Agentic Work-Integrated Learning: The RaD Engage Approach

White paper · 21 November 2026

Mahdi Abolghasemi, Senior Lecturer in AI and Data Science, Queensland University of Technology, and Founder of RaD Engage · mahdi@radengage.com

Executive summary

RaD Engage is an AI-powered, research-grounded work-integrated learning platform. AI agents turn an industry partner's real problem into a unit-ready student project and draft feedback on students' work, and a human expert decides at every gate.

Work-integrated learning is valuable but costly to design: partners lack time to scope projects, coordinators must check each project against unit outcomes, and students need relevant projects and timely feedback (Patrick et al., 2008; Jackson, 2015). RaD Engage addresses all three in one system.

- **AI that drafts, people who decide.** Specialised AI agents do the drafting; coordinators and industry partners approve.
- **Co-creation.** Partners answer a handful of plain-language questions, and the platform drafts a project fitted to the unit.
- **Quality control.** Drafts pass through coordinator review, controlled edits, validated data and an audit trail before any student sees them.
- **Relevance.** Students are steered toward projects that match their career interests, which the motivation research links to interest and performance (Hulleman & Harackiewicz, 2009).
- **Access and opportunity.** Projects are completed online around a student's own time and place, and a short set of questions lets even a small business host one, so remote, regional and less-advantaged students and small businesses can be connected (Section 10).
- **Efficiency.** The platform takes over drafting and first-pass work. Our working estimate, drawn from interviews with WIL coordinators, lecturers, university staff and industry professionals, is 100 to 150 hours of stakeholder effort saved per project, which we are setting out to measure.
- **Feedback from several perspectives.** An AI mentor, the unit coordinator and the industry partner each contribute before the student sees a result, in line with research on effective feedback (Hattie & Timperley, 2007; Nicol & Macfarlane-Dick, 2006).

Research-driven by design. Each component is grounded in peer-reviewed work on constructive alignment, work-integrated and project-based learning, feedback literacy and multi-agent language-model systems. We treat RaD Engage as a research platform as well as a product, and Section 9 sets out the studies through which we are building our own evidence base.

1. The problem: good WIL projects are expensive to design

Work-integrated learning (WIL) is valued by students, universities and employers, but a well-designed WIL experience takes scarce academic and industry time to set up. The Australian national scoping study of WIL, which drew on 35 universities and almost 600 participants, was the first large-scale mapping of the field and found wide variation in how WIL is defined, designed and supported (Patrick et al., 2008). Later work shows that the quality of design matters: students develop employability skills when activities are structured, assessed and supported, and the barriers they report often concern the experience itself rather than the idea of WIL (Jackson, 2015). Billett (2009) argues that the educational value of workplace experience is not automatic; it depends on how experiences are sequenced, supported and connected to curriculum.

The problem is also unevenly spread. Universities Australia's national audit of WIL found that students from regional and remote areas, low socioeconomic backgrounds and Indigenous backgrounds took part at lower rates than other students (Universities Australia, 2019). Equity students describe time pressures, financial responsibilities, caring commitments and geographic location as the barriers (Bell et al., 2022). On the employer side, small and medium enterprises have real problems to offer but are often held back by limited resources and supervision capacity (Jackson, Shan & Meek, 2022). The students who face the most barriers, and the smaller businesses with the least capacity, are therefore the least likely to be connected by traditional placements.

Three frictions recur for the people who run WIL programmes.

- **Industry partners** have real problems but rarely have the time to turn them into a scoped, assessable student project.
- **Academics and unit coordinators** must check that every project fits the unit's learning outcomes and assessment criteria, and that quality is consistent across a cohort.
- **Students** need projects that feel relevant to their career, and need timely, specific feedback on what they submit.

RaD Engage addresses these three frictions with one connected system rather than three separate tools.

2. Pedagogical foundations

Each major component of RaD Engage is anchored in an established body of educational research, summarised below. The design choice comes first; the evidence base for it follows.

