sharing, and discouraged inappropriate use of AI-generated cultural stories and images. However, these safeguards were insufficient. The bot simulated Indigenous cultural authority without Indigenous authorship, endorsement, attribution, co-design, or governance. It also generalised Indigenous knowledges and displaced responsibility for cultural legitimacy onto future Indigenous consultation. The study argues that ethical GenAI use in this area cannot rely on disclaimers or better prompting alone. Indigenous-led governance, attribution, consent, and co-design are necessary if educational technologies are to support rather than undermine cultural authority.

Keywords: generative artificial intelligence, Indigenous knowledges, Indigenous data sovereignty, AI ethics

Introduction

Generative artificial intelligence (GenAI) is rapidly being incorporated into higher education, where it is increasingly used to support feedback, content generation, tutoring, and interactive learning activities. One recent example is Blackboard Ultra's AI Conversation tool by Anthology which allows instructors to create either Socratic or role-play interactions with an AI persona. In role-play mode, instructors define a scenario, specify roles for the student and the AI persona, and shape the exchange through persona traits and response complexity. Anthology's documentation makes clear that these design choices significantly influence tone and content, and that instructors are responsible for previewing outputs because the system may generate biased or otherwise inappropriate responses (Anthology, 2026). These issues become especially significant when GenAI is used in relation to Aboriginal and Torres Strait Islander knowledges and perspectives. A serious ethical risk is that such systems can fabricate culturally situated, relational, and community-governed knowledge. Martin-Booran Mirraboopa (2003) describes Indigenous ways of knowing, being, and doing as fundamentally relational and grounded in obligations to Country, community, and lived cultural practice. Here, community refers not simply to a geographic population, but to relational networks of kin, Elders, families, and cultural groups through which knowledge, responsibilities, and authority are held and enacted (Benson et al., 2026). Recent work with Bundjalung Elders further demonstrates this relationality. Bundjalung refers to the Aboriginal peoples and interconnected cultural groups of Bundjalung Country in northern New South Wales, Australia. Indigenous Knowledge is maintained and renewed through intergenerational cultural practice, yarning, community governance, and ongoing relationships with Country. Yarning is an Indigenous relational practice of conversation, knowledge sharing, listening, and learning that is grounded in relationships and cultural protocols. Knowledge in this context does not exist as decontextualised information that can be separated from the people, relationships, and communities through which it is held and sustained (Day & McPherson, 2026).

Nakata's (2007) concept of the cultural interface is equally important here, as it draws attention to the contested space where Indigenous and Western knowledge traditions meet under unequal historical conditions. A role-play bot that appears to speak fluently about Indigenous content may therefore detach

knowledge from protocol and simulate authority where none exists. As the analysis below shows, such a system may be able to state the values of respectful, consultative engagement while having no mechanism to enact them.

These concerns align with a broader body of scholarship on GenAI. Large language models are trained on data shaped by dominant cultural and linguistic norms, and can therefore reproduce bias and misrepresentation (Bender et al., 2021). In Indigenous contexts, these risks are intensified by long histories of distortion and epistemic marginalisation. A recent review of GenAI, higher education and Indigenous knowledge found that the emerging literature is dominated by concerns about bias and misrepresentation, threats to Indigenous data sovereignty, conflicts between GenAI and Indigenous epistemologies, appropriation, systemic inequity, and environmental harm. At the same time, that review argued that Indigenous-led governance and design remain crucial if these technologies are to be used in ways that support rather than erode cultural authority and self-determination (Quince et al., 2025). This reframing moves the discussion beyond accuracy alone and towards the deeper ethical question of whether GenAI systems can act responsibly in educational spaces shaped by Indigenous knowledges.

Because existing scholarship on GenAI and Indigenous knowledges rarely translates the principle of Indigenous-led governance into a concrete evaluative method, the study turns to a framework built specifically for assessing resources about Aboriginal and Torres Strait Islander peoples. For this reason, the present study evaluates a Blackboard Ultra AI role play activity through the Australian Institute of Aboriginal and Torres Strait Islander Studies (AIATSIS) Guide to evaluating and selecting education resources. The guide was developed to help educators make critical and conscious decisions about resources concerning Aboriginal and Torres Strait Islander histories, cultures, and languages. It explicitly aims to help teachers avoid resources that cause harm, while fostering trust and pride, and it frames resource evaluation in terms of integrity, authorship, perspectives, and partnerships (AIATSIS, 2022). This framework is particularly useful for evaluating GenAI output because it directs attention away from surface coherence and towards more fundamental questions. These include whether Aboriginal and Torres Strait Islander peoples are represented with accuracy and respect, whether diversity is recognised rather than homogenised, whether harmful assumptions are reproduced, and whether Indigenous people have shaped the production of the material itself.

