AI Detectors in education

Associate Professor Mark A. Bassett

August 2025

Table of contents

Acknowledgement

This paper consolidates a series of articles on AI detectors in education based on collaborative work with Wayne Bradshaw, Hannah Bornsztejn, Alyce Hogg, Kane Murdoch, Bridget Pearce, and Colin Webber.

Part 1: Accuracy, deterrence, and fairness in practice

Educational institutions that utilise artificial intelligence text detection software (AI detectors) have grown accustomed to them. They appear to offer a straightforward, reportable, and visible solution to the academic integrity issues posed by unauthorised use of generative artificial intelligence (GenAI) by students in assessments. By using an AI detector, institutions appear to be taking action and doing *something*, often argued to be better than nothing, which is also frequently and erroneously framed as the only alternative. AI detectors provide institutional leaders with something to point to, an action to highlight as part of their approach to mitigating the academic integrity risks posed by GenAI. Spending $X on an AI detector appears to be an appropriate and effective measure to mitigate the risks associated with GenAI. In short, proponents of AI detectors argue that they are used because they are effective in identifying AI-generated text and deterring students from using GenAI in an unauthorised manner; they are used because they 'work'. But what does it mean for an AI detector to 'work'?

In simple terms, we want an AI detector to identify AI-generated text. But what percentage of AI-generated text is sufficient for our needs? If it identifies 80% of all AI-generated text in a paper, would that suffice? And how reliable is the AI detector? If it ignores every 20/100 AI-generated papers, but its accuracy within the other 80 is 100%, would that be preferred? As Phillip Dawson, author of *Defending assessment security in a digital world* noted, to create a 100% accurate AI detector, simply program the software to flag all text as AI-generated. The obvious issue with this approach is that the software will incorrectly flag human work as AI-generated. But this is also an issue with real AI detectors, so what percentage of these incorrect flags are we comfortable with? It's clear that a single metric for accuracy, such as "99% accurate", is not useful.

In Signal Detection Theory, the metrics related to these questions are False Positive Rate (FPR), True Positive Rate (TPR), False Negative Rate (FNR), and True Negative Rate (TNR). The FPR is the most widely published metric among developers of AI detectors. It represents the percentage of human-written papers that were *incorrectly* flagged as being AI-generated in training. A FPR of 1% indicates that the detector incorrectly flagged 1% of human-written papers as AI-generated. A low FPR is therefore considered desirable and is used by several companies to promote their AI detectors. The TPR indicates the percentage of AI-generated papers that were *correctly* flagged as AI-generated. The FNR represents the percentage of AI-generated papers that were *incorrectly* flagged as human-written, while the TNR

indicates the percentage of human-written papers that were *correctly* identified as such. However, these metrics are established in closed testing environments; they do not indicate accuracy in real-world environments.

Let's consider a hypothetical AI detector with a 100% TPR. During testing, it demonstrated high *sensitivity*, flagging every AI-generated paper as such. However, as a human can feasibly write *any* text that a GenAI system can produce, in the real world, a 100% TPR is meaningless. There is no AI-generated text that a human could not have written; human-written text and AI-generated text are not mutually exclusive sets.

Moving on to the heralded FPR, let's consider a different hypothetical AI detector with an FPR of 1% and a TPR of 90%. During testing, it incorrectly flagged 1/100 human-written papers as AI-generated, and correctly flagged 90/100 AI-generated papers as such. We have 100 essays to grade, and the AI detector has flagged the paper in front of us as 100% AI-generated. Given the AI detector's FPR of 1% and TPR of 90%, what are the chances that the paper in front of us is AI-generated? 99%? 10%? 100%?

The formula for determining this is:

$$\frac{\text{Number of papers flagged as AI that are actually AI-generated}}{\text{Number of papers flagged as AI-generated}}$$

which equals

$$\frac{\text{Number of papers flagged as AI that are actually AI-generated}}{\text{Total false AI flags + Total true AI flags}}$$

which equals

$$\frac{\text{TPR x number of AI-generated papers}}{\text{FPR x number of human-written papers + TPR x number of AI-generated papers}}$$

It's clear from the formula above that to determine the chances that a paper flagged as AI-generated is, in fact, AI-generated, we need to know three things:

1. The TPR
2. The FPR
3. The proportion of human-written and AI-generated papers in the sample of 100 papers.

We know the TPR and FPR. We don't know, and, *in practice, can never know,* the proportion of human-written and AI-generated papers in the sample. Therefore, **in practice, we can *never know* the likelihood that a paper flagged as AI-generated is AI-generated**.

