From detection to verification

Assessment towards 2030

UNIwise

From detection to verification: assessment towards 2030

Rasmus Blok, Chief Evangelist and co-founder, UNIwise · October 2026

Introduction

Almost four years after generative AI arrived in students' hands, higher education is still searching for a stable answer to a simple question: how do we know that a grade reflects what a student actually knows and can do?

Institutions have not stood still. They have tried detection software, rewritten rules on AI use, moved exams back into supervised rooms (in some cases back to pen and paper), increased the share of oral examinations, tightened lockdown and introduced online supervision. They have often done so at speed, under pressure from boards, examiners and the public, and within regulations and workload agreements that were never designed for this. Those responses were reasonable, and many of them still have a place.

But the ground keeps moving. Detection tools fall behind each new model. AI is moving from screens to glasses, earbuds, small pocket-sized assistants and other devices that are hard to see and harder to ban. And the legal and ethical expectations on institutions – around evidence, privacy, fairness and accessibility – are rising, especially in Europe.

This white paper is written for those who have to make decisions about assessment in that environment: academic leaders, heads of exams and quality, educational developers and the teaching staff who carry the change. It looks at where today's measures stand, what we expect to be true by 2030, and what a more durable approach might look like. It is not a call to tear everything down. It is an argument for knowing which direction to build in – and for taking the steps one at a time.

In brief

- No single response to AI in assessment so far – detection, process tracking, watermarks, lockdown, the exam hall, online supervision – offers a sustainable, long-term fix on its own. Each can play a part in a combined approach, but most have a shorter shelf life than expected.
- Any measure must produce evidence rather than suspicion, be fair and explainable, respect privacy and work for students who rely on assistive technology – especially in Europe, where GDPR and the AI Act make these legal requirements rather than good practice alone.

- By 2030, AI and the technology that carries it will be a fact of students' lives, often invisible, and not an exception to detect or ban.
- The products we assess – the essay, the report, the exam script – were once evidence of a thinking process and of understanding. Technology has eroded that evidence. The durable answer is to find other products and signs that demonstrate understanding, and to move from securing the finished product to verifying what lies behind it.
- There is no single answer and no single protection mechanism that covers every assessment. Institutions need to differentiate: match the level of protection to the stakes of each assessment and concentrate verification where they are highest – step by step, within the constraints they actually work under.

The glasses were still warm

In late May 2026, during a university entrance examination in Taiwan, invigilators noticed a candidate wearing a hooded top on a hot day. A first check found nothing. A second check found that the candidate had swapped to a different pair of black-rimmed glasses – and that the frame was still warm. They were AI glasses, able to read the questions and return answers.

The candidate was caught thanks to chance and attentive staff. That kind of luck will not last. Six months earlier, researchers in Hong Kong connected a pair of ordinary commercial smart glasses to a current language model and had the wearer sit a university exam. The answers on the lenses scored 92.5 per cent – in thirty minutes. Similar cases have since been reported in elsewhere, where several universities issued new bans before the summer exams.

This white paper is not about smart glasses. It is about what they reveal: the measures institutions have reached for since 2022 – trying to spot AI in submitted work once it has been handed in, or keeping AI and devices out of the exam room – are running into the same wall. The technology is becoming invisible, and measures that rely on seeing it are losing their grip.

Every fix has a half-life

Since generative AI arrived, institutions have done what institutions under pressure should do: they have acted. But each of the measures they have reached for has a half-life – a period in which it works well, or at least is the best we have, followed by decline as the technology moves on.

AI detection

The most natural question for many institutions has been: can't we simply detect AI, in the same manner we can check for plagiarism? It is an understandable ask. Similarity checking has been part of submission workflows for two decades, and AI detection is often sold as its natural extension. As we will see, the two are conceptually different – but first, what the tools can and cannot do.

AI detection tools analyse a piece of text and estimate how likely it is to have been written by a machine. The tools have improved, and some of them are better than anticipated. Under ideal conditions, the best of them can achieve a low rate of wrongly flagging human writing as AI. But ideal conditions are not the conditions of a real exam season for most institutions:

- **Detection is a moving target.** When a major new language model was released in March 2026, an independent study found that the best-performing commercial detector classified only around a third of its output as AI, while remaining near-perfect on older models. When the researchers rewrote AI text specifically to avoid detection, the same tool's detection rate fell from 96 per cent to 27 per cent. Meanwhile, human writing itself is changing as people read and write alongside AI every day, slowly moving away from the pre-2022 text detectors were trained on.
- **Test conditions are not real cohorts.** Test results describe the tools at a given moment, against the models available then. The human texts in many studies were mostly polished, published or already-approved writing – not the hesitant first-year essay or the carefully proofread work of a second-language writer.
- **The risk of bias is real.** In a widely cited Stanford study, seven detectors misclassified more than 61 per cent of essays written by non-native English speakers as AI-generated, while being much better on essays by native speakers; almost all of the non-native essays were flagged by at least one tool. Today's best tools perform better, but there is still little independent evidence of how they behave on second-language writing in real cohorts – a serious gap for European institutions where many students write in their second or third language.
- **Low is not zero.** Across an institution marking tens of thousands of submissions a year, even one per cent means hundreds of students wrongly flagged. At ten per cent – not implausible for some groups of writers – it means thousands.

Why AI use is not plagiarism

Plagiarism and AI use are often treated as two versions of the same problem. They are not.

- **Plagiarism has a source.** It is the use of someone else's existing work without acknowledgement. A similarity check can find the source, show the matching

passages side by side, and let both assessor and student see exactly what was copied and from where. The judgement still belongs to a person, but it rests on evidence that can be checked.

- **AI-generated text has none.** It is produced fresh each time, so there is nothing to match it against and nothing to show. A detector can only estimate whether a text looks statistically like machine output.
- **Plagiarism is always a breach; AI use is not.** Students may be permitted, or even required, to use AI for brainstorming, structuring, translation, feedback or language polishing, and the rules vary from course to course and assessment to assessment. The question an assessor needs answered is therefore not whether AI was involved, but whether its use was permitted – and whether the student understands and owns what they have submitted. No detector can answer that.