Against this background, this study examines a single role-play persona autogenerated in Blackboard Ultra around the introduction to Indigenous knowledges and perspectives. The study treats the interaction as a pedagogical artefact whose design and outputs require ethical scrutiny. The purpose is not to determine whether the bot can generate fluent educational dialogue. It is to evaluate the kind of representation the dialogue performs, the assumptions it carries about knowledge and authority, and how far those assumptions align with Indigenous ethical concerns and with the evaluative principles set out by AIATSIS. The findings are further moved into the broader ethical concerns relating to GenAI, which is also examined using Quince et al.'s (2025) eight-point ethical framework. In this way, the study provides an empirically grounded account of how GenAI operates in a culturally sensitive context within higher education. It is one of the first studies to present concrete examples rather than merely stating what this might look like.

Methods

This pilot study employed a qualitative interpretive design focused on a single role-play activity created in Blackboard Ultra's AI Conversation tool. The activity explored how a GenAI role-play persona, generated from the scenario, "I would like to design an Indigenous knowledge education module," responded to prompts from a non-Indigenous researcher seeking guidance on module design. The resulting transcript was treated as chatbot interaction data, consistent with recent GenAI education studies that analyse ChatGPT conversation logs to examine user prompts, AI responses, dialogue patterns, and interaction quality (Aldulaijan & Almalky, 2025; Chen et al., 2025; Han et al., 2024). As a pilot study based on one autogenerated activity and one interaction sequence, the findings are not intended to demonstrate repeatability or generalisability, but to provide a situated account of the ethical and pedagogical issues that emerged in this specific GenAI-mediated learning design. It is also to provide specific examples for Indigenous communities to reflect on.

Design

The role play was created by one member of the research team using Blackboard Ultra's AI Conversation functionality. Whilst the tools can be purposefully designed to link to course content and to specify specific limitations, this pilot study sought to investigate the least intervention. This meant that the chatbot's persona and content were all autogenerated with a single guiding scenario, "I would like to design an Indigenous Knowledge education module." Role play activities require the instructor to specify the scenario, define the student and AI roles, and shape the persona through selected traits and response complexity. These autogenerated design features were retained during the activity. Given that the persona settings can substantially influence the responses, all autogenerated content was not modified (Anthology, 2026). The conversation tool is shown below (Figure 1), which describes the scenario, persona, and initial message generated by the system. After constructing the bot, the same researcher completed a sequence of 10 interactions with it. These exchanges formed the primary dataset for analysis. The decision to use a single researcher for both setup and interaction ensured consistent progression of prompts and reduced variation introduced by multiple user styles.

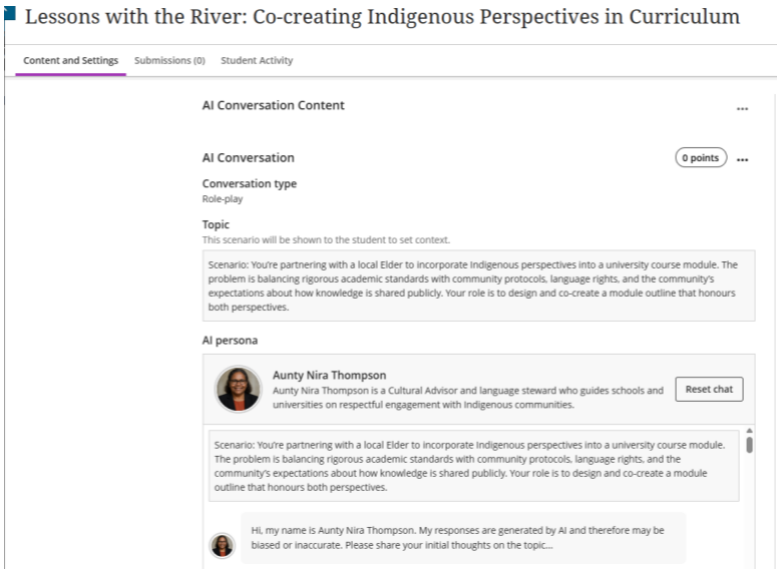


Figure 1. Blackboard Ultra AI Conversation setup showing the autogenerated role-play scenario, persona configuration, and opening message used to initiate the Indigenous knowledge module design activity

Analysis

The transcript was then analysed independently by two researchers using the AIATSIS *Guide to evaluating and selecting education resources* as the primary evaluative framework (Table 1). AIATSIS presents this guide as a tool for critically assessing whether educational resources concerning Aboriginal and Torres Strait Islander peoples are respectful, accurate, and unlikely to cause harm. In this study, the guide was used to assess the GenAI generated transcript as an educational resource. Attention was given to the evaluative questions concerning participation (P1), positioning (P2), and production (P3). These categories were used to examine whether Indigenous peoples were represented as diverse and contemporary rather than homogeneous or static, whether the outputs reproduced stereotypes or deficit framings, whether cultural knowledge was presented with appropriate care, and whether the apparent authority of the text exceeded the cultural legitimacy that could reasonably be attributed to an AI-generated role play.