Even if we knew the FNR and TNR, we'd be no closer to an answer. This highlights an all-too-common myth in the AI detector space: that these metrics—FPR, TPR, FNR, and TNR—are meaningful in practice, outside of the controlled training environments in which they were

derived, where the number of AI-generated and human-written papers is known. **In practice, these metrics alone tell us nothing about the chances that a flagged paper is AI-generated**. The paper in front of us could have anything from a 0%, 2%, 10%, 80%, or 99% chance of being AI-generated, but there's no way for us to know which. Institutions that use AI detectors due to their advertised accuracy are placing trust in figures that offer no practical value—such use should be abandoned on principle.

To establish the metrics above, developers of AI detectors need a set of documents that is guaranteed to be human-written. Turnitin tested its detector using 700,000 student papers submitted before 2019 to ensure that none of the papers contained any AI-generated content. This methodology, however, relies on an unsupported assumption: that student work from pre-2019 is directly comparable to student work today. Consider a student research project aimed at developing an AI detector that doesn't flag current student work by training it on student work from over five years ago. Any competent supervisor would advise the student that they need to first establish, using *evidence*, that student work from before 2019 is a valid comparison to current student work. Without establishing this, the testing of the student's AI detector is methodologically flawed. Developers of AI detectors have provided no such evidence.

AI detectors *do* expose more assessments to increased scrutiny. Given that we know from surveys, for example, that students are using GenAI in their assessments, the more assessments we scrutinise, the more academic misconduct we will find. If we subject the assessments of every student born on a Thursday to increased scrutiny, we will likely uncover more misconduct—misconduct that has absolutely nothing to do with being born on a Thursday. But where is the harm in exposing more student work to scrutiny? If they haven't committed misconduct, they have nothing to worry about, right? If only it were that simple.

Let's assume a hypothetical AI detector has a 100% TPR; every single assessment flagged as AI-generated is indeed AI-generated. Every student whose work is flagged by the detector has unequivocally used GenAI, and in this example, its use was unauthorised. Every student flagged has a case to answer. The FNR of this detector, however, is 20% as it incorrectly flags 20% of AI-generated papers as human-written. **But the papers that this and all other AI detectors *don't* flag are not random.** They are written by students with access to the best (most expensive) LLMs, by students with premium paraphrasing tool accounts, by students with computers at home, by students who are more technologically literate and who most likely come from a higher socio-economic status. It would be immeasurably fairer if we implemented the 'born on a Thursday' rule.

The additional scrutiny of assessments that is enacted by using AI detectors is akin in some ways to conducting a search of the student's assessment, looking for evidence of

misconduct. In a just system, the ethics of a search don't hinge solely on how much wrongdoing it uncovers. For example, we would not tolerate a policy that singles out welfare recipients for random searches, even if doing so uncovered more stolen goods than searching no one at all. The discovery of wrongdoing is clearly not inherently tied to being a welfare recipient; the same logic would apply to any selectively targeted group, including students whose work is flagged by AI detectors.

In education, this logic holds. The more assessments we subject to AI detection, the more misconduct we will find, but that doesn't justify a tool that disproportionately targets certain groups of students. AI detectors offer no reliable standard of evidence, no meaningful transparency, and no guarantee of fairness. While education doesn't require the criminal standard of probable cause, it does demand methods that are consistent, accountable, and trustworthy. AI detectors fail on all counts. What they do achieve—unequal surveillance—is not defensible simply because it sometimes produces results.

Some institutions use AI detectors as deterrents, based on the belief that students are less likely to use GenAI in an unauthorised manner if they believe that software exists that can detect its use. This is true in a narrow, instrumental sense, but deterrence alone does not justify this approach, especially when the mechanism used to deter students is as unreliable, opaque, and unequally applied as AI detectors. The mere fact that something works does not make it right; authoritarianism 'works'. Arrest everyone and only let them go when they've proven their innocence 'works' if the sole goal is reducing crime, but ensuring academic integrity *at all costs* is a recipe for injustice. A system that deters through fear of being (mis)judged, rather than clarity of expectation and appropriate enforcement, is not ethically defensible and shifts the burden of proof onto students without offering procedural safeguards in return. Even if AI detectors prevent student misconduct, they do so by compromising the values that should underpin assessment: fairness, transparency, and mutual respect.