This is also why similarity checking still matters. Copying from other students and from published sources has not disappeared because AI arrived; in many cases it has simply been disguised by paraphrasing. Checking for it and for paraphrased or (washed) parallels remains sound, evidence-based practice. Treating AI detection as the same kind of check is where the confusion begins.

A probability, not a verdict

The bigger problem is what a detection result actually is.

- **It cannot prove anything.** A score is a probability, not a verdict. It cannot point to a source, it cannot be checked by the assessor or challenged by the student, and it cannot prove that misconduct took place. At best it is a reason to look more closely – never evidence on its own.
- **It cannot clear anyone.** Most students now use AI somewhere in their work. In a 2026 survey of UK undergraduates, 94 per cent said they used AI to support assessed work, and around a quarter said they had included AI-generated text that they had edited first. Add spellcheckers, grammar tools and translation, and few texts come back with a clean zero. When researchers in Belgium ran the leading detector over 1,163 master's theses, 45.5 per cent were flagged as containing some AI content – and a flag says nothing about whether that use was permitted. Rather than separating the innocent from the guilty, detection casts a shadow of suspicion over everyone and leaves no student able to prove that their work is their own.
- **It changes how students behave.** In a 2026 YouGov survey of more than 2,300 UK students, three in four of those using AI reported significant stress about being wrongly flagged, and international students were twice as likely as others to report a lot of stress. There are growing reports of honest students deliberately leaving errors in their work, simplifying their language, running their own writing through detectors before submitting, or recording every keystroke as proof – in effect,

learning to write worse in order to look human. Some turn to AI "humaniser" tools to protect themselves, the opposite of what detection was meant to achieve.

- **It does not stand up as evidence.** When a detection score is used against a student, the student has nothing concrete to respond to: no source, no passage, no reasoning they can examine or contest. That puts their legal position under pressure and leads to appeals, complaints and protests. Many cases cannot be upheld because there is no evidence beyond the score itself. Many institutions have therefore decided not to use AI detection at all. Some have also been concerned by detection providers seeking the right to use submitted student work to improve their own AI models – which many see as a breach of students' intellectual property.

Humans are no better at this than software. In one recent study, subject experts asked to tell human from AI text performed no better than chance. Better detection of AI involvement is not better evidence of misconduct.

Process tracking

A newer group of tools records how a text was written – keystrokes, pauses, pastes and revisions – rather than judging the finished text. It is an understandable idea: if the product can no longer be trusted, look at the process. But recording a process is not the same as assessing it, and the side effects are significant.

- **It changes the process it measures.** For many people, writing is a way of thinking: ideas are tried, abandoned and rewritten. Writers who know they are being watched tend to play it safe, experiment less, or draft elsewhere and retype the result to produce a log that looks right. The record becomes a performance rather than a window on thinking.
- **It turns the log into a KPI.** Once the log is what counts, students optimise the log – steady typing, visible revisions, well-placed pauses – rather than the quality of their thinking. It rewards looking diligent over being thoughtful.
- **It distorts collaboration.** Good group work is uneven by nature. One student drafts, another restructures, a third challenges the argument in a conversation that never appears in the document. When contribution is measured in keystrokes, the natural division of labour becomes a contest over who typed most, and groups start organising their work around the measurement rather than the task.
- **It penalises those who work differently.** Dictation software and other assistive technology can look like large pasted blocks. Neurodivergent students and second-language writers may not match the "normal" rhythm the tools expect. Students who need these tools risk having to justify their accommodations.

- **It collects more data than it needs.** Continuous capture of how a person types is behavioural data – keystroke patterns can even be used to identify individuals – and must be necessary and proportionate under GDPR.
- **It can be fooled.** One sector expert built a browser plugin in a couple of hours that types AI-generated text in a convincingly human rhythm. Tools that rely on version history or metadata from Google Docs or Word are even easier to sidestep: the history can be cleared, the text pasted into a fresh document, or the work submitted as a PDF – and then there is no process data to examine at all.

The result is another layer of surveillance that does little for learning and trades trust for compliance.

Watermarking

Watermarking takes a different approach from detection. Instead of guessing after the fact, the AI provider embeds an invisible statistical pattern in what its model produces, which can later be checked with the right key. For images, video and audio, the technique is now widespread. For text, it has only just begun. Since 2 August 2026, the EU AI Act has required providers of generative AI systems to ensure that their output is marked in a machine-readable format and detectable as AI-generated, as far as technically feasible. Systems already on the market before that date have until 2 December 2026, and tools that only assist with standard editing are exempt. In the autumn of 2026, major providers began watermarking text from their chatbots.

It is tempting to see this as the technical answer detection never was. For assessment, it is not:

- **It is easily removed from text.** Light edits may leave the pattern somewhat intact, but substantial editing or paraphrasing – for example by a second AI model – or translation into another language removes it. That is precisely what a student who wants to hide AI use would do.
- **It is unreliable on short texts.** The signal builds up over many words. In one recent study, at a one per cent false-positive rate, only around 40 per cent of watermarked 200-token passages were detected. Short answers, code and highly predictable text carry little or no signal.
- **Coverage is patchy.** Each provider uses its own scheme and keys, and some apply watermarks to text only in the EU. Open-source models running on a laptop carry none. A text without a watermark may simply come from a tool that does not use one.
- **Checking is not built for marking.** Public checkers exist for images, video and audio, but they work one file at a time. For text, detection access is so far limited to

selected researchers and organisations, and nothing plugs into a submission or marking workflow at the moment.