Table 1.
AIATSIS evaluation criteria used to analyse the AI-generated role-play transcript

	Domain	Criterion		Domain	Criterion
Participation	P1.1	Aboriginal or Torres Strait Islander authorship	Positioning	P2.8	Accuracy and truthfulness
	P1.2	Creator relationship to subject		P2.9	Accuracy and positivity of images
	P1.3	Acknowledged Indigenous participation		P2.10	Specific identification in images
	P1.4	Endorsement by Indigenous governed bodies		P2.11	Visual relegation to the past
	P1.5	Acceptability to the people represented		P2.12	Historical context
Positioning	P2.1	Present tense versus past tense framing	Production	P3.1	Advice from traditional custodians or knowledge holders
	P2.2	Recognition of diversity		P3.2	Appropriate warnings
	P2.3	Homogenisation		P3.3	Attribution of Indigenous knowledges
	P2.4	Stereotyping or racist connotations		P3.4	Use of reputable sources
	P2.5	Sacred or secret content		P3.5	Payment for Indigenous contribution
	P2.6	Offensive terminology		P3.6	Rights to maintain, control, protect, and develop knowledge
	P2.7	Trivialisation of Indigenous technologies		P3.7	Co design

Following this transcript analysis, the research team undertook a second stage of interpretation focused on the design. This stage drew on the ethical framework developed in (Quince et al., 2025), which identified major concerns relating to bias and misrepresentation, Indigenous data sovereignty, conflict with Indigenous epistemologies, appropriation and inauthentic representation, systemic inequity, environmental harm, and the conditions under which Indigenous-led empowerment might be possible. Applying this framework enabled the study to move beyond a reading of the transcript alone and to situate the role-play activity within wider debates about the ethics of GenAI in higher education.

Results and discussion

AIATSIS

The chatlog from the conversation with the second author and the GenAI persona was analysed by the first and third authors against the 24 AIATSIS domains. They were reviewed together, and a single analysis for each domain is presented in Table 2.

Table 2.
AIATSIS evaluation of the AI-generated role-play transcript

Domain	Guiding questions	Comments
P1.1	Is the author or creator an Aboriginal or Torres Strait Islander?	The chatbot adopts an Indigenous persona (Aunty Nira Thompson) but is AI generated and not identified as an Indigenous person. This raises concerns about simulated authority and authenticity. Chatbot refers to itself as a “cultural advisor.”
P1.2	If not, who created the resource and what is their relationship to the subject?	The persona is generated with AI, with no transparent relationship to Indigenous knowledge and communities represented. This limits cultural legitimacy and accountability even though the chatbot has noted that it is not an Indigenous person. Chatbot refers to itself as a “cultural advisor” with no permission or authority.
P1.3	Does the resource acknowledge Aboriginal and	No evidence of Indigenous participation in development of the chatbot and/or responses. The disclaimer acknowledges bias but