When presented with the above, most institutions that use an AI detector will offer the predictable rebuttal that they're only used in conjunction with other evidence or act, as a "red flag", "part of a wider approach", or to "inform discussions with the student". This argument is disingenuous. In practice, the output of an AI detector does not sit passively in a list of concerns. It drives suspicion, initiates investigations, and implicitly reshapes the threshold for guilt. Once a piece of writing is flagged, it becomes tainted, even if the institution insists that decisions are made "holistically". In effect, the detector is not just a tool among many, it becomes the lens through which all other information is interpreted.

Part 2: Confirming the meaningless

Educators use several methods to attempt to corroborate the results of artificial intelligence text detectors (AI detectors). Using multiple detectors is common and based on the rationale that if multiple detectors agree that the text is AI-generated, this holds more evidentiary weight than the findings of a single detector. However, instead of providing additional, independent evidence of AI use, this approach merely reinforces the myriad issues associated with AI detectors as detailed above. Even if all detectors found that a paper was AI-generated, it would be no more validating than if a group of phrenologists agreed on a diagnosis.

Aside from failing to validate the original detector's result, submitting assessments to third-party software with which we don't have an agreement or which our IT department hasn't vetted is like playing Russian roulette with our students' work. Pick the wrong AI detector, and those assessments might end up in an essay mill. Imagine an innocent student being called to an academic integrity hearing to explain why their work ended up on a contract cheating website, when we are responsible for it being there.

Another common approach is to use the output of a generative artificial intelligence (GenAI) tool to verify suspected AI use. This involves submitting the assessment outline or question to a chosen Large Language Model (LLM), then comparing the output to each student's work. While often framed as a form of triangulation, this approach is shaped as much by confirmation bias as by evidence. When educators generate LLM outputs to compare with students' work, they often do so with the presumption of AI use based on the flawed results of an AI detector. LLMs generate text by drawing on patterns in their training data, often producing predictable responses to standard prompts. Students' work may follow similar patterns, not because it's AI-generated, but because students respond to the same task within the same academic conventions. Structural or thematic similarity may simply reflect the constraints of the assignment and is not evidence of GenAI provenance. Trying multiple LLMs until we end up with an output that resembles a student's work is an exercise in retrofitting suspicion to match a predetermined conclusion.

Comparisons of students' work to the output of a GenAI tool are often based on "similarity", "general structure", and "similar presentation". Interpretations of similarity will, however, inevitably vary, with some staff convinced of GenAI use, while others remain unconvinced. Such judgments are inherently subjective and inconsistent without predefined benchmarks, turning what appears to be evidence into little more than intuition shaped by suspicion. Such arbitrary judgments have no place in an academic integrity investigation. And if institutions did implement benchmarks, what would a student think if their efforts were flagged as "insufficiently human"?

Asking a GenAI tool if it wrote a student's assessment is another method that's been observed. Putting aside the issue that by entering a student's intellectual property into a GenAI tool, we are very likely breaching its terms of service and potentially the student's rights, LLMs can't identify AI-generated text. They *will* occasionally tell us with conviction that they wrote the student's text we uploaded, but these claims are baseless and must be ignored. Likewise, denials of authorship by GenAI tools are equally unfounded.

Falsified references are a favourite among educators as they provide seemingly concrete evidence of GenAI use. However, educators don't need to go down the "unauthorised AI use" route, as falsifying references is itself academic misconduct. It's irrelevant how these falsified references were generated; the offence is the student's submission of them, not how they were procured. If we're angling to penalise a student for *both* falsified reference*s* and unauthorised GenAI use (based on falsified references), unless each offence involved *distinct, separate actions*, this is the opposite of an educative approach to academic integrity, as we're effectively charging them twice for the same act. Pursuing unauthorised GenAI use in this case is unnecessarily punitive, adds little and risks conflating the issue, where the actual breach is already clear, demonstrable, and sufficient.