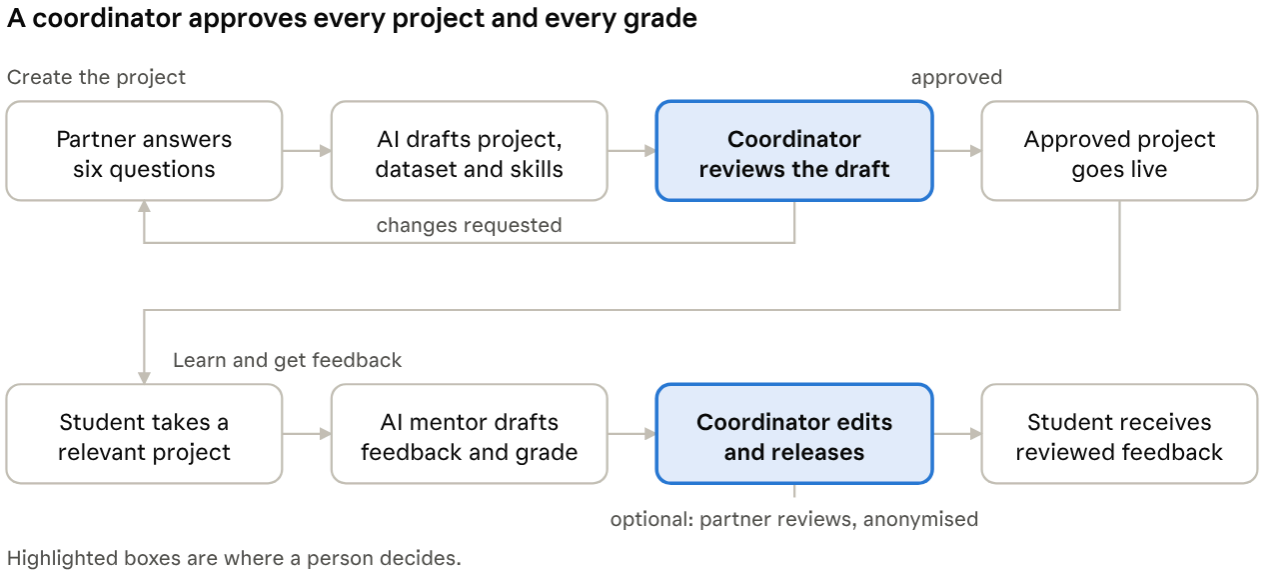
System component	Pedagogical principle	Key sources
Projects fitted to unit learning outcomes and assessment	Constructive alignment: outcomes, learning activities and assessment are designed together	Biggs (1996); Biggs & Tang (2011)
Authentic industry projects as the learning vehicle	Work-integrated and project-based learning	Patrick et al. (2008); Billett (2009); Blumenfeld et al. (1991)
Recommending projects relevant to a student's goals	Perceived relevance (utility value) raises interest and performance	Hulleman & Harackiewicz (2009)
Feedback from several perspectives, with a human review chain	Feedback is most useful when it is specific, actionable and supports self-regulation	Hattie & Timperley (2007); Nicol & Macfarlane-Dick (2006); Carless & Boud (2018)
Several specialised AI agents plus review stages	Multi-agent critique and iterative refinement of language-model output; model-based evaluation with known biases	Du et al. (2023); Madaan et al. (2023); Zheng et al. (2023)
Human experts in the loop at every approval point	Human oversight and critical engagement in educational AI	Zawacki-Richter et al. (2019); Kasneci et al. (2023)
Projects completed online, around the student's time and place	Online WIL can ease cost, time and location barriers for equity students	Bell et al. (2022); Universities Australia (2019)

Constructive alignment is the organising idea. Biggs (1996) describes it as combining constructivist learning theory with the alignment of intended learning outcomes, teaching activities and assessment tasks. Because RaD Engage supplies the unit's outcomes to the generator and has the coordinator review each draft against them, alignment is part of project creation rather than an afterthought.

