

Streamlining Pre-Placement Compliance with GPT-Based Tools: A Case Study of the BLAKE Framework in Work-Integrated Learning

MICHAEL ROLDAN¹

University of Melbourne, Victoria, Australia

Abstract

As Work-Integrated Learning (WIL) programs expand, universities face increasing complexity in managing pre-placement compliance. This case study describes the development and implementation of a generative pre-trained transformer (GPT)-based internal assistant within the Faculty of Medicine, Dentistry and Health Sciences (MDHS) at the University of Melbourne, to support a small team in drafting policy-aligned compliance communication. A five-component governance framework, BLAKE, was designed to ensure responsible and human-in-the-loop use. Drawing on verified institutional and government sources, the tool generates draft responses that reflect the team's tone and values that is then reviewed by staff before dissemination. Early indications suggest improvements in turnaround times, staff confidence, and communication consistency, were achieved without compromising student privacy or regulatory requirements. This case highlights the potential for carefully governed generative AI tools to support administrative functions in WIL. It may offer a practical and transferable approach for institutions operating in similarly resource-constrained environments.

Keywords: Generative AI, Work-Integrated Learning, Student Pre-Placement Compliance, Higher Education Administration, Clinical Professional Education, BLAKE Framework.

Work-Integrated Learning (WIL) is central to the preparation of health professionals. It provides students with the clinical experiences necessary to develop practice-ready competencies. At the University of Melbourne, the Faculty of Medicine, Dentistry and Health Sciences (MDHS) coordinates approximately 3,500 student placements annually across six schools and more than twenty courses, spanning medicine, dentistry, allied health, psychology, and population health. The breadth of this program reflects the scale and diversity of compliance responsibilities managed by a small faculty-level administrative team. This creates a clear need for scalable solutions that can sustain staff capacity, reduce operational strain, and ensure the reliability and consistency of compliance processes.

¹ Corresponding author: michael.rolدان@unimelb.edu.au

The complexity of pre-placement compliance in clinical professional education is illustrative of the broader administrative burden faced by WIL programs across many disciplines, though the specific requirements vary significantly by field. Each placement is governed by rigorous pre-placement compliance requirements designed to ensure patient safety and meet regulatory standards. Students are required to satisfy up to 13 mandatory compliance checks before being cleared for placement. These span eight immunisation requirements including Hepatitis B, Influenza, Measles Mumps and Rubella, Pertussis, Tuberculosis, and Varicella and five non-immunisation checks, including a National Police Check, Victorian Working with Children Check, National Disability Insurance Scheme (NDIS) Worker Screening, Overseas Police Record, and Mask Fit Testing. The complexity of these requirements is considerable: some checks, such as a full Hepatitis B vaccination course, can take up to eight months to complete, and students must provide specific documentation aligned with the Australian Immunisation Handbook and relevant regulatory standards. Errors or delays in documentation can result in a student being unable to commence their placement, with direct consequences for their academic progress.

Managing this compliance load across approximately 3,500 students each year places significant pressure on administrative teams. At MDHS, a team of two to three full-time staff members is responsible for pre-placement compliance across the entire Faculty. As enquiry volumes rise due to growing student numbers and there is increasing complexity of regulatory requirements, staff face the challenge of maintaining timely, accurate, and consistent communication with students while simultaneously managing processing workflows, stakeholder relationships, and compliance documentation. Without proportional increases in staffing or digital infrastructure, this creates an unsustainable workload and risks to compliance integrity.

This case study responds to that challenge by documenting the development, implementation, and early outcomes of a GPT-based compliance assistant, and by proposing a governance framework, BLAKE, that may offer a practical model for WIL departments seeking low-resource, ethical approaches to AI integration. While developed in a clinical professional education context, the framework is designed to be adaptable by any small administrative team managing high volumes of policy-answerable enquiries.