The presence of linguistic markers is arguably the most pervasive approach educators take to corroborate the results of an AI detector. As discussed above, there is no AI-generated text that a human could not have written; human-written text and AI-generated text are not mutually exclusive sets. Despite this, linguistic markers such as em dashes, semicolons, colons, lists, paragraphs starting with firstly, secondly, thirdly, formulaic prose, the use of specific words like "delve" or "tapestry", using a "z" instead of "s", and American spelling continue to be used as evidence of GenAI use. These elements occur in AI-generated writing because they are present in the human texts on which the models were trained. They are not valid indicators of AI-generated text, on their own, or in combination with an AI detector's result. Staff who seek to penalise students based on the presence of these elements rarely pause to consider whether exposure to the GenAI outputs increasingly found in advertisements, social media, and even journal articles has influenced their writing.

Getting hit by a car when crossing the road after reading a horoscope that warned you to "tread carefully today" doesn't legitimise astrology. Likewise, when a student admits to using GenAI, it doesn't corroborate the results of an AI detector. To take a student's admission as evidence of the detector's accuracy is a dangerous case of confusing correlation with causality. Students may also confess for reasons unrelated to misconduct. Academic integrity investigations can be intimidating for students, with a clear power imbalance that must be carefully managed. It's not unreasonable to believe that a student might feel pressured to admit to something that "we know you did" (based on the AI detector's flawed results, when, in fact, we don't know) to make what can be an ordeal end, particularly when reduced penalties are on offer for early admissions.

If we seek a subjective method to judge students unfairly, look no further than comparing their work to “past writing style”. Primarily, this approach is led by confirmation bias. Differences between current and past work, when viewed with suspicion through the lens of an AI detector’s result, are taken to indicate misconduct, not natural variation. A rushed in-class reflection will read differently from a carefully revised capstone project, and a second-language writer’s phrasing will mature across semesters. Comparison to past writing styles reverses the burden of proof onto students, requiring them to prove that the result of genuine progress, stress, or illness is not machine-generated.

For the cavalier looking to depart from the principles of mutual respect, trust, and the same ethical standards that we expect from students, hiding prompts within the instructions of an assessment—whether as invisible text, cryptic phrases, or contrived references designed to appear only in an AI-generated answer—offers an opportunity to entrap students. If we denounce cheating as dishonest, we cannot smuggle hidden prompts into assessments and maintain any moral authority.

Many institutions now ask students to show evidence of their writing process to prove they have not used GenAI. If drafts or similar materials may be required, the assessment brief must state this. Penalising students who cannot produce documents they were never told to keep is procedurally unfair. Using software that tracks students' writing, including the number of edits, typing speed, deletions, construction time, etc., without clearly defined benchmarks is a recipe for procedural unfairness and successful appeals. How fast is too fast? How few edits are permitted? These benchmarks must be published, otherwise, the result will be arbitrary judgments that vary with each staff member’s interpretation. We should use this software if, instead of focusing on their writing and expressing their voice, we want students to distract themselves with questions like “Am I writing too fast?” and “Have I done enough edits?” This technology is poison in the well of education.

Some staff claim to be able to tell if students have used GenAI in their assessments. Whether they can tell or not isn't the issue, however. A *belief* (even a confident one) that a student used GenAI holds no evidentiary weight. There will be situations where staff are *positive* that a student has used GenAI, but this isn't evidence; it doesn't tip the scales towards misconduct, not even slightly.

This critique is not an argument against investigating potential misconduct, but a call for those investigations to rely on genuine evidence rather than hunches, subjective judgments, or practices that undermine trust and fairness.

Part 3: Threadbare analogies

People often justify the use of AI detectors by comparing them to everyday safety measures like airport metal detectors, smoke alarms or door locks. These comparisons try to make AI detection seem simple, fair and effective. But they're misleading. This article looks at these three popular analogies and explains how their supposed similarities with academic integrity processes break down when examined closely, exposing significant differences in transparency, evidence and fairness.

Airport metal detectors

There are many parallels between airport security[1] screening processes and academic integrity investigations, including the principles of natural justice and the right to have a support person present in certain circumstances. However, the comparison to AI detectors falls apart when we examine the similarities more closely.

All passengers entering an airport's secure area must pass through a metal detector (or body scanner, etc.). By walking into the screening area, you agree to be searched. Everyone is scanned, and when the metal detector raises an alert, you must remove any metal items you forgot to submit for x-ray, then walk through again. This process is repeated as needed and may involve using a handheld metal detector and/or a frisk search. You can be asked to remove clothing items to enable a proper search, and for searches of sensitive or private areas, you will be offered a private room. For these private searches, you can elect to bring a support person and must provide written consent (or not enter the secure area / get on the plane). You can continue once the security officer is satisfied you're not carrying any weapons or prohibited items. You may also be subjected to an explosives trace test. You can refuse the screening process, of course, but you can’t then enter the secure area or board a flight.

