ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

AI as interlocutor, not author: A three-phase assessment design for reflective professional learning in business education

Bruce Mitchell, Mary Jesselyn Co
Monash University

This paper reports preliminary findings from a redesigned capstone assessment in a first-year Management unit at an Australian university involving over 300 students. Rather than treating generative AI as a content-generation tool, the assessment positions it as a simulated interview partner within a three-phase architecture: a pre-interview strategy memo, an AI-conducted adaptive interview, and a post-interview reflective analysis. Grounded in Schön's reflective practice framework and Biggs and Tang's constructive alignment, the design surfaces the gap between how students plan to communicate and how they do. Preliminary thematic analysis of 21 student reflective submissions reveals four patterns: (a) the strategy memo functions as a metacognitive anchor, making the gap between intention and execution visible and diagnosable; (b) students consistently demonstrate framework knowledge but revert to habitual communication patterns under adaptive pressure; (c) the structured reflection phase produces sophisticated self-diagnostic reasoning that connects transcript evidence to communication theory; and (d) students who completed multiple simulation attempts used iterative practice as a developmental mechanism rather than a remedial exercise. These findings suggest that AI-simulated interviews, when embedded within a reflective assessment architecture, generate qualitatively richer professional learning than traditional scripted interview preparation.

Keywords: generative AI, professional communication, reflective practice, interview simulation, AI-augmented assessment, constructive alignment, scholarship of teaching and learning

Introduction

Generative AI has created a qualitatively new problem for higher education assessment, distinct from earlier challenges around plagiarism or contract cheating. Tools that produce polished professional outputs risk decoupling what students submit from what they truly understand. Research on cognitive offloading suggests this risk is not hypothetical; increasing reliance on AI for routine intellectual tasks has measurable effects on retention and critical thinking (Yan et al., 2024). In professional communication education, this tension runs deeper still. Cover letters, resumes, reflective statements, and interview scripts—the very artefacts students are asked to produce—are now precisely what large language models generate well. The pedagogical question has shifted: it is no longer whether students can produce professional documents, but whether our assessment designs can make the *thinking behind* those documents visible and assessable. Yan et al's (2024) Theory of Agentivism frames this challenge directly: when AI systems contribute to task completion, learning can no longer be inferred from performance alone. Assessment must instead make visible the allocation of cognitive and epistemic agency between learner and system.

This paper describes a redesigned capstone assessment in a first-year Management unit (approximately 300 students per semester) at an Australian university. The redesign repositions AI from content-generation assistant to simulated interlocutor within a three-phase assessment architecture grounded in Schön's (1983) reflective practice framework and Biggs and Tang's (2011) constructive alignment. Students articulate a communication strategy, perform under adaptive pressure in an AI-simulated job interview, and then produce a critical reflective analysis comparing what they planned with what they did.

The redesign grew out of a structural limitation in the previous assessment, which asked students to generate interview questions using a commercial AI tool and write scripted answers supported by academic references. This appeared to integrate AI meaningfully, but it rewarded rehearsed performance rather than adaptive communication. Students selected their own questions and composed answers without time pressure,

conditions that bore little resemblance to the professional contexts the assessment claimed to prepare them for. The redesign draws on emerging work in AI-integrated assessment: Perkins et al.'s (2024) AI Assessment Scale; Xia et al.'s (2024) scoping review on assessment transformation for self-regulated learning; and Sarikose et al.'s (2025) finding that AI-guided interview simulations reduced anxiety among students. This paper extends that work by reporting preliminary findings from the first implementation cycle, drawing on thematic analysis of 21 student reflective submissions.

Assessment design

The assessment retains its job search portfolio structure—job selection, cover letter, resume, and LinkedIn profile—but transforms the interview component into three interconnected phases, each targeting a different dimension of professional communication competence. All three phases use text-based interaction. Students were instructed not to use AI tools to produce their strategy memo or reflective analysis; the assessment architecture itself guards against this since phases one and three require analysis of the student's own transcript and pre-interview reasoning, making AI-generated content incompatible with the specificity of the task demands. Submissions were also checked through the university's similarity detection software.

Phase 1: Pre-interview strategy memo (500 words). Before engaging with the AI chatbot, students write a strategy memo that identifies the key competencies the role demands, maps evidence from their resume to each competency, and sets out which communication frameworks (e.g., STAR, AIDA, ACE) they intend to deploy. Phase 1 represents an extension of Schön's (1983) work into proactive reflection.