Background and Literature Review

Literature on AI in higher education largely focusses on student-facing tools or teaching applications, such as adaptive learning platforms, intelligent tutoring systems, and automated feedback (Crompton & Burke, 2023). However, administrative use cases, particularly in the compliance and governance space, are emerging as viable pathways for meaningful AI integration (Luckin et al., 2022; Ocen et al., 2025). Luckin et al. (2022) highlight that AI tools hold significant potential to reduce administrative burden and support educator capacity - provided that governance structures are in place. Ocen et al. (2025) identify the automation of administrative tasks as one of the most consistently reported benefits of AI in higher education yet note that empirical case studies documenting such implementations in compliance-specific

contexts remain scarce. A key gap in the literature is the absence of staff-led frameworks that guide responsible AI integration in WIL administration.

Compliance management in clinical professional education represents a particularly pressing administrative challenge. WIL programs are time and resource intensive due to their administrative requirements, with institutional systems focussed on student management, monitoring, and compliance (Fleming & Hay, 2021). Risk areas perceived as high priority by WIL staff include health and safety, breach of confidentiality, and compliance with legislation and policy, all of which are directly implicated in pre-placement compliance processes in clinical professional education (Fleming & Hay, 2021). Australian institutions have been urged to adopt structured, well-informed approaches to mitigate both legal and operational risks in WIL programs (Cameron, 2017), yet the administrative infrastructure required to support this remains under-resourced at many universities (AHPRA Accreditation Committee, 2023). Existing literature also does not document how AI tools might be deployed to address this specific resource gap in health WIL compliance.

The integration of generative AI into administrative workflows introduces both opportunity and responsibility. Central to ethical deployment is the concept of human-in-the-loop (HITL) design, in which AI systems augment rather than replace human judgement, with staff retaining accountabilities for all outputs (Mosqueira-Rey et al., 2022). This approach has been shown to support trust, transparency, and institutional confidence in AI-assisted decision-making (Salloch & Eriksen, 2024). In compliance-sensitive environments such as clinical professional education, the HITL model is particularly appropriate to ensure that policy knowledge and professional judgement remain central to all student-facing communications. However, neither Mosqueira-Rey et al. (2022) nor Salloch and Eriksen (2024) address how HITL principles might be operationalised within a small, resource-constrained compliance team—a practical gap this case study addresses.

This case study describes the development and implementation of a GPT-based internal assistant designed to address these pressures by supporting staff in drafting compliance-related communications. It examines the tool's design, governance, and early outcomes, and reflects on the implications for WIL departments seeking low-resource, ethical approaches to AI integration in administrative practice.

Method

Study Design

This paper presents a case study of an administrative service improvement initiative designed within the MDHS Pre-Placement Compliance Team at the University of Melbourne. A case study approach was adopted to enable an in-depth and context-specific examination of how a GPT-based assistant could be integrated into compliance workflows in a real-world setting (Yin, 2018). The assistant was piloted from November 2024, with adoption expanding progressively through into 2026 as compliance staff were trained on its appropriate and

consistent use.

Tool Design and Technical Setup

The assistant was configured using the ChatGPT platform and designed exclusively as a staff-facing tool, with no student-facing interface or access. It operates as a prompt-based query system that generates draft email responses for staff to review, edit, and send. The assistant does not integrate directly with institutional student management systems and holds no access to student records.

Content Curation and Source Governance

A curated set of verified institutional and government sources was selected to guide the assistant's responses, including:

1. MDHS Pre-Placement Compliance Handbook
2. Australian Immunisation Handbook
3. Australian Technical Advisory Group on Immunisation (ATAGI) guidance
4. Standardised Student Induction Protocols (SSIP)
5. Internal Standard Operating Procedures (SOPs)

These sources were selected on the basis of their authority, currency, and direct relevance to pre-placement compliance requirements. Ongoing quality assurance requires periodic review of source documents to ensure alignment with updated policy and regulatory guidance. While the sources used here are specific to clinical professional education, the same configuration process applies to any team with verified internal documentation whether that be a faculty handbook, a university policy, a course guide, or any other authoritative institutional resource.

Governance and Privacy Approvals

Prior to deployment, the tool was reviewed by the University's legal team to confirm that the de-identified approach—whereby no student information is entered into the system at any point—satisfied the University's privacy and confidentiality obligations. Staff were advised that all enquiries must be submitted using generic, de-identified language. No formal human ethics approval was required given the absence of student data collection or processing. In this implementation, the process began with a single conversation with a line manager, followed by a request to the University's legal team both of which took less than one month in total. For teams considering a similar approach, this experience suggests that starting the conversation is often the hardest step, but the process itself moves quickly once initiated.