By comparison, students agree that by submitting their assessments, they consent to have their work sent to (and stored by) third-party academic integrity monitoring software, including AI detectors. This must be written in policy, and ideally, students are informed of this before each submission. Students can refuse to have their work reviewed by an AI detector, but cannot submit the assessment or receive a grade or feedback. At this point, the similarities between metal detectors and AI detectors are overtaken by the differences, which are many.

Proponents of this analogy highlight that both AI detectors and metal detectors generate false positive alerts, which is true. They also argue that, like metal detectors, one can verify

[1] Based on Australian government regulations.

the results of an AI detector quickly and non-invasively, so these false positives aren't a problem. This is false. **It is not possible to verify the results of an AI detector.** Metal detectors are also publicly tested, regulated, and their limitations are well known. They are built to international standards. AI detectors are "black boxes", proprietary, with no agreed-upon independent standards for accuracy, fairness, or error rates in educational use.

Metal detectors are built based on well-understood, transparent physics. They detect actual physical properties of metals. The relationship between the presence of metal and the alarm is direct, reliable, and independently testable. You can see and measure the metal object. AI detectors, however, use probabilistic *estimates* based on statistical similarity, not direct evidence of AI use. If you attempted to pass through security with a knife, they would show the knife at your trial, not play the metal detector's beeping sound.

Metal detectors are calibrated with clear, objective thresholds. AI detectors are not. AI detectors use arbitrary thresholds that are inconsistent across tools and unvalidated in real-world conditions. For detection purposes, the set "metal" overlaps by a minuscule amount with the set "no metal", as thresholds need to exist in practice. For example, a 0.0001g metal filing won't set off the metal detector, but it is certainly metal. However, >1g of metal may be classified as "metal" by the detector. Conversely, as argued above, there is no AI-generated text that a human could not have written; human-written text and AI-generated text are not mutually exclusive sets. No amount or type of "AI-detected text" allegedly identified moves an assessment from the "human" category to the "AI" category.

Metal detectors do not discriminate between people. They do not ignore metals on the person of people who can afford the best metallurgists, people with metal forgeries at home, or people who are more metallurgically literate. You can also easily and quickly prove you don't have metal on your person. The same cannot be said for proving you didn't use AI.

The metal detector analogy fails because it oversimplifies and misrepresents what AI detectors actually do, how they work, and how reliably they can support disciplinary claims.

Smoke detectors

Smoke detectors work by sensing physical particles or heat that indicate combustion. The link between what is measured (smoke, heat) and what is inferred (smoke, fire) is clear, testable, and grounded in science. In contrast, AI detectors try to infer text's *origin* from statistical patterns in language, a fundamentally indirect, probabilistic, and error-prone task.

When a smoke detector goes off, it is treated as an alert to investigate further, not as proof of arson or negligence. However, institutions often treat AI detector flags as *evidence of cheating* rather than as a marker for inquiry.

False positives in smoke detectors are benign, annoying, but benign. Proponents of this analogy argue the same for AI detectors, that the AI-generated flag is also benign. However, as discussed above, the output of an AI detector does not sit passively in a list of concerns. It drives suspicion, initiates investigations, and implicitly reshapes the threshold for guilt. It must be stated that this is not a flaw of the AI detector, rather, it is a flaw of the policy / procedure of the institution. Notwithstanding this, it is demonstrably true that false positives in AI detectors continue to lead *directly* to charges of academic misconduct against students at some institutions; reprehensibly, it can be the *only* evidence used to convict students of academic misconduct.

Smoke detectors are also easy to test and verify—create smoke at a given threshold and see if the detector works. Smoke detectors that cannot detect smoke at a minimum threshold cannot be sold (in Australia, at least). There are no such standards for AI detectors.

Smoke detectors' performance is also publicly regulated and standardised. Their false alarm and detection rates are well-documented, and devices are certified to meet safety standards. AI detectors have no comparable regulatory oversight. Their error rates are often unknown, claimed without supporting evidence, or misrepresented, especially when used outside controlled training data.