Phase 2: AI-simulated interview (transcript submitted as appendix). Students interact with a purpose-built AI chatbot configured as a hiring manager who conducts a structured five-question interview with follow-up probes. The chatbot reviews the student's uploaded job advertisement and resume, adopts an appropriate persona, and works through a prescribed sequence: motivational, behavioural, situational, skills-based, and a curveball designed to test adaptive thinking. The chatbot was built using a structured system prompt deployed through a commercially available large language model (OpenAI's GPT-4o, accessed via the ChatGPT interface). The prompt instructed the chatbot to adopt a professional but approachable tone, ask probing follow-up questions, and avoid providing model answers or coaching during the interview. Each simulation generates a unique interview trajectory because the chatbot adapts its follow-up probes to the student's responses and the specific uploaded job advertisement. Importantly, the transcript is assessed for completion only, not performance quality—a deliberate choice to encourage authentic practice over rehearsed performance. Students can run the simulation multiple times, and no time constraint was imposed on reflective submission.

Phase 3: Post-interview reflective analysis (1,000 words). Students review their transcript and write a critical analysis comparing their strategy memo with their actual performance. They identify moments where they applied or failed to apply communication frameworks, evaluate their strengths and weaknesses, and propose a revised approach for one answer. This phase represents Schön's (1983) reflection-on-action and carries the greatest evaluative weight.

Methodology and data analysis

This study is part of an ongoing project examining the impact of the three-phase assessment redesign. Ethics approval for this project was granted by the Monash University Human Research Ethics Committee (Project ID: 46893). The data reported here are drawn from Phase 3 reflective submissions completed by students in one semester. From the cohort of approximately 300 students, a purposive sample of 21 de-identified reflections was selected to capture variation across performance levels, disciplinary backgrounds, and culturally and linguistically diverse student profiles.

The reflections were analysed using reflexive thematic analysis following Braun and Clarke's (2006) six-phase process: familiarisation through repeated close reading of all 21 submissions; systematic generation of initial codes across the dataset; collation of codes into candidate themes; review and refinement of themes against the coded extracts and the full dataset; definition and naming of final themes; and production of the analytic narrative presented below. Coding was conducted manually by both authors without the use of qualitative

data analysis software, a decision consistent with Braun and Clarke's emphasis on researcher-driven reflexive engagement with the data rather than software-mediated code management. Coding was conducted inductively—that is, themes were driven by the data rather than mapped onto a predetermined framework—though the analysis was inevitably informed by the theoretical framing of the assessment design (Biggs & Tang, 2011; Schön, 1983). The four themes reported represent patterns that were both prevalent across the dataset and analytically significant in relation to the study's research question: whether the three-phase design generates the reflective professional learning it was intended to produce.

Preliminary findings

Thematic analysis of 21 de-identified student reflective submissions (Phase 3) reveals four patterns relating to the execution of strategy, the application of frameworks, self-reflective diagnosis, and developmental iteration. While drawn from a purposive sample, they offer initial evidence that the three-phase design is generating the metacognitive learning it was intended to produce.

The strategy-execution gap as a learning mechanism

All the students identified substantive gaps between their strategy memo and their actual interview performance. What matters is that these gaps were not experienced as assessment failures—they became diagnostic evidence. As one student observed: "I knew what kind of response I was aiming for, yet I did not always sustain that standard once the pressure for specificity increased." Students described planning to use STAR but reverting to generalisations under the pressure of unanticipated follow-up probes or preparing evidence-rich competency maps that collapsed into narrative-driven value alignment when confronted with motivational questions. Several drew precise distinctions about what had gone wrong. One noted: "STAR was available to me, but the conditions of the interview activated a more instinctive communication pattern before I could apply it"—framing the issue as a performance failure, not a planning one. Another reflected: "this was a structural gap more than a knowledge one. The content was there; I just did not address the answer properly." A third recognised that deviating from planned structure under a scenario-based question actually "provided opportunity to produce a more natural answer," showing emerging judgement about when frameworks help and when they get in the way. This distinction—between knowing a framework and deploying it under adaptive pressure—is precisely the gap Schön (1983) described between espoused theory and theory-in-use.

Framework knowledge versus framework deployment

Across all submissions, a consistent pattern emerged: students knew their frameworks but could not consistently use them when it counted. They demonstrated strong conceptual knowledge of STAR, ACE, AIDA, and impression management, yet their transcripts told a different story. As one student put it: "I was more consistent in performing the style of fit than in delivering the evidence of fit." Another captured the tension succinctly: "I prioritised sounding natural over being evidence-based, and specificity was the casualty." The most common structural pattern was over-contextualisation—students spent disproportionate time on the "Situation" element of STAR while under-developing "Action" and "Result." Several noted that situational questions (requiring forward reasoning) produced more structured responses than behavioural questions (requiring retrieval of experience)—a counterintuitive finding that aligns with Culbertson et al.'s (2017) distinction between these question types.