Operational Workflow

Enquiries are submitted by staff and then the AI assistant generates draft responses aligned with established compliance protocols. Typical enquiries include questions about immunisation documentation requirements, serology thresholds, police check processes, and overseas document translation standards. Prior to distribution, these responses are reviewed

and edited by staff, maintaining a human-in-the-loop model (Mosqueira-Rey et al., 2022; Salloch & Eriksen, 2024). No AI-generated response is distributed without human review. Staff apply their professional judgement and compliance expertise when reviewing and refining drafts, consistent with their role and responsibilities. While no formal audit checklist was used, full accountability for every response sent rests with the reviewing staff member, consistent with existing professional responsibilities within the compliance function. Enquiries are managed through a shared team inbox; however, each response is sent with the individual staff member's signature, providing a clear and identifiable record of who reviewed, approved, and sent every communication.

The BLAKE Framework

Recognising that simply switching on an AI tool without structure or governance would be insufficient, the BLAKE framework was developed to ensure every aspect of the implementation was deliberate, ethical, and student-centred. Each component addressed a distinct aspect of responsible AI integration in a compliance-sensitive environment, and the framework was subsequently used to train compliance staff on the tool's appropriate and consistent use.

B — Baseline Identification. The implementation began with a review of existing operational pain points. The primary baseline problem identified was volume: a high number of repetitive, routine enquiries were consuming significant staff time and energy and left limited capacity for complex queries and higher-value work. This baseline assessment informed the decision to focus the assistant specifically on high-frequency, policy-answerable enquiries rather than attempting to automate the full breadth of compliance communications.

L — Language Development. Once the baseline was established, the prompt structure that governs how the assistant receives and responds to enquiries was developed. Drawing on real student questions, structured prompts were crafted that reflected the team's preferred tone of clear, concise, and accessible communication. Critically, it was important that the language used remained a representation of the team's personality and values, to ensure that students would feel they were still communicating with the people they knew, rather than AI. For example, prompts specified that responses should avoid clinical jargon, use plain English, and follow a consistent structure aligned with the Faculty's communication standards. The assistant was explicitly instructed to produce responses that were easy for students to follow and free from unnecessary jargon.

A — Alignment with Verified Sources. To ensure accuracy and prevent the generation of unverified information, the assistant was configured to draw exclusively from approved institutional and government sources, including the MDHS Pre-Placement Compliance Handbook, the Australian Immunisation Handbook, ATAGI guidance, Standardised Student

Induction Protocols, and internal SOPs. Web search functionality was disabled to ensure that all responses remained grounded in verified, policy-aligned content.

K — Knowledge Guardrails. Clear boundaries were established for the scope of the assistant's responses. Most critically, the assistant was explicitly instructed that it is not a medical professional and must not provide medical advice under any circumstances. This guardrail was essential in a clinical professional education context where immunisation and health screening queries might otherwise prompt clinically inappropriate responses. These boundaries ensured the assistant remained a compliance communication tool, not a clinical decision-support system. Every response, regardless of how it was drafted, continued to reflect the professional judgement and human presence of the compliance team. While the specific guardrail in this case addressed clinical advice, the principle applies equally across disciplines. A WIL officer in a law faculty might instruct the tool not to provide legal advice. A business faculty team might specify that the tool must not make financial recommendations. In each case, the guardrail should reflect the professional boundaries that govern the team's role and the regulatory context in which they operate.

E — Evaluation. Ongoing evaluation was built into the operational model from the outset. Draft responses were reviewed on a weekly basis, with feedback drawn from staff observations, student responses, and identified inaccuracies in generated drafts used to refine prompts, update source materials, and improve response quality over time. This iterative evaluation process has continued beyond initial deployment, reflecting a commitment to continuous improvement and responsible AI governance. Staff retain full accountabilities for all communications sent, with the assistant functioning as a drafter rather than a decision-maker at all times.