- **It cannot answer the real question.** Even where a watermark is found, it shows only that a particular model was involved at some point. The providers themselves say it cannot tell whether the AI wrote the text or merely edited it, how much of the document it produced, or whether its use was permitted.
- **It is not neutral.** Researchers have shown that watermark robustness varies with the statistical properties of the content, which differ across languages – another fairness question for multilingual cohorts.

The presence of a watermark says something; its absence says very little. Watermarking is valuable for transparency in society at large, and useful for subjects that work with images and media and seldom do manual edits. As a basis for academic misconduct decisions on written work, it repeats the problem of detection: a signal without an answer.

The return to the exam hall

Many institutions have moved high-stakes assessment back into supervised rooms. Some have even returned to pen and paper. At Oxford, the Faculty of Medieval and Modern Languages decided to bring back closed-book, handwritten exams for almost all papers, citing generative AI as the reason. In New Zealand, Massey University is phasing in-person exams back in by 2028, after concluding that it had no reliable way to authenticate online students' work. Others keep exams digital but move most of them them onto locked-down devices in a supervised room; a lockdown browser restricts what a student can open on their own device during an exam. In Denmark, Aarhus University reported in late 2025 that many of its degree programmes had replaced take-home exams with on-site exams – often run offline in a locked browser – with oral exams, or with a combination of the two; Public Administration, for example, replaced a three-day written exam with an oral exam. In Norway, the number of students sanctioned for suspected AI-assisted cheating at the country's ten largest institutions doubled from 95 in 2024 to 190 in 2025, and a government-appointed expert committee has recommended fewer uncontrolled exam formats and more controlled ones.

The move is understandable, but it also comes with a price:

- **Cost and capacity.** Rooms, invigilators and timetabling return as a fixed cost every exam season – and many institutions simply do not have the capacity to meet the demand. At NTNU in Norway, exam volume has grown by 30 per cent since 2022, partly because AI has made take-home exams harder to run, and more than 6,500 students will sit their autumn exams in January 2027 – in the following semester. Afternoon and weekend sessions have been added to make fuller use of the exam halls. At OsloMet, just over 1,500 students will sit exams on Saturdays in December

2026, as former take-home exams become supervised written exams. The exam season stretches – into the evening, into the weekend and into the next semester.

- **An extra step.** Paper scripts can be scanned and brought into a digital workflow, so marking, moderation and archiving can stay digital. But digitising paper adds a step of its own – logistics, quality control and the risk of pages that are missing, misordered or unreadable.
- **Handwriting.** Handwritten answers are harder and slower to read, and legibility can influence the marks given – a bias that has nothing to do with what the student knows. Writing speed also becomes an unintended part of what is assessed. And many students now rarely write by hand at all: after years of typing, a three-hour handwritten exam is itself an unfamiliar task. Returning to paper asks students to perform in a format that neither their studies nor their working lives prepare them for – a short-sighted answer to a long-term change.
- **Accessibility.** Special arrangements are harder to deliver in a hall than on a student's own, configured device. On paper, they largely disappear: students with dyslexia or other needs lose the spellchecking, text-to-speech and other assistive tools they rely on every day, and those who need them may have to be moved to separate rooms with extra staff – adding further to the pressure on capacity.
- **The student is not secured.** Paper and locked-down devices control what is on the desk and on the screen. They do not control what the student brings into the room on their body, whether legitimately, as an approved aid or assistive technology, or as an attempt at misconduct. From the outside, the two can be impossible to tell apart. The room can be secured; the student cannot.

That is where wearables come in. Smart glasses, smartwatches, earbuds, pendants, rings and pens are acquiring AI features, and they are becoming cheaper and less conspicuous every year. Alongside them, a new category of small, screenless AI devices is emerging – clipped to clothing, worn around the neck or simply kept in a pocket – that you talk to and that talk back, typically through an earbud. Several of the leading AI companies are developing devices of this kind.

The numbers are already significant. Around seven million pairs of Meta's AI glasses were reportedly sold in 2025 alone. In England, the exams regulator Ofqual reports that phones and smart devices accounted for 44 per cent of student malpractice cases in school and college exams in summer 2025. Researchers at Deakin University describe device bans as an enforcement illusion: a rule that cannot be enforced does not make an exam AI-free; it only makes it look that way. They also point to where enforcement leads – what they call bodily adjudication, with invigilators deciding whether a hearing aid, a pair of glasses, a headscarf or other head covering, or a pair of headphones is suspicious. Many students have entirely legitimate reasons for wearing them: religious practice, medical needs, or ear defenders and noise-cancelling headphones to shut out sound and

concentrate, as many neurodivergent students do. The burden falls hardest on disabled students and those whose dress is not optional.

The return of the oral exam

Alongside the exam hall, many institutions have reintroduced oral examination or increased its share. In the United States, Associated Press reports that universities such as Cornell, Penn, NYU and UC San Diego are adding oral defences to written work, with Penn's teaching centre describing "a massive shift toward in-person assessments". In Ireland, the Higher Education Authority's national framework, published in December 2025, encourages institutions to consider an institution-wide oral safeguard, in which the outcome of a spoken authorship check takes precedence over the written work – oral verification, it notes, can ensure authenticity "without recourse to unreliable detection technologies". In Denmark, upper-secondary students will have to defend major written assignments orally. And in much of continental Europe, and in the Nordic countries in particular, oral exams never went away; the question there is how to use more of them.

A conversation remains one of the strongest ways to establish understanding, and we return to it below. But scaling it has a price – one that makes it impossible to use as the only defence against AI across a whole institution:

- **Time and scale.** A written exam can seat 500 students at the same time. An oral examines them one by one, like beads on a string: each student needs a slot of their own, a room and assessors present throughout, and marking happens continuously, candidate after candidate, rather than in one batch. With twenty minutes of examination and ten minutes for the examiners to confer and set the grade, a cohort of 500 students means 250 hours of assessor time – for every examiner in the room. The cost in assessor time, scheduling and documentation rises in step with cohort size.
- **Consistency.** Different examiners, questions and dynamics can lead to different outcomes for similar performances.
- **Fairness.** Orals can reward confidence and fluency over understanding – a concern for anxious students and second-language speakers alike.
- **Not immune.** The Deakin researchers argue that interactive orals are now vulnerable to wearable AI as well, if less clearly than invigilated exams.