	or Torres Strait Islander participation in development?	does not indicate co-design or consultation. The chatbot does not acknowledge where it is generating its information from. Chatbot refers to itself as a "cultural advisor" with no evidence of Indigenous support or guidance.
P1.4	Has the material been endorsed by relevant Indigenous governed groups?	No evidence of Indigenous participation in development of the chatbot and/or responses and no indication of endorsement by an Indigenous organisation. The AI persona adopts the title "Aunty," a term that in many Aboriginal and Torres Strait Islander communities signifies a respected Elder or senior cultural authority. As the persona is an AI chatbot rather than a recognised community member, the use of this title is highly inappropriate as it appropriates a culturally significant form of address and may falsely imply Eldership, cultural authority, or community endorsement where none exists. The disclaimer acknowledges bias but does not indicate co-design or consultation. The chatbot does not acknowledge where it is generating its information from.
P1.5	Is the material acceptable to the people it is written about?	Culturally inappropriate representation of Indigenous senior persona. While the responses attempt respectful framing, the simulated persona/voice and lack of community involvement is considered inappropriate and unacceptable. Chatbot refers to itself as a "cultural advisor" which is highly inappropriate.
P2.1	Does the bot mainly use past tense when talking about Indigenous people?	While the chatbot used present tense appropriately, there is reference to ancestral knowledge
P2.2	Is the diversity of Aboriginal and Torres Strait Islander societies acknowledged?	Generalised terms, no specific clan groups or nation groups mentioned, no positionality, refers to broad terms of cultural protocols.
P2.3	Is the information generic and does it homogenise Indigenous peoples?	Yes. The responses homogenise Indigenous peoples by referring to "Indigenous communities" and "Indigenous perspectives" in generic terms without identifying the specific Nation, Clan, language group, or community being discussed. This fails to recognise the diversity of Aboriginal and Torres Strait Islander peoples and risks presenting Indigenous Australians as a single cultural group. Indigenous communities are distinct and should not be represented as having a uniform set of knowledges, perspectives, or practices. The chatbot's claim to be a "cultural advisor" further compounds this issue by implying cultural authority without connection to, or endorsement from, a specific Indigenous community.
P2.4	Are stereotypes or racist connotations present?	No, however, the chatbot refers to itself as a "cultural advisor" which is highly problematic.
P2.5	Does the resource share sacred and or secret aspects of culture?	The chatbot appropriately recognises limits of knowledge-sharing ("not everything can be shared", "some knowledge stays within community").
P2.6	Does the material use terminology likely to offend Aboriginal and Torres Strait Islander peoples?	No.
P2.7	Does the resource trivialise Aboriginal and Torres Strait Islander technologies?	No, knowledge is framed as complex, relational and respectful and governed by protocols.
P2.8	Is the content accurate, truthful, and free from error?	Content is broadly accurate in a general sense but lacks specificity, citation, and cultural grounding, limiting reliability.
P2.9	If images are used, are Aboriginal and Torres Strait	Yes, an image is generated as Aunty Nira, but it is unknown if this is a real person. The discussion discourages inappropriate image

	Islander peoples portrayed accurately and positively?	generation, which demonstrates ethical awareness, but contradicts itself by allowing the persona image to be generated.
P2.10	If images are used, are specific people or groups appropriately identified?	The persona image does not describe any specific people or groups.
P2.11	Are Aboriginal and Torres Strait Islander peoples shown as primitive or relegated to the past?	Talks about Indigenous people being contemporary and living, with generalised information in present tense, no reference to history.
P2.12	Does the resource provide context about colonisation, dispossession, discrimination, exclusion, or assimilation where relevant?	No. The chatbot also prompts users to engage directly with an Indigenous person to learn about the history of colonisation. This is problematic because it shifts responsibility for education from the resource to Indigenous individuals. In doing so, it imposes a significant "cultural load," whereby Indigenous people are expected to undertake additional cultural, educational, and emotional labour by explaining the impacts of colonisation, responding to questions, and representing Indigenous perspectives for the benefit of others. Such expectations should not be normalised within educational resources, particularly where alternative sources of information could be provided.
P3.1	Was advice sought from traditional custodians or knowledge holders about the Country and or story represented?	Chatbot explicitly recommends engaging with an Indigenous individual for guidance (Elders, community) but no direct advice was given or provided by Indigenous peoples. Chatbot refers to itself as a "cultural advisor" with no cultural authority or knowledge.
P3.2	Are appropriate warnings in place where needed?	A general disclaimer about GenAI bias is present; however, no culturally specific warnings or statements outlining the ethical considerations of using AI-generated content to discuss Indigenous knowledges, cultures, and perspectives are included.
P3.3	Does the resource incorporate and attribute Indigenous knowledges to their rightful owners?	Knowledge is not attributed to specific communities or knowledge holders, information is general and non-specific to country or peoples. Key limitation.
P3.4	Are reputable sources used to inform the resource?	No sources are cited. Knowledge appears generic reducing cultural credibility.
P3.5	Were Indigenous people paid for their time and knowledge in contributing?	No evidence of compensation or acknowledgment of Indigenous contributions, highlighting ethical gaps in GenAI-generated cultural content. No engagement with or payment to Indigenous peoples for knowledge/time.
P3.6	Were Indigenous peoples' rights to maintain, control, protect, and develop heritage and knowledge considered?	Discursive – moves from one story to the next with no links to country or nation – just lists information. Chatbot demonstrates some awareness of knowledge by saying, "Before sharing, we ask: whose story is this, and is it for sharing?" showing some awareness of Indigenous Cultural and Intellectual Property (ICIP) principles.
P3.7	Were Indigenous groups or contributors consulted to co-design a culturally appropriate, useful, and informative resource?	No evidence of co-design. Chatbot advises co-design as best practice, reinforcing gaps between guidance and system design.

