Generative AI and Academic Integrity:

Moving towards program level assessment design

CODE Webinar, 30 April 2025

The second of four webinars organised by the Centre for Online and Distance Education (CODE) during the summer term of 2025 covered a problem close to the top of almost every HE professional's priority list: how to assess whilst maintaining academic integrity in the age of generative AI (GenAI). It was held on Wednesday 30th April and the speaker was a long-standing, if far-flung friend of the Centre: **Professor Steven Warburton**, Pro-Vice-Chancellor for Education Innovation at the University of Newcastle, Australia and a CODE Fellow. Not surprisingly, the webinar drew a large crowd of enthusiastic delegates.

The webinar was ably facilitated by the CODE director, **Linda Amrane-Cooper**. She introduced Steven as an innovator with many years' experience of working in distance and online education in the UK, New Zealand and now Australia. He would be sharing his thoughts on GenAI and academic integrity and, intriguingly, introducing a 'two-lane approach' to the problem. Steven, in his own introduction, added that he would present an Australasian perspective on the topic: the university sector in Australia and New Zealand is relatively small, with around 50 institutions, and they collaborate closely enough that it is possible to talk of an 'Australasian way' of tackling academic challenges. And he ended his introduction, as Australian academics tend to, by acknowledging the First Nations custodians of the land and the Awabakal people who are indigenous to the place where the University of Newcastle now sits.

He then set out what the talk would not be about. It would not touch on the predicted future of AI or discuss exactly how powerful it is becoming. Anyone interested in that, or just keen to be amused, can investigate the 'blueberry muffin versus Chihuahua challenge' which is used as a test of pattern-matching ability. Instead, he would be discussing what students can and should be doing with the GenAI tools that are now readily available to them, how educators can work with it to prove student learning, and how unauthorised or unethical use can be detected.

GenAI: An Australasian Perspective

He summarised the current position on GenAI in higher education from an Australasian perspective, pointing out that

- The launch of the free-to-use GPT v.4 tool marks a 'democratisation' of GenAI, opening powerful tools to all, including students.
- Data privacy is taken seriously, and his university has chosen to endorse Microsoft Copilot, which is available to staff and students as a secure 'on network' platform. Other universities are also partnering with OpenAI to produce in-house tools.
- 'Stress testing' different types of assignment shows that GenAI can produce, for example, a creditable critique of a research article within 10 minutes.

- The Australian body that regulates the quality of higher education, TEQSA, recently asked all HE providers to submit a credible action plan to 'address the risk that GenAI poses to award integrity'.
- It is crucial to experiment with and understand the AI tools that our students now have in their hands.

Taken generally, the rise of AI in higher education can be said to be causing two basic problems, one chronic and one acute. The chronic problem is summed up as '*What and how do we teach when AI is everywhere*?'. We will need to think carefully about what should be taught in each discipline – skills over knowledge – and how to prepare students for an unknown, uncertain future. The acute, and urgent, problem is one of ensuring academic integrity: *how can we ensure that the assessments our students take, which determine the qualifications they receive, are fair and valid?*

Steven then pointed to survey data indicating that about three-quarters of HE staff in Australia are already using GenAI in and/or outside their work, and we know that students are as well. Access to suitable tools and training, clear guidelines and case studies have been the main drivers of this take-up. Ethan Mollick, a professor at the University of Pennsylvania in the US, has produced an engaging and frequently updated guide to 'which AI to use now?' in which he gives each tool a 'personality': describing Grok 2 from Elon Musk's X.ai, for instance, as 'sarcastic and fun'.

AI Detection: Academic Integrity Workflows

In order to manage students' use of AI in their assessed work, and therefore to maintain academic integrity, it is first necessary to detect whether, and when, it is being used. Even before GenAI had been thought of, many universities had set up workflows involving software like Turnitin as a similarity checker to detect potential plagiarism. Turnitin's developers, however, recently switched on a new tool for suggesting how much of each assignment may have been written by GenAI. Its use remains somewhat contentious as students are unable to see these results, as they can with Turnitin's standard similarity reports. Some Australian universities, including Newcastle (where it is under review) are keeping this tool running while others have disabled it.

Steven then presented a useful summary of four ways in which students appear to be using GenAI in their work:

- '**AI informed'**, using (e.g.) ChatGPT rather like Wikipedia, as a fact-checker;
- '**AI produced**', where AI output is copy-and-pasted with little modification;
- '**AI patterned'**, where the assignment is written in the student's own words to an AI-based 'scaffold';
- '**AI fixed'**, where the writing has been improved by GenAI.

The most popular of these, and the one which causes least concern, is the last: students without English as a first language can find this type of AI use particularly useful, and it can be thought of as a step up from Grammarly or even Word. Straight cut and paste is the easiest type to detect and is obviously unacceptable, but the other two are interesting. Students using AI in these ways can probably work ethically and also learn skills that will be important in the

workplace. Of course, any student use of AI in assessments, even when it is allowed in the course rubric, must be fully referenced.

Course coordinators need to know when a student's use, or abuse, of GenAI is serious enough to constitute actual academic misconduct, particularly as it is almost impossible to de-escalate a case that has been escalated to the academic misconduct office. At Newcastle the less serious cases are treated simply as poor academic practice and, taking an educative approach, these students are given extra training. Penalties are only applied where there has been a proven breach of academic integrity. And this only stresses the importance of clear communication to set expectations. Students must be told exactly when, and in what context, they are allowed to use GenAI in each assignment.

