ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Cogniti agents in practice: AI as academic infrastructure in a foundation business course

Nuddy Pillay, Una Lightfoot
University of Auckland

Generative artificial intelligence (GenAI) is increasingly entering universities not only as a student-owned productivity tool but as part of the digital infrastructure through which learning support, feedback and assessment guidance are mediated. This concise paper reports on an exploratory pilot of Cogniti AI agents in an Aotearoa New Zealand foundation business course, with a focus on human-centred learning design. The pilot involved five role-specialised agents: a course buddy for administration support, a case-discussion agent for Māori entrepreneurship learning, and three virtual team members focused on task coordination, relationship management and value-adding contribution. Rather than positioning AI as a generic chatbot, the paper presents a design model for embedding AI agents within the course ecosystem. Drawing on design records, early usage patterns and practice reflections, the paper argues that AI agents can extend human-centred pedagogy when they are co-created with students, shaped by student input, assessment-aligned, and circumscribed by teacher oversight. The contribution is a practical model for integrating AI agents into learning management systems in ways that prioritise connection, mentorship, evaluative judgement, feedback literacy and collaborative learning.

Keywords: Cogniti, AI agents, academic infrastructure, teamwork, human-centred learning design, feedback literacy, mixed methods

Introduction

The rapid uptake of generative artificial intelligence (GenAI) in higher education has moved institutional attention beyond questions of whether students will use AI towards questions of how educators can design AI use for learning. Early debate has often focused on academic integrity, individual productivity and the risks of students outsourcing cognitive work (Bittle & El-Gayar, 2025; Kasneci et al., 2023). A complementary question is now emerging: what changes when AI is designed as part of the academic environment itself? This paper takes up that question through the concept of Cogniti AI agents as academic infrastructure.

By academic infrastructure, we refer not simply to a digital tool but to a situated arrangement of technology, curriculum, assessment, policy, data practices, and teacher oversight that shapes how students seek help, interpret expectations, and collaborate. This framing draws on infrastructure studies, where infrastructure is understood as relational, embedded in practice and visible through use, rather than as technology alone (Star & Ruhleder, 1996). It also aligns with recent higher education GenAI literature, which argues that meaningful AI integration requires attention to governance, data privacy, teaching design and equity (Jin et al., 2025; McGrath et al., 2025).

Cogniti provides teacher-designed generative AI agents that can be embedded into Canvas and steered by educator instructions and course resources (Cogniti, 2024, 2026). This creates opportunities for educators to move from generic AI advice to bounded, course-specific support. This concise paper reports on an exploratory pilot in a foundation business course with approximately 140 students. The course serves pre-degree learners and includes agents for course navigation, case-based learning and team-based assessment. The paper is framed as research-informed innovative practice. It presents the design model; explains why it matters and identifies implications for educators considering AI agents as part of course infrastructure.

This paper is guided by two evaluative questions: (a) how can role-specialised Cogniti agents be designed as course-based academic infrastructure in a foundation business course? and (b) what early design insights emerge from student co-creation, usage patterns and teacher practice reflections?

From chatbot to course infrastructure

AI in higher education has a long history, but reviews have highlighted the need for stronger educator involvement in designing pedagogically meaningful AI applications (Zawacki-Richter et al., 2019). The arrival of large language models has intensified opportunities and challenges, including personalised explanation, feedback support, bias, hallucination, privacy and changes to assessment design (Kasneci et al., 2023). Recent work in tertiary education argues that GenAI requires educators and institutions to rethink assessment, learning design and governance, rather than treating AI as a peripheral technology issue (Lodge et al., 2023; Weng et al., 2024).

Course-embedded agents are one response to these concerns. Instead of positioning students as individual users of open-ended AI systems, educator-designed agents can be framed around particular learning activities, disciplinary language and assessment purposes. They can also be anchored to course documents and teacher-defined behaviours. Cogniti documentation notes that agents may be embedded in Canvas so students access them where learning occurs, and that uploaded resources can contextualise responses (Cogniti, 2024, 2026). At the same time, resource curation, prompt design and monitoring become part of the infrastructure because generative systems require careful management of the material from which they respond.