Evaluation Approach

Outcomes were assessed informally through reflection and direct feedback received from students and both academic and professional staff, including observations about response quality, timeliness, and student understanding. During weekly reviews, draft responses were assessed against three informal criteria: factual accuracy relative to verified source documents, consistency with the team's established communication tone, and appropriateness for the student's specific compliance situation. Responses identified as inaccurate, inconsistent, or requiring significant rewriting were used to refine the prompt structure and update source materials. System usage data that included the number of enquiries resolved with AI-assisted drafting was tracked operationally rather than through a formal research data collection protocol. No formal baseline data were collected prior to implementation, as the decision to develop the tool was made in response to an operational pressure rather than a planned research intervention. Staff feedback was gathered informally through direct observation, team conversations, and unsolicited student responses received by academic and professional staff. These limitations are acknowledged in the discussion.

Results

Implementation of the GPT-based assistant was associated with observed improvements across response times, staff capacity, communication consistency, and stakeholder experience. These outcomes were observed across the full implementation period from November 2024 through to the time of writing.

Response Times

Prior to implementation, response times for student compliance enquiries could extend to up to four business days, particularly during high-demand periods such as at semester commencement. Following deployment of the assistant, average response times improved to same-day turnaround under standard conditions. Critically, even during peak periods that were historically the most resource-intensive times for the compliance team, the response times were maintained at under two business days.

Staff Capacity and Time Savings

During peak periods the compliance team received between 80 and 120 student enquiries per week, the vast majority of which were addressed with AI-assisted drafting. The assistant resolved 1,064 enquiries in 2025 and a further 443 enquiries in April 2026. Based on an estimated average drafting time of ten to fifteen minutes per enquiry, this corresponded to approximately 177 to 266 hours in 2025 and a further 74 to 111 hours in April 2026. These estimates were not formally tracked through time-recording systems.

This reclaimed capacity was redirected toward higher-value activities, including compliance workshops, stakeholder presentations, and the management of complex student queries that required professional judgement beyond the scope of the assistant. The impact was particularly significant given that pre-placement compliance for the entire Faculty is managed by a team of two to three full-time staff—a resource profile that has not grown proportionally with student numbers or regulatory complexity.

Communication Consistency

Prior to implementation, there was a risk of variability in responses across staff members, with different team members potentially providing answers with differing structures, tone, or levels of detail. The assistant introduced a standardised drafting baseline, ensuring that all student-facing communications followed a consistent structure, tone, and policy-aligned framing. This consistency reduced confusion among students and created a more coherent communication experience across the compliance function. Staff reported that the team was now effectively speaking the same language in all student interactions.

Staff Wellbeing

Staff reported perceived improvements in workload management, with less time spent on repetitive drafting tasks and a greater capacity to focus on complex enquiries. For example,

staff noted that the reduction in routine drafting allowed more time for direct student engagement and proactive compliance guidance during orientation periods.

Student and Stakeholder Experience

Student feedback, received directly by both academic and professional staff, indicated that compliance communications had become more timely, informative, and easier to understand following implementation. While this feedback was not collected through a formal survey instrument, its reach across both academic and professional staff cohorts suggests a consistent and broadly recognised improvement in communication quality. Stakeholder relationships improved as staff members were able to extend the compliance function's reach to a greater number of students despite the team remaining small. No privacy breaches or compliance concerns were recorded during the implementation period, affirming the effectiveness of the tool's governance safeguards.

Discussion

AI as Augmentation, Not Replacement

This case illustrates that meaningful AI integration in higher education administration need not involve large-scale technological transformation. A well-scoped tool that augments but does not replace human workers can address real administrative pain points while preserving institutional accountability (Luckin et al., 2022). The BLAKE framework demonstrates that the value of generative AI in compliance administration lies not in the sophistication of the technology itself, but in the governance structures, human oversight, and continuous improvement practices that surround it. This finding aligns with broader literature suggesting that human-AI collaboration is most effective when roles are clearly delineated and human judgement remains central to decision-making (Mosqueira-Rey et al., 2022; Salloch & Eriksen, 2024). Critically, this case provides a practical example of how HITL principles identified in the literature as essential but underspecified in practice can be operationalised in a small compliance team context. A notable dimension of this implementation was the deliberate preservation of the team's voice and personality in all AI-assisted communications. By crafting prompts that reflected the team's tone, values, and communication style, students continued to experience responses as coming from people they knew rather than an automated system. This human-centred approach to language design is an underexplored but important dimension of ethical AI deployment in student-facing administrative contexts.