The relevance principle explains why matching students to projects matters. In a randomised field experiment, Hulleman and Harackiewicz (2009) found that helping students see the connection between coursework and their own lives increased interest in the subject, with the strongest effect for students who doubted their competence. RaD Engage applies the same idea at project level by steering students toward projects relevant to their studies and goals.

3. System overview: two stages, two human gates

RaD Engage has two stages, and a person decides at the end of each. In the first, an industry partner's answers become a project that a coordinator approves. In the second, students work on approved projects and receive AI-drafted feedback that a coordinator reviews and releases.



RaD Engage flow · two stages, two coordinator gates

The AI agents do the drafting in both stages. The coordinator approves the project before any student sees it, and edits and releases the feedback before any student receives it. A partner who asks for changes sends a draft back around the first loop, and the partner can also add an anonymised practitioner's review to a submission before release.

4. Project co-creation: subject-matter experts in, aligned drafts out

AI agents do the drafting; subject-matter experts keep control. An industry partner does not write a brief. They answer a few plain-language questions in three steps, and the platform turns the answers into a draft project for the unit coordinator to review.

1. **The business situation.** What is the real-world problem, and what role will the student play and report to?

2. **Data and challenge.** What data will the student work with, and what makes the problem harder than it looks?
3. **Outcomes and deliverables.** What would a strong result look like, and what should the student hand in and for whom?

Every question is optional, so a partner with little time can still start a draft. The answers are stored with the project as its problem statement, data context and expected outcomes, and sent to the generator as the partner's own words.

When the unit coordinator shares a co-design link, the partner's session also carries the unit's requirements into generation: its learning outcomes, level, marks, duration, team size and required skills. This is how a partner's real problem is fitted to a unit's assessment design without the partner needing to know that design.

The draft includes tasks, a dataset specification, the skills the student will gain and, where requested, marking criteria. The platform builds and validates the dataset itself rather than trusting model-written data. The partner can ask the AI to revise tasks or data (for instance, "make it harder"), and nothing is saved until the partner saves.

The platform uses a multi-agent approach, in which specialised agents each contribute a different kind of expertise to the brief, such as analysis of the problem, curriculum design and data strategy. This follows research on multi-agent critique, where several model instances reason over the same problem and improve the result (Du et al., 2023).

What alignment means today. The unit's learning outcomes are supplied to the generator as context, and the coordinator checks alignment when approving the draft. Storing outcomes as structured records, mapping each task to an outcome and reporting coverage is on the research roadmap (Section 9). That mapping is the step that would turn constructive alignment (Biggs, 1996) from a design intention into something measurable.

5. Quality control: no AI-generated project reaches students unreviewed

Quality control in RaD Engage is a set of review gates and controlled-change rules around the AI, so that generated content is always checked by a person before a student sees it.

Control	What it does
Draft status	A draft project is never shown to students. Only an approved, active project is.

Control	What it does
Coordinator review	A partner's draft is submitted to the unit coordinator, who approves it or requests changes with a written note.
Partner-controlled edits	AI revisions to tasks or data are proposals. Nothing is saved until the partner saves it.
Held rewrites	An AI rewrite of an existing project is held as a proposal for an administrator to apply or reject.
Data validation	The platform generates and validates datasets rather than trusting model-written rows.
Version history and audit trail	Prompt configurations are versioned, and review decisions and AI usage are logged, so any result can be traced.

The platform also includes an offline test harness for checking prompt changes. It is a regression check, not a validation study, and Section 9 says how this evidence will be strengthened.

The design choice that matters most for institutions is the order of control. The AI proposes, the subject-matter expert and the coordinator decide, and the system records who decided what. This reflects the finding in the educational-AI literature that human oversight and critical engagement must stay central (Kasneci et al., 2023; Zawacki-Richter et al., 2019).

6. The student experience: relevant projects, real feedback

Students work on projects drawn from real industry problems, and the platform steers each student toward projects that fit their career interests.