Door locks

Proponents of this analogy argue that although the locks on the doors of our house are relatively easy to defeat, we still use them and lock our doors. Ergo, it is argued, that simply because it's child's play to defeat AI detectors (it is), we shouldn't stop using them.

Locks on doors are simple, transparent security measures designed to prevent access. When you lock a door, you create a real, physical barrier. While not preventing access outright, locks certainly slow access, make it noisy in some scenarios, and can leave evidence of tampering. AI detectors don't *prevent* students from using AI and don't make the use of AI more difficult. Instead, AI detectors attempt, after the fact, to *guess* whether text was AI-generated.

Locks deter some intruders, but this deterrence is benign, contrary to AI detectors. As argued in the first article, the mere fact that something works does not make it right. A system that deters through fear of being (mis)judged, rather than clarity of expectation and appropriate enforcement, is not ethically defensible and shifts the burden of proof onto students without offering procedural safeguards in return. Even if AI detectors prevent student misconduct, they do so by compromising the values that should underpin assessment: fairness, transparency, and mutual respect.

Door locks have well-understood, standardised failure rates, and their effectiveness is publicly tested and measurable. The amount of force needed to kick a locked door in with a specific thickness, weight, frame, lock, etc., is *a known standard.* Consumers choose locks knowing their rated security level. AI detectors have no standards. Locks protect property and create accountability by forcing most break-ins to leave evidence. AI detection, in contrast, claims to *find* evidence based on probability alone.

These comparisons are lazy, intellectually dishonest fig leaves for unregulated, opaque, error-prone tools unfit to inform educational decision-making.

Part 4: Procedural fairness

In academic integrity investigations, the institution must show that it is more likely than not that the student has committed academic misconduct. The burden of proof rests with the institution, not with the student to disprove the allegation. It's also not the student's responsibility to prop-up a weak case against them (one based solely on the result of an AI detector, for example). The student's innocence must be the default starting position and remain until evidence—not suspicion, vibes, feelings, or "I just know"-type declarations—meets the **balance of probabilities standard**.

Students under investigation for academic misconduct have the right to silence. They are not required to speak on their behalf. Critically, enacting this right does not provide the institution with evidence against the student. It does not tip the scales against them. *If* the totality of evidence consists of the result of an AI detector, and the student under investigation chooses not to respond, the institution has **no case**.

This means that where a staff member is certain that a student has used AI and an AI detector confirms this, if the student under investigation chooses not to respond, finding that student guilty of academic misconduct is a baseless, procedurally unfair decision that does not meet the balance of probabilities standard.

On the other hand, there will be cases where other evidence is available/discoverable that could meet the balance of probabilities standard. This is why it is still critical for staff to continue to submit allegations when they suspect a breach of academic integrity. Note, however, that in many cases, there will be insufficient evidence to meet the standard. This outcome doesn't necessarily mean/indicate/declare that the student didn't use AI in an unauthorised way; they may have. But given that we operate in a system of procedural fairness, without sufficient credible, relevant, and probative *evidence*, we simply can't find against the student in such cases.

Simply telling a student that an AI detector has flagged their work as partially or entirely AI-written is problematic. Students don't know the litany of issues outlined above and likely think that a positive AI detector result constitutes evidence against them (it does not). Being leveraged with (non-existent) evidence, students may feel pressured to defend themselves but end up implicating themselves instead. If this sounds appropriate, serious reflection on your behalf regarding the principles and values of education is warranted.

Doesn't this all mean that, in many cases, students can use AI to avoid learning, and we can't take action?[2] Yes.

But this shouldn't come as a surprise. It's the underlying reason why Australia's Higher Education regulator, the Tertiary Education, Quality and Standards Agency (**TEQSA**), issued a **Request for Information** (RFI) to all Australian Higher Education providers asking them how they are responding to the risks (and opportunities) posed by GenAI. If AI detectors worked, it's arguable that TEQSA would never have issued an RFI.

Associate Professor Mark A. Bassett is Co-Director, Academic Quality, Standards, and Integrity, and Academic Lead (Artificial Intelligence) at **Charles Sturt University**. He is the author of the **S.E.C.U.R.E. GenAI Use Framework for Staff**.

Opinions expressed are the author's own.

[2] Academic integrity-related action. We can and must take action to change summative unsupervised assessments, which are the problem.