The pilot also responds to recurring challenges in collaborative assessment. Group work can support learning, but students often experience unequal contribution, communication problems and conflict (Chang & Brickman, 2018). AI agents may offer scalable, just-in-time prompts for team processes and formative guidance, but their role should be to help students interpret expectations, judge quality and decide when to seek human or peer support rather than replace peer accountability or teacher intervention. This aligns with the conference sub-theme of human-centred learning design because the central design challenge is not how to automate learning support but how to strengthen connection, reflection and relational support in increasingly digital learning environments.

Course context and design model

The pilot was implemented in a Tertiary Foundation Certificate business course delivered through Canvas. There were 140 enrolled students in the course. The course included three connected uses of Cogniti: administrative course support, case-method teaching and team-based assessment support. Rather than designing a single general-purpose chatbot, five role-specialised agents were created to reflect distinct student needs.

The design model had five elements: each agent had a clear pedagogical role; each agent was located within Canvas where students already accessed course materials; the agents were shaped by course-specific resources and teacher instructions; the team-support agents were co-created with students in response to their identified needs; and the agents were framed as supports for learning, reflection and decision-making, not as substitutes for teacher expertise or peer responsibility. The team support agents were co-created to support students who were working in teams of five to six to create a digital artefacts on a course topic of their choosing. Table 1 summarises the agent architecture.

Table 1

Role-specialised Cogniti agents in the course pilot

Agent	Student need	Design focus	Observation/refinement
Course buddy	Course navigation and administration	Warm course-coordinator persona; answers from course resources only	Useful just-in-time support; needs one curated resource pack
Case-discussion agent	Māori entrepreneurship case learning	Reflective questioning and conversation history around a case	High engagement; requires careful cultural positioning
Terri	Task progress	Planning, goals, workload and assessment milestones	Supports coordination and time management
Ray	Team relationships	Communication, conflict and relational judgement in teams	The agent with the highest number of student interactions

Val	Contribution quality	Peer accountability, evidence of contribution and value-add thinking	Supports fairness and standards in team assessment

Design insights

Because this is an exploratory pilot, the insights presented here are framed as practice reflections rather than claims of measured learning impact. The design insights were informed by four sources: agent design records created during development, aggregate Cogniti usage patterns, teaching-team practice reflections, and student input gathered through the co-creation process and anonymous course feedback. Design records included agent prompts, role descriptions, resource decisions and notes from the student co-creation activity used to inform the three team-support agents. Usage patterns were reviewed in aggregate to identify which agents attracted the most student interaction; individual student-level data were not analysed. Teaching-team reflections were based on observations of how the agents functioned in Canvas, how students engaged with them during assessment preparation, and what refinements were needed for the next iteration.

Ethics approval was granted by the University of Auckland Human Participants Ethics Committee. Participation in the research component was voluntary and separate from students' access to the Cogniti agents as part of normal course learning. Student usage data were reviewed only in aggregate form to inform course design and were not used to assess individual students. Anonymous survey responses were analysed without names, student IDs or identifying details. Participation or non-participation had no bearing on students' grades, access to learning support or relationship with the teaching team.

Five design insights are informing the next iteration.

Course embedded agent

Placing support inside Canvas changed the visibility and accessibility of help. The course buddy functioned as a Canvas-based support agent for routine questions about course processes. This made help available at the point of need. However, the pilot also highlighted the importance of data integrity. If course information is spread across multiple files or Canvas locations, the agent may draw on outdated or inconsistent material. The next iteration will use a single curated document containing all approved course information as the primary resource base, supported by clearer version control.

Role specialization

Role specialisation made the agents easier for students to understand. Students did not have to approach one general chatbot for every issue. Instead, they could select the agent whose role matched their immediate need: administrative guidance, case reflection, task planning, relationship support or value-adding contribution. This made the agents more like an academic support ecosystem than a single AI tool.

Co-creation

The co-creation of the team-support agents strengthened the human-centred design of the pilot. Before the agents were developed, students were asked during a class workshop what kind of AI support would help them work more effectively in teams. Their responses highlighted three areas: organising and completing the task, managing relationships and conflict, and adding value to the quality of their work. These student-identified needs became the basis for Terri, Ray and Val (see Table 1). The most used virtual team member was Ray, the relationship-focused coach. This suggests that students may see value in AI-mediated support for interpersonal dimensions of collaboration, particularly when they are unsure how to raise concerns with peers.