Privacy, Legal Obligations, and Cybersecurity

Privacy controls, source restrictions, and review protocols were vital not only to user trust and institutional approval, but also to meeting the University's legal obligations and managing cybersecurity risks inherent in AI-assisted workflows. These governance decisions were critical in enabling institutional trust. The step seeking legal review before deployment was not onerous and took less than three weeks but proved essential in building institutional confidence in the tool.

Factors Enabling Success

Several contextual factors contributed to the success of this implementation and merit consideration by institutions seeking to replicate the approach. First, the existence of strong, well-maintained compliance documentation prior to implementation provided a reliable foundation for source alignment. Without authoritative internal documentation the assistant's outputs would have been significantly less reliable. Second, and perhaps most critically, the implementation was driven by a genuine commitment to prevent staff burnout and maintain the highest standard of student service. Between November 2024 and May 2025, the compliance function operated with only two full-time staff managing faculty-wide responsibilities. This resource pressure was not merely an operational challenge but also the motivating force behind the innovation. A continuous improvement mindset, willingness to research emerging tools, and personal drive to act on identified problems were instrumental in translating the idea into practice. The qualities of curiosity, persistence, and a genuine commitment to improving the student experience are not rare. What this case suggests is that they are sufficient to drive meaningful AI innovation, even without technical expertise or institutional support.

Limitations and Unintended Consequences

This case study has several limitations that warrant acknowledgement. First, as a single-institution implementation, the findings cannot be generalised beyond the MDHS context without further research. Second, outcome data were gathered informally rather than through structured instruments, which limits the robustness of before-and-after comparisons. Third, feedback was largely gathered from staff who were early adopters of the tool, which may introduce a degree of positivity bias in reported outcomes.

In terms of unintended consequences, the assistant is not infallible. There have been instances where generated drafts contained inaccuracies or required significant correction. This reinforces rather than undermines the case for the human-in-the-loop model: when errors occur, staff identify them, correct the response, and use the feedback to refine the assistant's prompts and source materials. This iterative correction process is a core feature of the BLAKE framework's Evaluation component and has been essential to maintaining response quality over time. Institutions adopting similar tools must recognise that ongoing human oversight is not optional and it is the mechanism by which AI-assisted workflows remain accurate, trustworthy, and safe.

A further risk worth monitoring is the potential for over-reliance on the tool over time, which could gradually erode staff capacity for independent policy interpretation. This is mitigated through regular team engagement with source materials and by ensuring staff remain active reviewers rather than passive approvers of AI-generated content.

Future research should explore more formal evaluation designs to strengthen the evidence base for AI-assisted compliance tools, including longitudinal studies of staff

wellbeing, response quality audits, and multi-site case studies across different institutional and disciplinary contexts. From a practice perspective, the BLAKE framework could be extended to other compliance domains, including partner site communications, clinical supervisor onboarding, and interprofessional placement coordination. These extensions are particularly relevant anywhere that high enquiry volumes and regulatory complexity create pressure on small administrative teams. Future research might also explore student perspectives on AI-assisted compliance communication, including whether response quality and tone are perceived differently when AI drafting is involved.

Transferability and Implications for WIL Practice

The transferability of the BLAKE framework has already been demonstrated in practice across both institutions and functions. Following its first public presentation at the NAFEA conference in November 2025, staff at the University of New South Wales (UNSW) were trained in the application of the framework. UNSW has since applied BLAKE across multiple administrative functions that include student communication, data analysis, and reporting to demonstrate that the framework extends well beyond its original context of student-facing enquiry management. At UNSW, one staff member reported a saving of approximately three to four hours per week that were previously spent on manual data processing. This cross-institutional and cross-functional adoption provides early external validation that the BLAKE framework is not limited to student-facing communication but can be applied to a range of administrative tasks where structured governance, verified sources, and human oversight are required.