Recommendations. The platform matches projects to the career interests and skills a student has declared. It leaves out projects the student has already completed or whose skills they already hold at expert level, and it explains each suggestion as a new skill, a skill to improve, or progress toward a chosen career. The matching is rule-based rather than machine-learned, which makes each recommendation easy to explain and audit. The research case for it is the relevance effect: when students see how learning connects to their own goals, their interest and performance rise (Hulleman & Harackiewicz, 2009).

Access through the unit. Projects linked to a unit are available to students who redeem the unit's access token, so coordinators control who works on what.

Individual and team work. A project can be run individually or in teams with set minimum and maximum sizes. Team submissions go through the same review chain as individual ones.

Feedback. Students receive feedback on their submissions, but only after it has been reviewed and released (Section 8). Industry partners see students only under anonymised labels during review.

The project-based learning literature holds that authentic problems and sustained, supported work are what make projects motivating and educational (Blumenfeld et al., 1991). The platform's contribution is to make authentic projects available at the scale of a unit, rather than for a few students with the right connections.

7. Efficiency: taking set-up effort out of the process

RaD Engage's working estimate is that the platform removes roughly 100 to 150 hours of stakeholder effort from setting up and running a WIL project. The estimate is drawn from the author's interviews with professional WIL coordinators, lecturers and other university staff, and with industry professionals, about the time each stage takes in a conventional process. It is an informed estimate rather than a measured result, and we are validating it through a before-and-after time study (Section 9).

The table lists where the effort goes in a conventional process and what the platform does instead. It states no hours per activity, because none have been measured.

Activity	Conventional process	In RaD Engage
Scoping the problem into a student project	Partner and academic negotiate a brief over several exchanges	Partner answers six questions; the platform drafts the project
Creating a dataset	Partner or academic prepares and cleans data by hand	The server builds and validates a dataset from an AI specification
Fitting the project to the unit	Academic rewrites the brief against outcomes and marks	Unit requirements are supplied to the generator; the coordinator reviews
Drafting tasks and marking criteria	Written from scratch by the academic	Drafted by the AI, then edited by the partner and coordinator
Reference solution	Written by the academic	Generated on demand for lecturers as a Word file

Activity	Conventional process	In RaD Engage
First-pass feedback on submissions	Marker reads every submission	AI drafts feedback and a grade; the coordinator reviews and releases

The saving is in drafting and first-pass work. The platform does not remove the expert judgement: the coordinator still approves each project and releases each grade (Sections 5 and 8).

The platform's records of review activity support that time study across several units, which is the next step (Section 9).

8. Feedback from several perspectives, in a reviewed chain

The AI drafts the feedback; people review it. A student's submission is looked at from up to four perspectives before the student sees a result: an AI reviewer acting as an industry mentor, the unit coordinator, the industry partner and, finally, the student as the recipient who acts on it.

1. **AI reviewer.** The model is prompted to act as a senior industry practitioner and a good mentor. It calibrates expectations to the project's level and grades against the task's marking criteria when they exist.
2. **Unit coordinator.** Every AI grade is held for review by default and is invisible to the student until released. The coordinator can approve it as it stands, or edit it and release it. When they edit, the platform keeps the original AI grade and review next to the final version.
3. **Industry partner.** The coordinator can send a submission to the partner for a practitioner's view. The partner sees the student under an anonymised label, can approve or reword the feedback (but never change the grade) and can send it back with a note.
4. **Student.** The student receives feedback that has been through this chain.

This follows what the feedback literature asks for. Hattie and Timperley (2007) show that the effect of feedback depends on its type and how it is given, and Nicol and Macfarlane-Dick (2006) set out principles for feedback that builds students' ability to regulate their own learning, including clarifying what good performance is and giving information that helps students close the gap. Carless and Boud (2018) add that feedback only works when students can appreciate it, judge their work and act on it. Putting an industry practitioner's perspective beside the academic's gives students a view of quality that comes from outside the classroom, which is the purpose of work-integrated learning (Billett, 2009).