Used as one part of a design, the oral is powerful. Used as a blanket replacement, it is expensive and still not watertight.

Online supervision

Online supervision, or remote proctoring, moves the exam hall onto the student's own device. Instead of gathering students in one room, the institution supervises them where they are – via webcam, microphone, screen recording and automated flagging, with

recordings reviewed by people. Many institutions first adopted it during the pandemic. Since then it has become a way to keep home and distance exams credible as AI has made unsupervised formats harder to trust, and to avoid the cost and capacity limits of the exam hall. It has real strengths:

- **Supervision without the room.** Students sit the exam on their own device, wherever they are, without the cost of venues, invigilators and timetabling. It also spares students the journey: distance learners, students on exchange or placement abroad, and candidates for national or professional certifications no longer need to travel to a test centre in another city, region or country to sit an exam.
- **Accommodations as configuration.** Extra time, assistive software and other special arrangements can be set up on the student's own device rather than negotiated in a hall.
- **A documented record.** Recordings and logs replace an invigilator's recollection, so any concern can be reviewed later by a person, with the evidence in front of them.
- **Proportionate levels.** Supervision can range from light identity checks to device monitoring to full live proctoring, matched to the stakes of each exam.

But it faces the same pressures as everything else:

- **A camera sees what a camera sees.** A well-hidden earbud, a voice assistant in a pocket, a second device out of frame or a future pair of discreet glasses will challenge online supervision just as it challenges the exam hall.
- **A flag is not a finding.** Automated flags – a glance away, a voice in the room, an unusual movement – are probabilities, just like a detection score. They must be reviewed by a person, and under the EU AI Act systems that monitor and detect prohibited behaviour during tests are classed as high-risk, with requirements for human oversight from December 2027.
- **It reaches into the student's home.** Cameras, room scans and recordings are personal data captured in a private space. They must be necessary and proportionate under GDPR, with clear rules on retention, access and where the data is stored. In Norway, where the legal basis for monitoring home exams has been actively debated, the door is now opening for trials under defined conditions – a sign that the question is less whether online supervision can be used than how, and where.
- **It is not equally easy for everyone.** It assumes a reliable device, a stable connection and a quiet room. Automated analysis of faces and behaviour has been criticised for working unevenly across groups, and students whose movements or behaviour differ from the expected pattern – including many disabled students – are more likely to be flagged.

- **Being watched has a cost.** For some students, the awareness of constant observation adds anxiety that has nothing to do with what is being assessed.

For an occasional high-stakes exam, these are a reasonable price for assurance. As the preferred or only form of examination, used as the standard defence against AI, they stack up – creating new problems for students and straining the institution's relationship with them. And online supervision remains a costly form of examination: licences, technical support and the human time needed to review flagged sessions all grow with every exam that is supervised.

Taken together

Seen side by side, the measures fall into a pattern. Those that try to judge the finished work after the fact – AI detection, process tracking and watermarking – have the shortest lives. They chase a technology that changes faster than they can be updated, they can be sidestepped with little effort, and at best they produce a signal rather than evidence.

The measures that control the setting – the lockdown browser, the exam hall and online supervision – last longer, and each does something useful. A lockdown browser on a closed-book exam is sensible; supervising a high-stakes online exam can be proportionate. But all of them secure the room or the device rather than the student, and as AI moves onto the body, that gap widens. They also come at a growing cost in space, staff and time.

The oral exam stands apart. It looks directly at understanding and is likely to remain valuable – but not as a blanket replacement, because it does not scale to every student and every exam.

None of this means these measures are wrong. Each can be a sensible layer; none is a foundation. An assessment strategy built on any one of them will need rebuilding every time the technology moves. Supervision buys time; design buys durability – and it is design we turn to in the sections that follow.

The preconditions: evidence, fairness, privacy and accessibility

Before choosing any measure, European institutions also have to ask whether it is allowed – legally, and in terms of their own values. Four preconditions apply to everything that follows.

Evidence, not suspicion

A misconduct decision can affect a student's degree and career, so it must rest on evidence the student can see and respond to. A similarity report points to a source that

both assessor and student can check. An AI detection score is a probability without a source. A student cannot prove a negative, and when a score becomes the basis for an accusation, the burden of proof quietly shifts to the student. That erodes their legal position and, over time, trust in the whole system.

National bodies are reaching the same conclusion. Ireland's Higher Education Authority states that detection tools should not be treated as conclusive proof of misconduct, and Norway's expert committee has advised against using them to identify unauthorised AI use at all, citing their unreliability, the difficulty of explaining their results and the weakening of students' legal protection.

Fairness

Measures must work equally well for everyone. As described above, detection carries a documented risk of bias against second-language writers, and research on watermarking points to similar differences across languages. For European institutions with multilingual cohorts, any measure whose reliability varies with a student's language background fails this test.

Privacy and data protection

Every measure that monitors students – their writing, screen, face or room – processes personal data and must be necessary and proportionate under GDPR, and decisions with significant effects cannot rest solely on automated processing. The EU AI Act classifies AI systems used to evaluate learning outcomes, or to monitor and detect prohibited behaviour during tests, as high-risk. That is likely to cover automated proctoring flags and AI detection used in misconduct cases, as well as AI used in marking, and it brings obligations on risk management, human oversight, data quality and transparency, which apply from 2 December 2027 following this year's amendment of the Act.