This implementation offers a practical and replicable model for WIL departments seeking to integrate AI into compliance administration. Importantly, the BLAKE framework is not specific to the ChatGPT platform. Each component; Baseline, Language, Alignment, Knowledge guardrails, and Evaluation, can be applied to any generative AI tool, including institutional platforms such as Microsoft Copilot, Google Gemini, or other university-preferred AI environments. This platform flexibility strengthens the transferability of the framework across institutions with different technological infrastructures or data governance requirements. Institutions wishing to adopt a similar approach should consider the following practical steps:

1. Begin with a baseline audit of high-volume, repetitive enquiry types to identify where AI-assisted drafting will have the greatest impact. For example, reviewing frequently asked questions (FAQs) to identify the three or four topics that generate the most enquiries.
2. Invest in the quality and currency of internal documentation before configuring any AI tool, as the reliability of outputs is only as strong as the sources provided. Outdated or incomplete documents will produce unreliable responses.
3. Establish clear governance structures prior to deployment, including legal or privacy review, de-identification protocols, and defined staff review responsibilities. Ensure no personal or sensitive data is entered into the system at any point.

4. Develop structured prompts that reflect institutional tone, communication standards, and compliance requirements. Specify the voice, format, and boundaries of every response the tool will generate.
5. Build evaluation into the operational model from day one, with regular review cycles to identify errors, refine prompts, and update source materials. Treat the first month of outputs as a calibration period rather than a finished product.
6. Invest in staff training and trust-building, recognising that technical implementation alone is insufficient without accompanying change management. Staff need to understand both how to use the tool and why human review remains essential.

While these steps are illustrated through a health compliance context, they apply equally to any WIL office, student services team, admissions office, or administrative function managing high volumes of policy-answerable enquiries. The BLAKE framework offers a structured starting point for this process and may be adapted to suit the compliance contexts of other faculties, institutions, or disciplines particularly those operating in regulated health, legal, or social services environments, as well as any administrative context where high enquiry volumes and the need for accurate, consistent communication create operational pressure.

Further recognition of the framework's sector relevance is reflected in its presentation at the Work-Integrated Learning New Zealand (WILNZ) conference in April 2026, underscoring the growing interest among WIL practitioners in governance-led approaches to AI integration in administrative practice.

Conclusion

This case study suggests that generative AI can be integrated into WIL compliance administration in a manner that is ethical, scalable, and responsive to operational demands regardless of the AI platform adopted and without requiring significant financial investment or technical infrastructure. The GPT-based assistant developed and implemented at MDHS has delivered measurable improvements in response times, staff capacity, communication consistency, and student experience, while maintaining the privacy safeguards and human oversight essential in a compliance-sensitive administrative context.

This case study also illustrates a broader principle: that meaningful innovation in higher education administration is often driven not by large-scale institutional programmes, but by individual practitioners motivated by a commitment to service excellence and staff wellbeing. As WIL programs continue to grow in scale and complexity bringing with them increasing compliance burdens, regulatory requirements, and student support responsibilities human-AI collaboration may become essential to sustaining safe, timely, and student-centred support systems. Any professional staff member with access to a generative AI tool, a set of verified internal documents, and a clear understanding of their team's most frequent enquiries has everything required to begin.

Credit Statement

Sole responsibility is confirmed for the following: study conception and design, tool development and implementation, data collection, analysis and interpretation of results, and manuscript preparation.

Acknowledgements

This project did not receive any external funding. No conflict of interest is declared; this case study describes an initiative developed and implemented within a professional capacity. Blake Knapping, a valued colleague of the MDHS Pre-Placement Compliance Team, is gratefully acknowledged. Sincere thanks are also extended to the National Association for Field Experience Administrators (NAFEA) and the Work-Integrated Learning New Zealand (WILNZ) community for the opportunity to share this work at a national and international level, and to the University of New South Wales for their openness in adopting the BLAKE framework. Finally, heartfelt gratitude is expressed to Erica Cheng, whose love and support sustained this work from beginning to end.