A measurable by-product. Because the original AI grade is kept when a coordinator edits, the platform holds a record of how often and by how much human reviewers change the AI's judgement. That is direct evidence of agreement between the AI and expert markers, and it can be reported to institutions without extra data collection.

Where the design goes next. Today the AI side is a single calibrated reviewer rather than a panel. A candidate next step is a self-critique pass in which a second model call reviews the first reviewer's feedback before a human sees it. This builds on work showing that iterative self-feedback improves model output (Madaan et al., 2023) and on evidence that model judges agree closely with human preferences while showing known biases that need checking, such as position and verbosity bias (Zheng et al., 2023).

9. Evidence, responsible AI and research roadmap

RaD Engage is built on published research, and its own evidence base is at an early stage. This section says plainly what has been shown and what is still to be shown.

Where the evidence stands. The design of each component draws on peer-reviewed work (Section 2), and the platform includes an offline regression harness for its prompts (Section 5). Evaluation of our own outcomes is at an early stage, and the research programme below is how we are building that evidence.

Responsible AI. Three design choices limit the risk of relying on AI output.

- Humans decide: coordinators approve projects and release every AI grade by default, and partners cannot alter a grade.
- Privacy: industry partners see students only under anonymised labels. The platform's privacy policy names the third-party AI providers it relies on, says processing may take place outside Australia, and states that AI outputs are not necessarily accurate and should not be treated as definitive assessments.
- Traceability: prompt versions, review decisions and AI usage are logged, so a given result can be traced to the configuration that produced it.

Limitations.

- Learning outcomes are supplied to the generator as text, so alignment depends on the coordinator's review rather than on an automated check.
- AI grading uses one reviewer persona, and its agreement with expert markers has not yet been reported.
- Project recommendations are rule-based, so they depend on how fully students describe their interests.
- The time-saving estimate has not been measured.

- Language models can produce plausible but wrong content, which is why every output passes through a person.

Research programme. We treat RaD Engage as a research platform as well as a product. The studies we plan are:

Study or build	What it would show	Status
Structured outcome mapping: each task linked to a unit learning outcome, with a coverage report	Constructive alignment that can be checked and reported	Not started
AI–coordinator agreement analysis from the retained original AI grades	How closely AI grades match expert markers, by unit and level	Data already captured; analysis to do
Before-and-after time study across several units	A measured figure for the effort saved	Not started
Self-critique pass before human review	Whether a second model call improves feedback quality	Candidate
Student and partner outcome surveys, with ethics approval	Perceived relevance, learning gain and partner satisfaction	Not started
Access and equity analysis by region, socioeconomic background and caring status, including the mix of partners	Whether the platform widens access to WIL and to small-business partners	Not started

The second row matters most for credibility. The data to answer it is already produced by the review chain, so the first published result can come from the platform's normal use.

10. The opportunity: connecting students, small businesses and universities

The larger aim of RaD Engage is to make university–industry connection routine, and to reach the students and businesses that conventional placement models reach least.

Reach that does not depend on location. Projects are completed online and around a student's own time, so a student in a regional or remote area, or one balancing paid work or caring responsibilities, can take on an industry project without relocating. Research on online work-integrated learning found that students valued its affordability and flexibility, and also showed what must be designed carefully: digital access, a private place to work,

the loss of workplace interactions and the difficulty of giving good feedback remotely (Bell et al., 2022). RaD Engage addresses the last of these directly, with structured feedback that a coordinator and a practitioner review.

Small businesses as hosts. Small and medium enterprises have real problems worth solving, but resource and supervision limits keep many of them from hosting students (Jackson, Shan & Meek, 2022). A short set of plain-language questions, AI drafting and light-touch review lower the cost of taking part: a partner needs no desk space, no day-to-day supervision and no teaching expertise to offer a well-formed project.

Employability for students, talent for companies. Students complete authentic projects on real problems, receive feedback from an academic and a practitioner perspective, and build a record of the skills each project develops. Companies receive fresh analysis of real problems and a low-risk view of how emerging graduates approach them. Partners see student work anonymised today. Opt-in introductions between students and partners, designed with student consent at the centre, are a natural extension.