Explainability matters just as much. Explainability means being able to say, in terms a person can understand, why a system produced the result it did. The AI Act requires high-risk systems to be designed so that those using them can interpret their output, and gives people affected by decisions based on such systems the right to a clear and meaningful explanation of the role the system played. GDPR already gives students the right to meaningful information about the logic behind automated processing that affects them. In practice, an institution must be able to tell a student why their work or their exam session was flagged – not just that a tool returned a score. A detector that reports "87 per cent AI" without being able to show why, or a proctoring flag based on patterns no one can explain, will struggle to meet that standard.

Institutions should also know where student data goes, and whether it is used to train anyone's models.

Accessibility

Smart glasses that caption speech, earbuds that translate and apps that read text aloud are, for some students, assistive technology. Ear defenders and noise-cancelling headphones help many students shut out distractions, and head coverings are worn for religious and medical reasons. A blanket ban on devices – or routine inspection of what students wear – risks cutting disabled students off from the tools that make assessment fair for them, singling out students for their faith, and handing invigilators the impossible job of deciding what is legitimate. Any new measure should be reviewed through an equality lens before it is introduced.

Five questions to ask before adopting any measure

The preconditions above can be turned into a simple test. Before introducing a new measure – or keeping an existing one – ask:

- **Durability.** Does it still work when the technology is invisible?
- **Evidence.** Does it produce evidence a student can see and challenge?
- **Explainability.** Can we explain its result to the student?
- **Fairness.** Is it fair across languages, backgrounds and abilities?
- **Proportionality.** Is the data it collects necessary and proportionate?

A measure that fails one of these questions is not necessarily ruled out, but it should not carry the weight of a decision on its own.

What we expect to be true in 2030

Predictions age quickly, so we offer these as expectations rather than certainties. But each follows directly from developments we can already see.

- **AI will be worn, not just used.** Glasses, earbuds, pendants and small pocket devices that you simply talk to will carry AI assistants that are hard or impossible to distinguish from ordinary accessories – or are not visible at all. Keeping AI out of a room will become a matter of trust, not control.
- **AI will act, not just write.** AI agents – tools that operate a browser or an application on a person's behalf – can already complete online quizzes, build spreadsheets and produce presentations. Any digital task that can be done can increasingly be delegated – which leaves any unsupervised assessment that ends in a digital product exposed – from online quizzes and multiple-choice tests to reports, code, spreadsheets and presentations.
- **Mixed authorship will be the norm.** Most written work will be produced with some AI involvement, in drafting, structuring, translating or polishing. The question will

not be whether AI was used, but whether the student understands and owns what they submit.

- **Employers will expect graduates who can work with AI.** Employers already rank AI and big data as the fastest-growing skills, and two-thirds plan to hire people with specific AI skills, according to the World Economic Forum. Graduates will be expected to use AI well: to judge its output, question it, combine it with their own expertise and take responsibility for the result. Assessment that only ever measures students in conditions where AI is absent will increasingly test something the labour market no longer asks for. Authentic assessment – tasks that reflect how knowledge is actually used in practice – will have to include AI, not exclude it.
- **Detection will remain a race.** Detection tools will keep improving, but the gap will widen: each new generation of models, and the growing share of work written with some help from AI, will move faster than detection can follow. The reference material is eroding too. Detectors have learned what human writing looks like from texts written before generative AI. If they keep relying on that pool, it grows ever more outdated; if they learn from newer texts, AI-generated material seeps into what they treat as human. Either way, the baseline for "human" weakens – and honest work risks looking suspicious. A score may be of interest to a curious assessor, but it will not become a sound basis for sanctions.
- **The value of a grade will depend on how it was verified.** Employers, accreditation bodies and the public will increasingly ask what a grade actually proves. Institutions that can show how they verified understanding will be in a stronger position than those that can only show what they prohibited.

The common thread is simple. By 2030, AI and the technology that carries it will be a fact of students' lives: something they carry, use and depend on, not an exception to be detected or forbidden. Assessment must be designed for that world, not against it.

AI in the assessment process

Alongside what students bring to their studies, we expect AI to change the assessment process itself. These are not new exam formats, but conditions that will shape how assessment is run – and, used well, they can strengthen evidence, fairness and transparency.

- **AI as quality assurance.** AI will increasingly support the checks that already surround assessment: acting as a second reader alongside the examiner, helping moderators and external examiners by surfacing patterns across a cohort, checking consistency between markers and against the assessment criteria, flagging outliers and selecting samples for human review. Norway's expert committee has already recommended that institutions trial AI as support for marking and grade

justification. What matters is that a human stays in the loop: AI can prepare, check and flag, but the examiner makes and owns every decision. Used that way, the result can be more consistent and better documented grades.

- **AI in feedback and assessment design.** AI will help assessors turn marking into learning: richer and more consistent feedback, concrete action points and feed-forward for the next task, clearer grade justifications, and support in drafting assessment criteria and rubrics from existing marking guidance. Done openly, this makes it easier for students to understand why they received a grade – and what to do next.
- **AI competence as a new dimension of learning outcomes.** Working with AI will become a competence in its own right, alongside the familiar cognitive levels in frameworks such as Bloom's taxonomy or SOLO. At a basic level, it means using tools appropriately; at higher levels, judging and improving AI output, deciding when not to use it, and taking responsibility for the result. The OECD will assess media and AI literacy in PISA from 2029, with students working with AI tools during the test. Learning outcomes and assessment criteria will increasingly need to describe what good use of AI looks like at each level.

None of this replaces human judgement. Used openly and within clear rules, each can help make assessment more consistent, more transparent and better evidenced.

From securing the product to verifying understanding

For most of the history of higher education, the finished product – the essay, the report, the exam script – has been evidence of something we could not see directly: a thinking process, and the understanding behind it. Producing a coherent argument required the student to have done the thinking. That is why the product was worth assessing.

Technology has eroded that evidence. A polished text can now be produced without the thinking it used to prove, and no amount of policing the product will restore the link. The product has not become worthless, but on its own it no longer demonstrates understanding.