References

- Australian Health Practitioner Regulation Agency Accreditation Committee. (2023). Information paper: Good practice approaches to embedding clinical placements, pedagogical innovations and evidence-based technological advances in health practitioner education. Australian Health Practitioner Regulation Agency. <https://www.ahpra.gov.au/documents/default.aspx?record=WD23%2f33307%5bv2%5d&dbid=AP&chksum=Xut6gosIEZJxExW5OAHtpUo7G6bopnG4MFG%2bIUA5RB0%3d>
- Cameron, C. (2017). The strategic and legal risks of work-integrated learning. *Asia-Pacific Journal of Cooperative Education*, 18(2), 87–99.
- Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, 20(1), Article 22. <https://doi.org/10.1186/s41239-023-00392-8>
- Fleming, J., & Hay, I. (2021). Understanding the risks in work-integrated learning. *International Journal of Work-Integrated Learning*, 22(2), 167–181. https://www.ijwil.org/files/IJWIL_22_2_167_181.pdf
- Luckin, R., Cukurova, M., Kent, C., & du Boulay, B. (2022). Empowering educators to be AI-ready. *Computers and Education: Artificial Intelligence*, 3, Article 100076. <https://doi.org/10.1016/j.caeai.2022.100076>
- Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Bascarán, J., & Fernández-Leal, Á. (2022). Human-in-the-loop machine learning: A state of the art.

Artificial Intelligence Review, 56, 3005–3054. <https://doi.org/10.1007/s10462-022-10246-w>

- Ocen, S., Elasu, J., Aarakit, S. M., & Olupot, C. (2025). Artificial intelligence in higher education institutions: Review of innovations, opportunities and challenges. *Frontiers in Education*, 10, Article 1530247. <https://doi.org/10.3389/feduc.2025.1530247>
- Salloch, S., & Eriksen, A. (2024). What are humans doing in the loop? Co-reasoning and practical judgment when using machine learning-driven decision aids. *The American Journal of Bioethics*, 24(9), 67–78. <https://doi.org/10.1080/15265161.2024.2353800>
- Yin, R. K. (2018). Case study research and design methods (6th ed.). SAGE Publications. <https://www.sagepub.com/shop/buy-a-book/case-study-research-and-applications-6-250150>



Work Integrated Learning

in Practice

About the Journal

WIL in Practice exists to disseminate good practice and learnings in work integrated learning (WIL), forming a valuable resource for NAFEA members and the broader WIL community. Publishing research and other suitable publications from academics, administrators and others whose interests align with the practice and administration of WIL, the journal will add to, and enhance, the existing body of knowledge that currently exists about WIL in all its various forms.

Early-career and emerging researchers and writers are also encouraged to submit their work. WIL in Practice provides an outlet for professionals to publish in a variety of formats.

WIL in Practice is supported by the National Association of Field Experience Administrators Inc. (NAFEA). WIL in Practice aims to be a high-quality and internationally recognised journal, publishing research and other suitable manuscripts from academics, administrators and the broader community whose interests align with the practice and administration of WIL. The journal will:

- Provide a welcoming, supportive and educative environment for authors at all levels to publish quality work
- Meaningfully contribute to the distribution of knowledge and professional development of WIL practitioners and researchers
- Publish at least once annually
- Take no action or inaction that is contrary to the values or directions of the NAFEA Executive Management Committee
- Operate in accordance with all laws and relevant policies.
- WIL in Practice is proudly supported by the National Association of Field Experience Administrators Inc. (NAFEA).

Editorial Board

Editor-in-Chief: Ondine Bradbury, Deakin University

Managing Editor: Dr. Ricky Tunny, QUT

Associate Editors:

- Dr. Sheree Lloyd, Griffith University
- Dr. Judith Needham, Griffith University
- Associate Professor Matthew Winslade, Charles Sturt University
- Dr. Simon Thornton, RMIT University
- Dr. Anita Dewi, Charles Sturt University

Copyright Notice

In line with the philosophy of *WIL in Practice* to disseminate publications to a broad audience, the journal uses Creative Commons - Attribution International licensing. When submitting manuscripts for publication, authors agree to the Creative Commons - Attribution International licence. Publications in *WIL in Practice* are able to be distributed, adapted, and used without seeking permission, provided that the original authors are appropriately attributed. Copyright is retained by the authors.