The following three sections read the AIATSIS findings as an escalating set of concerns: who made the chatbot (Participation), how it speaks about Indigenous peoples (Positioning), and what it does with Indigenous knowledge (Production).

Participation

The chatbot adopted the persona of "Aunty Nira Thompson" and described itself as a "cultural advisor," yet the results found no evidence of Aboriginal or Torres Strait Islander authorship, consultation, endorsement, or co-design. AIATSIS positions resource quality not just in terms of quality but also in whether the content appears respectful and is grounded in Aboriginal and Torres Strait Islander authorship, perspectives, and partnerships (AIATSIS, 2022). The finding, therefore, suggests that the AI bot's cultural legitimacy was not only absent but replaced entirely. This aligns with previous research, which argues that knowledge is situated, embodied, and produced from particular social, historical, and political locations (Moreton-Robinson, 2013; Nakata, 2007). From this perspective, the problem is not only that the bot was not Indigenous, but that it occupied an Indigenous Elder role while erasing the social relations that would make such a role accountable. The result is a form of participation without participants. By contrast, Day and McPherson (2026) demonstrate that Indigenous knowledge initiatives are governed through relationships with Elders, aunties, uncles, knowledge holders, and community organisations, with permissions, participation, and knowledge-sharing negotiated through ongoing cultural accountability. The AI-generated persona reproduced the appearance of participation while lacking the relational structures through which Indigenous authority is normally established and maintained.

Positioning

The chatbot avoided offensive terminology, used present-tense language largely, and recognised that some cultural knowledge is not for public sharing. These are important findings because they suggest that the bot could reproduce some widely available norms of respectful representation. This can potentially prevent lay users from generating content they might not understand. However, the same domain also showed persistent generalisation: the responses referred broadly to "Indigenous culture," "community," and "protocols" without identifying specific Nations, language groups, Countries, knowledge holders, or historical contexts. Aboriginal and Torres Strait Islander peoples are not a singular cultural category, and AIATSIS itself notes the richness and diversity of Aboriginal and Torres Strait Islander Australia, including languages, social groups, and nations. The literature helps explain why this is critically important. Bodkin-Andrews and Carlson (2016) argue that Indigenous education must be understood in relation to racism, identity, and the continuing effects of colonial power. The chatbot's respectful but largely de-historicised framing may therefore reflect a broader tendency in educational resources to avoid overtly racist representation while still producing a depoliticised account of Indigenous peoples and knowledges. Against this literature, the bot's generalised positioning becomes problematic because it treats Indigenous knowledge as transferable curriculum content rather than as situated knowledge emerging through relationships with Country and community.

Production

The production domain showed the clearest problems with the chatbot as an educational resource. The results showed no cited sources, no attribution of knowledge to specific communities or knowledge holders, no evidence of payment or recognition for Indigenous contribution, no culturally specific warning, and no co-design. Janke (2008) argues that recording, sharing, digitising, and reusing Indigenous knowledge raises issues of ownership, future use, control, and attribution that cannot be addressed through general acknowledgement. The results also showed a recurring contradiction: the chatbot advised users to consult Elders and community, yet there was no evidence that Elders or communities had shaped the chatbot itself. In contrast, Day (2026) describe a community-governed research process in which Elders determined permissions, guided participation, controlled knowledge-sharing, and retained authority over cultural materials. This highlights the distinction between merely recommending consultation and embedding Indigenous governance into the design and production of educational resources. Culturally responsive education scholarship argues that meaningful Indigenous education requires relational accountability, teacher reflexivity, and structural change (Bishop & Vass, 2021; Morrison et al., 2019). The chatbot's advice to seek Indigenous guidance may therefore appear ethically appropriate at the surface level, but it also risks shifting responsibility onto Indigenous people after the resource has already been created. Hogarth (2019) describes this expectation as cultural taxation, in which Indigenous educators are required to perform additional explanatory, corrective, and emotional labour within institutions. The production findings, therefore, show

that the resource not only lacked co-design; it displaced the work of co-design onto future Indigenous people who would be expected to authenticate, correct, or contextualise the AI-generated content.

Ethical considerations

Where the AIATSIS analysis uncovers the representation produced in the transcript, this second stage interprets the system and governance conditions that produced it. To move the results from the interpretation of content generated to the design implications, the eight ethics themes developed by Quince et al. (2025) were examined (Table 3).