The durable answer is to find other products and signs that do: explanations, decisions, drafts, dialogue and the ability to apply and defend what has been submitted. In practice, that means designs that allow technology where it serves learning and reflects how the subject is practised, and then verify – by other means – that the student understands what they have handed in. Several forms are already well established:

- **Oral examination and oral defence.** A dialogue that probes reasoning and follows up on answers is far harder to outsource than a finished text – especially when

combined with other evidence. A short oral defence of a written assignment can verify authorship and insight in ten or fifteen minutes. Dublin City University's guidance plans for about fifteen minutes per student including marking, and suggests that a ten-minute oral can carry roughly the same weight as a 4,000-word written submission; Griffith University in Australia runs interactive orals with cohorts of more than 700 students. At UC San Diego, around 70 per cent of 560 engineering students who took 15-minute orals said the format increased their motivation to learn. The main barriers are practical – scheduling, assessor time, documentation and moderation – and they can be planned for.

- **Staged and sequenced assessment.** Assessing a piece of work at several points – a proposal, a draft with feedback, a final submission, a reflection – makes the learning visible over time, rather than resting everything on one finished product.
- **Combined and portfolio assessment.** Combining different formats lets each part verify the other: a group project followed by an individual reflection on the student's own contribution; a lab report paired with a short supervised task applying the same method to new data; a product or prototype presented and questioned by peers; or a portfolio built across a semester, in which students select their own work and explain how it shows their learning. The whole becomes harder to outsource than any single piece.
- **AI-permitted tasks with a verification step.** Students can be allowed, or required, to use AI, and then asked to explain, defend or apply what they produced with it: why they prompted the way they did, which suggestions they rejected and why, how they checked the output for errors and invented sources, and what they added that the tool could not. The point is to assess judgement, not to pretend the tool does not exist.
- **Short checks of understanding at the point of submission.** AI can read a submission, hold it up against the assessment criteria and generate a handful of quick questions about the student's own text – for example five multiple-choice questions – which the student answers as part of the hand-in itself. The answers can indicate whether they understand what they have submitted and identify who should be invited to a follow-up conversation. The check flags; it does not decide. Timing is the point: the questions come immediately, while the work is fresh and before there is time to prepare or seek help – not days or weeks later.
- **Tasks that assess judgement.** Instead of producing work from scratch, students evaluate, compare and improve work that already exists – an AI-generated answer, a peer's draft, two competing analyses or a published study. They must show what is good, what is wrong or missing, and how it could be better. Researchers argue that this ability to judge quality – evaluative judgement – becomes more important, not less, when AI can produce text on demand.

- **Practical, performance-based and problem-based assessment.** Students demonstrate competence by doing: simulations, practical and clinical exams, lab work, studio critiques and live problem-solving. Problem-based exams give students – often in groups – an unfamiliar, real-world problem to analyse and solve within a set time, followed by a presentation or defence of their solution. What is assessed is reasoning in action, which is hard to hand over to a tool.
- **AI-led oral examination.** An emerging option is for AI to ask the questions and conduct the examination, with a human examiner supervising or acting as second marker. At New York University, 73 students took a 30-minute oral exam run by an AI voice agent; the transcripts were marked by several AI models, with a share flagged for human review, and around two-thirds of the students felt the format tested their actual understanding. It could address the oral exam's scaling problem, but it is still early: its validity against human examiners has not been established, non-native speakers found it harder, and under the EU AI Act it would be a high-risk use requiring human oversight.

What these forms have in common is that they look at content and understanding rather than polish – and that they do not depend on keeping technology out of the student's hands. Verification should also be proportionate and part of learning, not a second trial: students should know in advance how their understanding will be checked, and why.

Differentiate: not every assessment needs the same protection

If every assessment were to be redesigned at once, or secured to the highest level, institutions would run out of rooms, staff time and money long before they ran out of exams. The answer is to differentiate – and to accept that different assessments can carry different levels of risk.

Some assessments exist mainly to support learning. If a student misuses AI on a formative task, the main person harmed is the student. Others decide whether someone progresses, graduates or is licensed to practise. There, the institution's responsibility to the public is direct. It makes sense to tolerate more risk in the first case, and to concentrate verification where the consequences are greatest.

This is not a call to abandon existing formats and switch wholesale to new ones. The formats described in the previous section are building blocks, not a menu from which to pick one. They can be combined with each other and with the exams institutions already run: a written exam with a short oral follow-up for a sample of students, a supervised test alongside a portfolio, an AI-permitted project with a check of understanding at hand-in. The task is to profile each assessment by its stakes and its learning outcomes, and then put together the combination of formats and safeguards that fits – keeping

what works, and adding verification where it is needed. This is also the direction taken by Norway's expert committee on AI in higher education, which has proposed that uncontrolled formats such as take-home exams be paired with a controlled component. That component need not be oral; it could, for example, be a supervised written exam based on questions drawn from the student's own submission.

Differentiation also works at two levels. Across a programme, verification can be planned as a whole: a few deliberately placed, secure checkpoints confirm the programme's learning outcomes, so that not every course has to carry the full burden of assurance. Within a course, assessment can be spread across several forms and moments – a staged assignment, a practical task, a short oral – rather than resting everything on a single final exam. Australia's higher education regulator, TEQSA, describes this choice between programme-wide, course-by-course and hybrid approaches in its 2025 guidance on assessment reform.

The table below shows how that might look in practice.

Risk profile	Typical assessment	Approach to AI	Proportionate verification
Lower stakes	Formative tasks, practice, early coursework	Allowed and taught; use declared	Feedback dialogue, reflection, peer review
Moderate stakes	Graded coursework contributing to a module result	Allowed within clear rules	Staged submissions, short checks of understanding, oral follow-up for a sample or for flagged cases
High stakes	Final, progression, professional and licensing assessment	Defined by the learning outcome	Individual verification: oral defence, supervised components, combined formats with layered security

Several institutions have already formalised this kind of differentiation. The University of Sydney's widely discussed "two-lane" model separates secure assessments – used sparingly, and ideally at programme level, to confirm learning – from open assessments in which AI is allowed and which make up most of a student's work. The University of Bath is introducing a similar model from 2026–27. Tools such as the AI Assessment Scale, which sets out five levels from no AI to full and exploratory use, help make the expected use of AI explicit for each task. In simplified form, the scale looks like this:

Level	Use of AI	How understanding might be verified
1. No AI	The task is completed without AI, under supervised conditions.	Supervised exam, oral examination or practical test.
2. AI planning	AI may support exploring the topic, outlining and early research. The work itself is the student's own.	Planning notes or a short reflection; a check of understanding at hand-in.