All this, however, assumes that it is possible to detect where GenAI has been used inappropriately, and this is getting harder as the tools improve. Steven quoted a study from the University of Reading (UK) that found that only one out of 33 AI-generated exam scripts was detected by the markers, and that the AI generally out-performed the students by half a grade boundary.

Assessment Design: Program Level Assessment

So, if it is so hard to detect unauthorised AI use, and if (as it seems) it is impossible to ban, how can assessments be designed to allow for it? In Newcastle, they have arrived at a strategy termed program level assessment design. If, as they estimate, about half the 6,000 or so individual assessment tasks set across the whole university are susceptible to the inappropriate use of GenAI, how can learning be assured? Programme-level assessment, or integrated programme assessment, is a form of assessment design that looks at an entire study programme, selects the key points where learning must be validated and makes sure that those assessments are secure. The strategy here is not to give students more and more invigilated exams, but to test learning outcomes at critical points so the university can be sure that, by the time they graduate, students are deserving of the award they are given.

The guiding principles for program level assessment design are underpinned by two publications: *Assessment reform for the age of artificial intelligence* (November 2023) and *GenAI strategies for Australian higher education: Emerging practice* (November 2024). The approach can be divided into four practical stages: **Map and analyse**, **Review and categorise**, **Assessment reform initiatives** and, finally, **Appropriate governance mechanisms**. The first phase involves mapping all assessment in a programme of study to identify those that must be made secure if the whole programme is to be valid. Then, in the second phase, these assessments are prioritised for redesign.

The University of Newcastle is taking a discipline-based approach with each discipline's norms, values and signature pedagogies being taken into account, in other words, this is not a silver bullet or 'one size fits all'. While secure invigilated exams will form a part of the new design – and at Newcastle, the business school really loves invigilated exams – programme leaders have a wide range of secure assessment strategies to choose from including interactive oral assessments, conditional grading, portfolios, online proctoring and competency-based grading. This approach can be built into an institution's governance cycles through continuous programme reviews and quality assurance, with secure assessments 'flagged' within the curriculum management system so they cannot easily be changed. In the short term, however,

as this work continues, academics must bear a few key considerations in mind: they should 'stress test' assignments, communicate expectations around GenAI use clearly and precisely and keep up to date in a fast-changing environment.

So is it Sorted? Frameworks and Culture Shifts

Many universities have developed a 'traffic light' system to grade assignments, where AI use may be forbidden (red), partially restricted (amber) or expected (green), but this hits problems as it tries to turn a complex situation into a simple linear one. If we say that AI use in a particular assignment is restricted, then what types of use are allowed or disallowed? A different approach, which might be more helpful, uses the analogy of a menu where lecturers select exactly what they want their students to use AI for. And this leads into a third analogy: the 'two lanes' that had been introduced at the beginning of the webinar. Steven explained that Lane One is for the assured assessments, including skills-based practice, interactive oral assessment and, of course, invigilated exams. Everything else is in Lane Two, where we should assume that AI will be used: in fact, we can think of this lane as 'human-AI collaboration, in assessment *for* and *as* learning'. There is no middle ground between the two.

Steven ended the presentation by highlighting the CRAFT framework for educators from work led by the University Sydney to help institutions work through developing a mature approach to teaching with and through AI:

- **Culture** (departmental, disciplinary, regional…)
- **Rules** (principles, policies, guidelines…)
- **Access** (equitable availability of AI and infrastructure)
- **Familiarity** (awareness, comfort, skill)
- **Trust** (between academics, students, leaders and AI)

He recommended that we assess our institutions, staff, researchers and students under each of these categories as an introduction to the culture shift that is required. We should aim to move from seeing AI as a servant to as a partner, from 'cheating' to learning, and from a threat to an opportunity; from policing it to understanding its possibilities; and from encouraging 'AI awareness' in our students to seeing its professional use as a graduate attribute.

Discussion

Linda thanked Steven for a brilliant, thought-provoking webinar and picked up on a few questions from the lively chat to begin the discussion.

- *What do you think the mandatory minimum percentage of controlled and assured assessments should be?* There are no strict figures available, but Steven suggested that a value of around 30% or one-third might be applicable in many cases. He is, however, wary of suggesting an exact figure as this will often be sensitive to disciplinary or accreditation requirements.

- *Have you started looking at the cost side of moving to this type of assessment?* Steven is particularly concerned about the cost of the invigilated exams that are often quoted as the 'gold standard' of assured assessments, for the students (particularly distance-

learning students) as well as for the institution. Online proctored exams are easier (if not always cheaper) to reschedule if required.

- *Must the 'second lane' (potential open use of GenAI) assessments be non-credit bearing?* Not necessarily. Steven illustrated this with an example of a lecturer in entrepreneurship at Newcastle who has embedded GenAI throughout his courses so much that students must use GenAI in their carefully designed assessments.

Linda then threw the discussion open to the floor. One delegate opened up the issue of assessment at the unit level: if students are able to switch institutions taking credits with them, as is the case in much of the EU, the programme-level approach will fall apart and each unit must include sufficient 'lane one' material. Steven acknowledged that this is a difficult question but suggested that a competency-based approach would be one way forward. Units of learning with similar assessment types can be matched.

The benefits of a competency-based approach were raised elsewhere in the chat, where delegates also emphasised the importance of observing learning and of setting assignments that test the upper levels of Bloom's taxonomy. There was also a short discussion of the value of AI in supporting neurodiverse students, with a shout-out for goblin.tools as a helpful platform for them, and other posters usefully stressed the importance of building trust.

Linda reluctantly drew the conversation to a close, highlighting that one of the remaining three CODE webinars this term, on 4 June, would focus explicitly on the issue of academic integrity.