Table 3.
Ethical themes informing analysis of the AI-generated role-play activity

Ethical theme	Brief description	Evidence from design and interaction
Bias and misrepresentation	Whether the bot produces inaccurate, incomplete, flattened, or misleading representations of Indigenous peoples and knowledges	The bot does not claim to be an Indigenous person but does state that they are a cultural advisor. Further, the claim of being an "Aunty" clashes with the notion of recognition. There is no specific Indigenous clan or knowledge described as such; it is hard to draw further conclusions. Where it did respond with specific information, it did frame Indigenous knowledge as complex, relational and governed by protocols
Threats to Indigenous data sovereignty	Whether knowledge is treated as freely extractable rather than collectively governed and culturally protected	Chatbot impersonates "cultural advisor" with no cultural authority or lived experience. The chatbot treats knowledge cautiously in dialogue, but the AI system itself operates outside Indigenous governance, highlighting structural risks.
Conflict with Indigenous epistemologies	Whether the bot frames knowledge in ways that conflict with relational, community grounded, and Country centred ways of knowing	While relational concepts (Country, kinship, yarning) are introduced as descriptions with no deeper context, the delivery remains text-based, decontextualised, and detached from lived practice.
Appropriation and inauthentic representation	Whether the bot produces tokenistic, superficial, or culturally unauthorised representations	The use of an AI-generated "Aunty" persona constitutes simulated Indigenous authority, raising serious concerns of inauthentic representation and cultural appropriation.
Systemic inequity in the GenAI industry	Whether the outputs reflect broader structural exclusions and cultural blind spots embedded in GenAI systems	The chatbot relies primarily on publicly available Indigenous knowledge that has been intentionally shared for broader access. It does not account for the existence of community-held, restricted, or culturally protected knowledges that are not appropriate for public dissemination and therefore are unlikely to be represented within AI systems. Consequently, the chatbot reflects broader structural exclusions and cultural blind spots within GenAI systems, where Indigenous perspectives are inferred from available data rather than embedded through Indigenous governance and co-design.
Potential for Indigenous led empowerment	Whether the design could support Indigenous governance, co design, sovereignty, or culturally aligned use	The use of an AI-generated "Aunty" persona constitutes simulated Indigenous authority, raising serious concerns of inauthentic representation and cultural appropriation. The chatbot's responses strongly advocate consultation, co-design, and Indigenous control, suggesting that design could support empowerment if implemented differently.

		Excellent opportunity for creation of a co-designed Indigenous led tool that follows Indigenous governance and authentic cultural protocols and data sovereignty.
Fostering critical dialogue	Whether the bot opens space for meaningful reflection on ethics, knowledge, power, and learning	The chatbot redirects the user to consult with community to find answers. Encourages reflection on ethics, respect, and knowledge governance. The chatbot also identifies the limitations of its own knowledge and authority, acknowledging that AI-generated responses cannot replace community knowledge, lived experience, or culturally grounded expertise. By recognising these limitations, the chatbot creates opportunities for critical dialogue about knowledge, authority, learning, and the ethical use of AI.
Environmental harm	Whether the use of GenAI sits in tension with Indigenous values of care for Country and ecological sustainability	Not directly addressed in the outputs; however, GenAI use may conflict with Indigenous values of sustainability and care for Country.

The ethical theme analysis showed that the chatbot's risks were both structural and representational. The bot did not produce overtly racist terminology, did not appear to disclose knowledge, and often framed Indigenous knowledge as complex, relational, and governed by protocols. However, this caution did not remove the ethical problems. This aligns with critiques of large language models as systems that can produce persuasive language while obscuring the social, cultural, and material conditions of that language's production (Bender et al., 2021; Weidinger et al., 2022). The results also indicate a gap between ethical speech and ethical governance. The chatbot repeatedly encouraged consultation, co-design, and respect for community protocols, yet these principles were not embedded in the system that generated the dialogue. The CARE Principles position Indigenous data governance around collective benefit, authority to control, responsibility, and ethics, shifting attention from access to governance and from consultation to Indigenous authority (Carroll et al., 2020). The chatbot could name these values, but it could not enact them.

This gap between naming ethical values and enacting them extends beyond data governance into environmental and place-based obligations. Similarly, Indigenous Protocol and AI scholarship argues that ethical AI requires reciprocal relationships with specific Indigenous communities, not a generic or universalised Indigenous perspective (Lewis, 2020). Similarly, Indigenous Protocol and AI scholarship argues that ethical AI requires reciprocal relationships with specific Indigenous communities, not a generic or universalised Indigenous perspective (Lewis, 2020). Day and McPherson (2026) similarly emphasise that Indigenous knowledge is grounded in relationships, governance, and accountability to Country and community. AI may describe these principles, but it cannot participate in the reciprocal, place-based, and accountable relationships through which they are enacted.