Cultural learning

The case-discussion agent demonstrated both the potential and limits of AI-supported cultural learning. Students appeared to engage strongly with the case activity, but the design also raised questions about cultural care, authorship, representation and authority. A human-centred approach requires that AI support for culturally located learning be

carefully positioned. The agent should support reflective engagement with a case; it should not be framed as a cultural authority or a replacement for Māori expertise, consultation or teacher judgement.

AI literacy and evaluative judgement

The pilot reinforced the importance of combining AI literacy with evaluative judgement. The following observations were made after a review of anonymised student Cogniti chats. Some students needed guidance on how to write effective prompts, evaluate AI responses and recognise when to seek human help. Students need to understand not only how to ask better questions, but also how to judge whether a response is useful, accurate, ethical and appropriate for the task. This aligns with research arguing that GenAI makes evaluative judgement more important because students must learn to assess AI outputs and the processes through which those outputs are produced (Bearman et al., 2024).

Discussion: Why the design model matters

The pilot suggests that the value of Cogniti agents lies less in novelty than in their infrastructural placement. Students did not need to leave Canvas, subscribe to a separate tool or translate generic AI advice back into local course requirements. The agents were embedded in the course ecology and aligned with learning and assessment moments. This matters because infrastructure shapes routine behaviour: where students go for clarification, how quickly they receive formative guidance and whether support is available outside scheduled teaching hours.

The design model also shifts the discussion from AI as individual productivity support to AI as shared academic support. In this pilot, the agents were connected to course documents, assessment expectations, teacher-defined purposes and student-identified needs. This does not remove the risks associated with GenAI. Accuracy, privacy, over-reliance, academic integrity and cultural safety remain important concerns (Kayalı et al., 2023). However, treating AI agents as academic infrastructure makes those concerns more visible and governable. The question becomes not simply whether students use AI, but how educators design, guide, monitor and explain AI-supported learning.

Three principles emerge. First, agents should be role specific. A course buddy, a case-discussion partner and a team relationships coach do different educational work. Second, agents should be grounded in curated resources and explicit instructions. Resource design is not an administrative afterthought; it is part of pedagogical quality assurance. Third, agents should be governed as socio-technical systems. Privacy, consent, data use, cultural safety, academic integrity and teacher oversight determine whether AI support is trustworthy.

The pilot also complicates narrow accounts of AI as a cheating risk or efficiency tool. The relationship-focused agent's higher use suggests that students may turn to AI for help with the hidden curriculum of university learning: interpreting expectations, navigating interpersonal uncertainty and developing confidence to contribute. Used carefully, AI agents can scaffold student learning by helping students make sense of expectations and decide how to act on guidance. In doing so, these tools can support student feedback literacy as conceptualised by Carless and Boud (2018), while also reflecting more recent accounts of feedback literacy in which learners must prompt, filter and evaluate automated or GenAI feedback sources (Carless, 2026; Zhan & Yan, 2025). However, this requires teachers to design agents that prompt reflection, comparison and judgement rather than provide ready-made answers.

Implications and next steps

For educators, the key implication is that AI agents should be designed from the learning problem, not from the tool. The most useful design move was not creating one chatbot for the course. It was identifying specific moments where students needed support and then designing agents with clear roles around those moments, including guidance on how to interpret, evaluate and act on AI-generated support. For learning designers and institutions, course-embedded AI requires new forms of quality assurance: prompt design, uploaded resources, version control, cultural review, privacy settings and student guidance all become part of the learning design process.

The paper reports an exploratory single-course pilot, so the insights should not be generalised as evidence of broad learning impact. Early usage patterns and practice reflections indicate promising areas for further investigation, but they do not show whether agent use improved assessment outcomes or long-term learning. The next design cycle will include a curated resource package, clearer prompt guidance, stronger cultural review of the case agent, escalation pathways for team conflict and fuller analysis of anonymous survey responses and aggregate usage patterns.