Students also noticed how the chatbot's interaction dynamics shaped their communication in unexpected ways. One recognised that affirming feedback led them to mirror the chatbot's tone, producing "inadvertent informality" that was not a deliberate, adaptive audience-centred strategy. This observation is consistent with research on sycophancy in large language models, which has shown that LLM-generated affirming responses can produce confidence bias in users and shift their communicative behaviour in ways they do not intend (Sharma et al., 2024). In this assessment context, the affirming tone was not explicitly programmed into the chatbot's system prompt; it appears to be a default behavioural tendency of the underlying model, which makes it both an unintended variable in the assessment and a pedagogically productive one—students who noticed it were able to analyse its effect on their own performance.

The most striking interaction-level finding concerned honesty under non-human judgement. Another found that the absence of human judgement made them "more willing to be honest," which paradoxically led to responses flagged as potentially arrogant. This finding warrants further interrogation. Several students elaborated on the dynamic: freed from the social risk of face-to-face evaluation, they disclosed weaknesses more readily and made bolder claims about their strengths. The paradox is that the very condition that encouraged honesty—the absence of a human assessor—also removed the social calibration mechanisms that moderate self-presentation in live interviews. For assessment design, this suggests that AI-simulated interviews create a distinctive communicative space that is neither a rehearsal of nor a substitute for human interaction, but a context with its own social-psychological dynamics that students can learn to recognise and analyse. These observations suggest that the AI simulation surfaces communicative patterns through a mechanism different from traditional interview practice because the social dynamics of interacting with a non-human interlocutor produce behaviours that students can identify only in retrospect.

Structured reflection as sophisticated self-diagnosis

The reflective analyses showed a level of self-diagnostic reasoning that went well beyond what one would typically expect from first-year students. They did not just describe what went wrong—they theorised *why* it went wrong. One student analysed a defensive response to a stakeholder scenario, recognising that they had "mistakenly defaulted to defensiveness by imposing personal bias that technology-based learning is highly effective, instead of acknowledging varying perspectives." Another identified a tension between interpersonal warmth and evidential rigour, writing that "the same style which helped them come across as natural and approachable may also have made some answers feel smoother than they were precise." Perhaps the sharpest observation came from a student who noted: "under stress, polished phrasing feels safer because it is harder to challenge in the moment, but for a trained assessor that safety is precisely what makes it unconvincing." A student interviewing for a consulting role distinguished between instinctive and deliberate framework adaptation, recognising that their successful improvisation "was not a planned call" and that "effective professional communication requires not just knowing frameworks but knowing when to set them aside." Several students from culturally and linguistically diverse backgrounds produced particularly nuanced reflections on navigating communication norms across cultural contexts. These students were able to articulate specific tensions between the directness expected in Australian professional settings and the communication norms of their home cultures, suggesting the design supports cross-cultural metacognitive development. However, since these observations arise from a text-based interaction, it is important to note that the assessment does not yet capture the additional complexities that would accompany verbal communication, including accent, intonation, and real-life comprehension.

Iterative practice as developmental learning

Several students completed the simulation more than once, and their reflections show that repeated practice worked as a developmental mechanism—not a remedial one. One student wrote: "this was my third attempt at the simulation. In my first two runs, I kept noticing the same problem"—opening behavioural questions too broadly—and entered their final attempt with targeted awareness of that habit. Another documented how their first attempt revealed a failure to apply STAR effectively, which they corrected the second time around: "in my second attempt, I executed both frameworks more intentionally." A third tracked specific improvements across three attempts, noting that "my main improvement was that I connected the resolution to my actions," including adding measurable KPI evidence and more deliberate audience-centred language. This reframing—repeated practice as evidence of learning rather than failure—is pedagogically significant, and the AI simulation creates conditions where iteration feels developmental. This finding connects to the broader literature on the "hidden curriculum" in higher education (Margolis, 2001), which states how unwritten expectations around professional conduct, self-presentation, and interview technique operate as barriers for students without access to professional networks. By providing a low-stakes, repeatable practice space, the AI simulation makes at least one dimension of this hidden curriculum visible and practicable.