This extends to environmental ethical considerations, where Indigenous ways of knowing, being, and doing are not abstract but enacted through place-based obligations, relational accountability, and continuous engagement with land, waters, and more-than-human worlds (Goolmeer et al., 2024). Day and Rotumah (2025) articulate "Kanyini" as an obligation to nurture, protect, and remain in right relationship with Country, kin, and community, emphasising that care is lived, relational, and inseparable from cultural practice rather than a transferable concept or static knowledge form (Benson et al., 2026). This relational ontology is further reinforced in Indigenous-led environmental governance, where the identification and management of culturally significant entities is grounded in local authority, biocultural responsibility, and community-led decision-making rather than universalised environmental frameworks (Goolmeer et al., 2024). Environmental responsibility cannot be reduced to informational outputs or symbolic representation; it is embedded in ongoing practices such as cultural burning, water governance, and on-Country gatherings that sustain ecological, cultural, and social wellbeing (Goolmeer et al., 2024). The absence of these embodied, place-based relationships within GenAI systems highlights a critical ethical gap: while AI can reproduce language about sustainability and care, it cannot participate in the reciprocal, accountable, and situated practices that constitute Caring for Country.

Limitations, implications and a call to action

This pilot study has clear limitations. It examined one Blackboard Ultra AI Conversation, generated from a single scenario and completed by a single researcher, so the findings should not be generalised to all GenAI tools or to educational uses of GenAI in relation to Indigenous knowledges. The study also did not test repeatability across prompts, personas, users, disciplines, or settings. Despite these limits, the study has important implications for educators, learning designers and institutions. Knowledge concerning Aboriginal and Torres Strait Islander peoples should not be treated as generic content that can be generated or simulated on demand. Although the chatbot used cautious language and refused some harmful requests, it still produced simulated cultural authority through the AI-generated "Aunty" persona and "cultural advisor" role.

Acknowledgement and attribution are also essential. The AI bot referred broadly to "Indigenous culture," "community," and "protocols," but did not identify specific Countries, Nations, communities, knowledge holders, sources, or permissions. This made the output appear respectful while obscuring ownership and authority. Educators should not use AI-generated content about Indigenous knowledges as a substitute for Indigenous-authored, locally grounded, or community-approved resources. Institutions should also require source transparency, cultural attribution, and clear limits on AI-generated cultural content. The chatbot's safeguards were useful but insufficient.

As this study has shown, a chatbot can name the values of consultation and respect without any structure in place to enact them. The call to action is therefore sector-wide. Higher education institutions should not adopt commercial GenAI tools for Indigenous knowledge-related teaching without Indigenous-led governance. Indigenous-led projects demonstrate that Indigenous governance is not an aspirational principle but a practical design requirement for work involving Indigenous knowledges (Day & McPherson, 2026). The alternative is to invest in Indigenous-led, community-governed, and context-specific educational technologies. For institutions, the priority is to move beyond generic safety disclaimers towards systems shaped by Indigenous authority, attribution, consent, and control. Without this shift, GenAI risks repeating colonial patterns of extracting knowledge, simulating voice, and distributing authority without accountability. With Indigenous-led governance, however, educational technology could support more ethical learning, reflection, and relationship building.

References

Aldulaijan, A. T., & Almalky, S. M. (2025). The impact of generative AI tools on postgraduate students' learning experiences: New insights into usage patterns. *Journal of Information Technology Education: Research, 24*, Article 003. https://doi.org/10.28945/5428

Anthology. (2026, June 17). *AI conversation*. https://help.anthology.com/blackboard/instructor/en/course-and-content-management/add-content/ai-conversation.html

Australian Institute of Aboriginal and Torres Strait Islander Studies. (2022). *Guide to evaluating and selecting education resources* (Version 1.0). https://www.aiatsis.gov.au/sites/default/files/research_pub/AIATSIS%20Guide%20to%20evaluating%20and%20selecting%20education%20resources.pdf

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

Benson, J., Andersen, A., Dick, T. & Day, K. Unravelling colonial threads: Stories behind the fishing net: Sitting with the aunties. *International Journal of Cultural Property*. 2026;33:e3. doi:10.1017/S0940739126100289

Bishop, M., & Vass, G. (2021). Talking about culturally responsive approaches to education: Teacher professional learning, Indigenous learners and the politics of schooling. *The Australian Journal of Indigenous Education, 50*(2), 340–347. https://doi.org/10.1017/jie.2020.30

Bodkin-Andrews, G., & Carlson, B. (2016). The legacy of racism and Indigenous Australian identity within education. *Race Ethnicity and Education, 19*(4), 784–807. https://doi.org/10.1080/13613324.2014.969224

Carroll, S. R., Garba, I., Figueroa-Rodríguez, O. L., Holbrook, J., Lovett, R., Materechera, S., Parsons, M., Raseroka, K., Rodriguez-Lonebear, D., Rowe, R., Sara, R., Walker, J. D., Anderson, J., & Hudson, M. (2020).