Conclusion

Cogniti agents offer a practical way to shift GenAI from an external, individualised tool into a course-based support infrastructure. In this pilot, the agents were most promising when they were aligned with authentic student needs, embedded in Canvas, designed for specific roles and governed through ethical safeguards. The design model shows that AI agents can support human-centred learning when they are co-created with students, circumscribed by teacher oversight and connected to assessment, feedback and collaboration. As academic infrastructure, AI agents should extend access to formative guidance while keeping teachers, peers and cultural responsibilities at the centre of learning design.

References

Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment & Evaluation in Higher Education, 49*(6), 893–905. https://doi.org/10.1080/02602938.2024.2335321

Bittle, K., & El-Gayar, O. (2025). Generative AI and academic integrity in higher education: A systematic review and research agenda. *Information, 16*(4), Article 296. https://doi.org/10.3390/info16040296

Carless, D. (2026). Feedback literacy. *International Journal of TESOL Studies, 8*(1), 1–16. https://doi.org/10.58304/ijts.260101

Carless, D., & Boud, D. (2018). The development of student feedback literacy: Enabling uptake of feedback. *Assessment & Evaluation in Higher Education, 43*(8), 1315–1325. https://doi.org/10.1080/02602938.2018.1463354

Chang, Y., & Brickman, P. (2018). When group work does not work: Insights from students. *CBE—Life Sciences Education, 17*(3), Article ar52. https://doi.org/10.1187/cbe.17-09-0199

Cogniti. (2024, September 12). *How do I embed a Cogniti agent into Canvas?* https://cogniti.ai/docs/how-do-i-embed-a-cogniti-agent-into-canvas/

Cogniti. (2026, February 25). *How are resources used by agents?* https://cogniti.ai/docs/how-are-resources-used-by-agents/

Huang, D., Huang, Y., & Cummings, J. J. (2024). Exploring the integration and utilisation of generative AI in formative e-assessments: A case study in higher education. *Australasian Journal of Educational Technology, 40*(4), 1–19. https://doi.org/10.14742/ajet.9467

Jha, M., & Atif, A. (2025). Reimagining pedagogy for the GenAI era: Frameworks, challenges and institutional strategies. *Australasian Journal of Educational Technology, 41*(5), 56–73. https://doi.org/10.14742/ajet.10645

Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence, 8*, Article 100348. https://doi.org/10.1016/j.caeai.2024.100348

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences, 103*, Article 102274. https://doi.org/10.1016/j.lindif.2023.102274

Kayalı, B., Yavuz, M., Balat, Ş., & Çalışan, M. (2023). Investigation of student experiences with ChatGPT-supported online learning applications in higher education. *Australasian Journal of Educational Technology, 39*(5), 20–39. https://doi.org/10.14742/ajet.8915

Lodge, J. M., Thompson, K., & Corrin, L. (2023). Mapping out a research agenda for generative artificial intelligence in tertiary education. *Australasian Journal of Educational Technology, 39*(1), 1–8. https://doi.org/10.14742/ajet.8695

McGrath, C., Farazouli, A., & Cerratto-Pargman, T. (2025). Generative AI chatbots in higher education: A review of an emerging research area. *Higher Education, 89*, 1533–1549. https://doi.org/10.1007/s10734-024-01288-w

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

Weng, X., Xia, Q., Gu, M., Rajaram, K., & Chiu, T. K. F. (2024). Assessment and learning outcomes for generative AI in higher education: A scoping review on current research status and trends. *Australasian Journal of Educational Technology, 40*(6), 37–55. https://doi.org/10.14742/ajet.9540

Zhan, Y., & Yan, Z. (2025). Students' engagement with ChatGPT feedback: Implications for student feedback literacy in the context of generative artificial intelligence. *Assessment & Evaluation in Higher Education*, 1–14. https://doi.org/10.1080/02602938.2025.2471821

ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education, 16*, Article 39. https://doi.org/10.1186/s41239-019-0171-0

Pillay, N., & Lightfoot, U. (2026). Cogniti agents in practice: AI as academic infrastructure in a foundation business course. 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. 932–937). https://doi.org/10.65106/apubs.2026.3336

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.

© Pillay, N., & Lightfoot, U. 2026