One connection instead of many negotiations. A university unit often builds its own industry relationships from scratch. A shared platform gives coordinators, partners and students one place to scope, approve, complete and review projects, with a common record of what was done. This serves aims already set in national policy: the national strategy on WIL aimed to expand participation, including by smaller employers (Universities Australia et al., 2015), and the Universities Accord final report sets equity targets for rural, regional and remote students and calls for stronger industry–education partnerships (Australian Government Department of Education, 2024).

What we will test. These are design intentions, not yet demonstrated results. The research programme in Section 9 includes a study of who takes part and how they fare, by region, socioeconomic background and caring status, and of the mix of partners, including small businesses.

References

Australian Government Department of Education. (2024). *Australian Universities Accord: Final report*. https://www.education.gov.au

Bell, A., et al. (2022). Student and educator perspectives on equity and online work integrated learning. *Australasian Journal of Educational Technology, 38*(6). https://ajet.org.au/index.php/AJET/article/view/7524

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education, 32*(3), 347–364. https://doi.org/10.1007/BF00138871

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

Billett, S. (2009). Realising the educational worth of integrating work experiences in higher education. *Studies in Higher Education, 34*(7), 827–843. https://doi.org/10.1080/03075070802706561

Blumenfeld, P. C., Soloway, E., Marx, R. W., Krajcik, J. S., Guzdial, M., & Palincsar, A. (1991). Motivating project-based learning: Sustaining the doing, supporting the learning. *Educational Psychologist, 26*(3–4), 369–398. https://doi.org/10.1080/00461520.1991.9653139

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

Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2023). Improving factuality and reasoning in language models through multiagent debate. arXiv:2305.14325; published in *Proceedings of the 41st International Conference on Machine Learning* (PMLR 235, 2024). https://arxiv.org/abs/2305.14325

Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of Educational Research, 77*(1), 81–112. https://doi.org/10.3102/003465430298487

Hulleman, C. S., & Harackiewicz, J. M. (2009). Promoting interest and performance in high school science classes. *Science, 326*(5958), 1410–1412. https://doi.org/10.1126/science.1177067

Jackson, D. (2015). Employability skill development in work-integrated learning: Barriers and best practice. *Studies in Higher Education, 40*(2), 350–367. https://doi.org/10.1080/03075079.2013.842221

Jackson, D., Shan, H., & Meek, S. (2022). Enhancing graduates' enterprise capabilities through work-integrated learning in co-working spaces. *Higher Education, 84*, 101–120. https://doi.org/10.1007/s10734-021-00756-x

Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences, 103*, 102274. https://doi.org/10.1016/j.lindif.2023.102274

Madaan, A., Tandon, N., Gupta, P., et al. (2023). Self-Refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems, 36* (NeurIPS 2023). https://arxiv.org/abs/2303.17651

Nicol, D. J., & Macfarlane-Dick, D. (2006). Formative assessment and self-regulated learning: A model and seven principles of good feedback practice. *Studies in Higher Education, 31*(2), 199–218. https://doi.org/10.1080/03075070600572090

Patrick, C.-J., Peach, D., Pocknee, C., Webb, F., Fletcher, M., & Pretto, G. (2008). *The WIL (Work Integrated Learning) report: A national scoping study* [Final report]. Australian Learning and Teaching Council, Queensland University of Technology.

Universities Australia. (2019). *Work integrated learning in universities: Final report*. https://universitiesaustralia.edu.au/wp-content/uploads/2022/03/WIL-in-universities-final-report-April-2019.pdf

Universities Australia, Business Council of Australia, Australian Chamber of Commerce and Industry, Australian Industry Group, & Australian Collaborative Education Network. (2015). *National strategy on work integrated learning in university education*.

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

Zheng, L., Chiang, W.-L., Sheng, Y., et al. (2023). Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *Advances in Neural Information Processing Systems 36* (Datasets and Benchmarks Track). https://arxiv.org/abs/2306.05685