Level	Use of AI	How understanding might be verified
3. AI collaboration	AI may help generate ideas, draft, give feedback and refine – but the task is designed so that AI alone cannot meet the standard.	An account of how AI output was evaluated and used; oral follow-up for a sample.
4. Full AI	Students are expected to direct AI to complete the task. What is assessed is their judgement and subject knowledge in doing so.	Defence of the choices made and evaluation of the result.
5. AI exploration	Students use AI creatively to tackle problems or develop new insights, sometimes designing the approach together with the teacher.	Presentation or defence of the approach and the findings.

Levels adapted from the AI Assessment Scale by Perkins, Furze, Roe and MacVaugh. The verification examples are our own.

The risk profiles above build on the same idea, with room for more than two lanes and for combinations within each.

Sampling matters here. Not every student needs an oral defence. A combination of random selection and follow-up where there is reason for doubt can give strong assurance across a cohort at a fraction of the cost. At fifteen minutes per student, following up a 20 per cent sample of a 200-student cohort takes around ten hours of assessor time – far less than examining everyone orally. Students who know they may be asked to explain their work have every reason to understand it.

Differentiation also makes the exam hall, lockdown and proctoring easier to justify. Used where the stakes demand them and the learning outcomes call for them, they are proportionate. Used everywhere, by default, they become expensive, intrusive and – as wearables spread – increasingly symbolic.

One step at a time

None of this happens in one exam season, and it would be unfair to suggest otherwise. Assessment is bound up in course descriptions, exam regulations, national legislation, accreditation requirements, external examiners, workload agreements and timetables. Changing it involves many people, and much of it is negotiated a year or more in advance. Institutions facing real dilemmas today deserve a path they can actually follow.

The pace will differ from institution to institution. What matters is the order: each step builds on the one before.

First: get an overview and set clear rules

- Map your assessments by stakes and identify where the integrity risk is highest.
- Make the rules on AI use clear for each assessment, and explain them – and the reasons behind them – to students.
- Stop treating AI detection scores as evidence – at most, a prompt for a conversation.

- Check that lockdown and supervision are proportionate, and that device rules meet your accessibility obligations.

Next: try new formats and build the conditions for them

- Pilot oral defences, staged designs and short checks of understanding in a few programmes, and learn from them.
- Build oral assessment into workload models and timetables, so it does not depend on goodwill.
- Adjust exam regulations so that new formats, sampling and follow-up conversations have a clear basis.
- Invest in staff development, both in using AI well and in designing assessment that verifies understanding.
- Build students' AI literacy, so that they learn to use AI critically and to account for how they have used it.

Then: make it part of the design

- Scale what works from the pilots, and combine formats according to each assessment's risk profile.
- Plan verification across whole programmes rather than module by module, and vary assessment forms within each course rather than relying on a single final exam.
- Make verification of understanding a design principle rather than an emergency measure.

What we believe – and what we build

We are often asked whether WISEflow and WISEflow Originality can detect AI, block it or keep it out. We understand why. But the more useful questions are how institutions can know, with confidence and fairness, what their students understand – and, where misconduct is suspected, how it can be documented with evidence that holds. Our view rests on five principles:

- **Verify understanding, not just products.** The finished product is no longer enough on its own. Assessment must look at what students know and can do.
- **Evidence before suspicion.** Decisions about students must rest on evidence they can see and challenge – not on probabilities without a source.
- **Differentiate and be proportionate.** Different assessments carry different risks. Security should be strongest where the stakes are highest, and lighter where learning is the purpose.

- **Design for the technology students actually have.** AI is part of students' lives and future work. Assessment should be designed for that reality, including for students who depend on assistive technology.
- **European by default.** Student data belongs to the institution and the student. It should stay in Europe and must never be used to train anyone's models without consent.

This is how those principles shape what we build at UNIwise:

- **WISEflow Originality** checks similarity against identifiable sources, with matches marked sentence by sentence so assessors and students can verify them – including paraphrased and AI-reworded copying. We have deliberately chosen not to offer AI detection, because it cannot provide that kind of evidence.
- **Layered, configurable security.** A lockdown browser, device monitoring and remote proctoring can be set exam by exam, so institutions can apply the level of supervision each assessment actually needs – no more, no less.
- **Oral exams as a process.** WISEflow supports the full process of oral examination – planning, scheduling, assessor coordination, marking and documentation – so that oral verification can work at scale.
- **Sequences and combinations.** With Sequence of flows and combined and portfolio formats, institutions can assess work at several points and combine written, oral and supervised components.
- **AI-assisted feedback, with the assessor in control.** AI can help assessors give richer, more consistent feedback and grade justification, while the human assessor remains responsible for every decision.
- **AI-assisted rubrics and translation.** Coming in WISEflow's December 2026 release: AI can help draft rubrics from existing marking guidance and translate assessment material, with the assessor reviewing and approving every result.
- **Accommodations as configuration.** Special arrangements and assistive technology are handled as part of everyday exam set-up, not as exceptions.
- **AI-supported paper submission.** Where institutions choose handwritten exams, students write a code in the top right-hand corner of each sheet. After scanning, AI recognises the codes and sorts and distributes every page to the right submission in WISEflow – so the process stays digital for marking, moderation and archiving, even when the exam is written on paper.
- **Data in Europe.** Student data is hosted in the EU and is not used to train third-party models.