Discussion and implications

These preliminary findings suggest that the three-phase assessment architecture produces learning that is qualitatively different from what scripted interview preparation delivers. The strategy memo gives students a metacognitive anchor; the AI interview applies adaptive pressure that surfaces habitual communication patterns; and the structured reflection converts that experience into theoretically informed self-diagnosis. The iterative practice enables a developmental cycle that traditional one-shot assessments simply cannot replicate.

Three implications stand out. First, generative AI can function as a pedagogical mechanism rather than just a productivity tool. The chatbot does not teach students to communicate better—it creates conditions under which they discover how they communicate. This aligns with Wollny et al.'s (2021) observation that chatbots can serve as accessible, non-judgemental practice partners, and is consistent with the growing body of work on AI as a pedagogical interlocutor (for example, the resources from the Cogniti research group's 2024 and 2025 symposia on AI-augmented learning). Second, there are early signs of equity benefits. The chatbot gives every student—regardless of background, or prior interview experience—a consistent, low-stakes, repeatable practice opportunity, potentially levelling an uneven playing field that previous assessment iterations inadvertently reinforced. This equity dimension has its corollary in the concept of the hidden curriculum, where interview technique can function as a gatekeeping mechanism that disadvantages students without prior exposure to professional norms. Third, the approach requires no custom software development: the chatbot was configured using a structured system prompt deployed through a commercially available large language model, making the design transferable to any professional communication unit.

Limitations should be noted. The sample of 21 reflections is purposive, not representative; however, preliminary analysis of the broader de-identified dataset (n ≈ 300) is consistent with these patterns. The chatbot is text-based and cannot simulate non-verbal communication; a future iteration of the design could extend to verbal interaction, though this would introduce additional variables including transcription accuracy for students whose primary language is not English. Variation in chatbot behaviour across platforms has been observed, and the sycophantic tendencies of the underlying model represent an uncontrolled variable that warrants systematic investigation in future implementations. Full mixed-methods analysis—including systematic coding of strategy-transcript alignment and thematic analysis across grade bands—is currently underway. Future work will also explore a peer-feedback component following the AI simulation, to examine whether and how students' communicative performance shifts when human judgement is reintroduced.

References

Biggs, J., & Tang, C. (2011). *Teaching for quality learning at university* (4th ed.). Open University Press.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Culbertson, S. S., Weyhrauch, W. S., & Huffcutt, A. I. (2017). A tale of two formats: Direct comparison of matching situational and behavior description interview questions. *Human Resource Management Review*, *27*(1), 167–177. https://doi.org/10.1016/j.hrmr.2016.09.009

Margolis, E. (Ed.). (2001). *The hidden curriculum in higher education*. Routledge. https://doi.org/10.4324/9780203901854

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): 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

Sarikose, S., Sengul, T., Uncu, B., Kirkland-Kyhn, H., & Kaya, N. (2025). AI-guided interview simulation to improve employability and reduce anxiety in final-year nursing and midwifery students: A quasi-experimental study. *Clinical Simulation in Nursing*, *109*, 101852. https://doi.org/10.1016/j.ecns.2025.101852

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rauber, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2024). Towards understanding sycophancy in language models. In *Proceedings of*

the 12th International Conference on Learning Representations. https://doi.org/10.48550/arXiv.2310.13548

Wollny, S., Schneider, J., Di Mitri, D., Weidlich, J., Rittberger, M., & Drachsler, H. (2021). Are we there yet? A systematic literature review on chatbots in education. *Frontiers in Artificial Intelligence*, *4*, 654924. https://doi.org/10.3389/frai.2021.654924

Xia, Q., Weng, X., Ouyang, F., Lin, T. J., & Chiu, T. K. F. (2024). A scoping review on how generative artificial intelligence transforms assessment in higher education. *International Journal of Educational Technology in Higher Education*, *21*(40), 1–22. https://doi.org/10.1186/s41239-024-00468-z

Yan, L., & Gašević, D. (2026). Agentivism: A learning theory for the age of artificial intelligence. *arXiv preprint*,. https://doi.org/10.48550/arXiv.2604.07813

Yan, L., Greiff, S., Teuber, Z., & Gašević, D. (2024). Promises and challenges of generative artificial intelligence for human learning. *Nature Human Behaviour*, *8*, 1839–1850. https://doi.org/10.1038/s41562-024-02004-5

Mitchell, B., & Co, M. J. (2026). AI as interlocutor, not author: A three-phase assessment design for reflective professional learning in business 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. 741–746). https://doi.org/10.65106/apubs.2026.3418

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.

© Mitchell, B., & Co, M. J. 2026