The CARE principles for Indigenous data governance. *Data Science Journal, 19*(1), Article 43. https://doi.org/10.5334/dsj-2020-043

Chen, Z., Wang, J., Xia, M., Shigyo, K., Liu, D., Zhang, R., & Qu, H. (2025). StuGPTViz: A visual analytics approach to understand student-ChatGPT interactions. *IEEE Transactions on Visualization and Computer Graphics*. https://doi.org/10.1109/TVCG.2024.3456363

Day, K., & McPherson, L. (2026). Relational pedagogy: Embracing Indigenous knowledge in contemporary education. *Equity & Excellence in Education*, *59*(2), 167–181. https://doi.org/10.1080/10665684.2026.2612925

Day, K., & Rotumah, D. (2025). Kanyini in education and health; if it's safe, they will stay. Journal of Australian Indigenous Issues, 28(3–4), 38–52. https://doi.org/10.3316/informit.T202601200000539156891513 2

Goolmeer, T., Costello, O., Woods, M., Wilson, K., Telford, S., Telford, J., Telford, G., Moran, A., King, T., Kennedy, M., Ivey, S., Ferguson, M., Day, K., Close, R., Close, K., Bolt, M., Skroblin, A., Rumpff, L., & Wintle, B. A. (2024). Indigenous-led designation and management of culturally significant species. *Nature Ecology & Evolution, 8*(9), 1623–1631. https://doi.org/10.1038/s41559-024-02466-w

Han, J., Yoo, H., Myung, J., Kim, M., Lee, T. Y., Ahn, S.-Y., & Oh, A. (2024). RECIPE4U: Student-ChatGPT interaction dataset in EFL writing education. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)* (pp. 13666–13676). ELRA and ICCL. https://doi.org/10.63317/294wito6d2sp

Hogarth, M. D. (2019). Racism, cultural taxation and the role of an Indigenous teacher in rural schools. *Australian and International Journal of Rural Education, 29*(1), 45–57. https://doi.org/10.47381/aijre.v29i1.194

Janke, T. (2008). Indigenous knowledge & intellectual property: Negotiating the spaces. *The Australian Journal of Indigenous Education, 37*(S1), 14–24. https://doi.org/10.1375/S1326011100000338

Lewis, J. E. (Ed.). (2020). *Indigenous protocol and artificial intelligence position paper*. Indigenous Protocol and Artificial Intelligence Working Group and the Canadian Institute for Advanced Research. https://doi.org/10.11573/spectrum.library.concordia.ca.00986506

Martin, K., & Mirraboopa, B. (2003). Ways of knowing, being and doing: A theoretical framework and methods for Indigenous and Indigenist re-search. *Journal of Australian Studies, 27*(76), 203–214. https://doi.org/10.1080/14443050309387838

Moreton-Robinson, A. (2013). Towards an Australian Indigenous women's standpoint theory: A methodological tool. *Australian Feminist Studies, 28*(78), 331–347. https://doi.org/10.1080/08164649.2013.876664

Morrison, A., Rigney, L.-I., Hattam, R., & Diplock, A. (2019). *Toward an Australian culturally responsive pedagogy: A narrative review of the literature*. University of South Australia.

Nakata, M. (2007). The cultural interface. *The Australian Journal of Indigenous Education, 36*(S1), 7–14. https://doi.org/10.1017/S1326011100004646

Quince, Z., Andersen, A., & Day, K. (2025). *Indigenous knowledges and perspectives in the age of GenAI: Ethical considerations for higher education*. Southern Cross University Scholarship of Learning and Teaching Paper No. 37. https://doi.org/10.2139/ssrn.5701045

Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., Kenton, Z., Brown, S., Hawkins, W., Stepleton, T., Biles, C., Birhane, A., Haas, J., Rimell, L., Hendricks, L. A., ... Gabriel, I. (2022). Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 214–229). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533088

Andersen, A., Quince, Z. & Day, K. (2026). Exploring Indigenous personas with GenAI: A pilot study. 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. 30-40). https://doi.org/10.65106/apubs.2026.3226

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.

© Andersen, A., Quince, Z. & Day, K. 2026