We are also exploring how AI can support more of the assessment process, always with the human in control – for example in creating questions and items. And we are

exploring new ways to verify understanding at the point of submission, building on the AI-assisted feedback we already offer, so that institutions can keep open, take-home formats with confidence.

Conclusion

The exam hall, the lockdown browser, the proctoring session and even the detection report all have their place. But each one is a defence against yesterday's version of the problem, and each one will need rebuilding as the technology moves. By 2030, AI will be worn, embedded and largely invisible. Assessment that depends on keeping it out will not hold.

What will hold is assessment designed to verify understanding: oral and written, staged and combined, open where learning benefits and secured where the stakes demand it. Getting there takes time, and institutions should take it in steps they can manage. The important thing is to know which direction those steps are going.

From detection to verification: that is where we believe assessment is heading – and where we intend to help our customers get to. If you are working through these questions at your institution, we would be glad to share what we see across the sector and how UNIwise can help – from mapping assessments by risk and planning pilots to setting up WISEflow for each step of the way.

Sources and further reading

Acompli (2026) AI Omnibus enters into force: high-risk deadlines move to December 2027 and August 2028, July 2026. https://acompli.ie/news/ai-omnibus-enters-into-force-july-2026/

Aformeziem, B. (2026) Catching the wrong students: AI detection, international students and the fairness crisis in UK universities. HEPI, 20 July 2026. https://www.hepi.ac.uk/2026/07/20/catching-the-wrong-students-ai-detection-international-students-and-the-fairness-crisis-in-uk-universities/

All Things Nordic (2026) Denmark orders schools to fight AI cheating with oral exams and screen monitoring, 9 August 2026. https://all-things-nordic.com/2026/08/09/denmark-orders-schools-to-fight-ai-cheating-with-oral-exams-and-screen-monitoring/

Associated Press (2026) Perfect homework, blank stares: why colleges are turning to oral exams to combat AI. Via WTOP, April 2026. https://wtop.com/education/2026/04/perfect-homework-blank-stares-why-colleges-are-turning-to-oral-exams-to-combat-ai/

Bearman, M., Tai, J., Dawson, P., Boud, D. and Ajjawi, R. (2024) Developing evaluative judgement for a time of generative artificial intelligence. Assessment & Evaluation in Higher Education. https://doi.org/10.1080/02602938.2024.2335321

BleepingComputer (2026) How Anthropic plans to watermark Claude's AI-generated text, 14 August 2026. https://bleepingcomputer.com/news/artificial-intelligence/how-anthropic-plans-to-watermark-claudes-ai-generated-text

Cheng, Lin, Reedy et al. (2025) Ability of AI detection tools and humans to accurately identify different forms of AI-generated written content. Advances in Simulation. https://pmc.ncbi.nlm.nih.gov/articles/PMC12752165/

China Daily (2026) More universities ban smart devices in exam halls, 15 July 2026. https://studychina.chinadaily.com.cn/s/202607/15/WS6a574366498e36855033e33f/more-universities-ban-smart-devices-in-exam-halls.html

Cloud Security Alliance (2026) OpenAI textGrain: watermarking text under the EU AI Act. Research note, October 2026. https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-text-watermarking-20261009-csa-s/

Corbin, T., Sharpe, S. and Dawson, P. (2026) On AI glasses and wearable AI in assessment. Assessment & Evaluation in Higher Education. https://doi.org/10.1080/02602938.2026.2661367

Davidson, J. (2026) OIA documents reveal why Massey brought back in-person exams – AI. Massive Magazine, 27 July 2026. https://www.massivemagazine.org.nz/articles/calx1oi0kpy89j84fhiszq1ue97umg-k7axs

Dublin City University (2026) Interactive oral assessment. Teaching Enhancement Unit. https://www.dcu.ie/teu/interactive-oral-assessment

Furze, L., Perkins, M., Roe, J. and MacVaugh, J. (2024) Updating the AI Assessment Scale, 28 August 2024. https://leonfurze.com/updating-the-ai-assessment-scale

Gibson Dunn (2026) EU AI Act Omnibus agreement – postponed high-risk deadlines and other key changes, May 2026. https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/

Ipeirotis, P. and Rizakos, K. (2026) Scalable and personalized oral assessments using voice AI. arXiv:2603.18221. https://arxiv.org/abs/2603.18221

Higher Education Authority (2025) Generative AI in Higher Education Teaching & Learning: Policy Framework, 22 December 2025. https://hea.ie/2025/12/22/hea-publishes-national-policy-framework-on-generative-ai-in-teaching-and-learning/

Journalen (2026) Studenter ved OsloMet må ha eksamen på lørdager, 29 May 2026. https://journalen.oslomet.no/nb/innenriks/2026/05/studenter-ved-oslomet-ma-ha-eksamen-pa-lordager

Khrono (2025) KI-utvalg vil teste bruk av kunstig intelligens i eksamenssensur, 16 December 2025. https://www.khrono.no/ki-utvalg-vil-teste-bruk-av-kunstig-intelligens-i-eksamenssensur/1023175

Khrono (2026) Krever at NTNU ordner opp i eksamenskaoset, 15 October 2026. https://www.khrono.no/krever-at-ntnu-ordner-opp-i-eksamenskaoset/1092233

Kingkade, T. (2026) To avoid accusations of AI cheating, college students are turning to AI. NBC News, 28 January 2026. https://nbcnews.com/tech/internet/college-students-ai-cheating-detectors-humanizers-rcna253878

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E. and Zou, J. (2023) GPT detectors are biased against non-native English writers. Patterns. Summary by Stanford HAI. https://hai.stanford.edu/news/ai-detectors-biased-against-non-native-english-writers

Littler-Jennings, P. and Robson, N. (2025) Exam conditions to change for languages students mid-course, AI blamed. Cherwell, 28 May 2025. https://cherwell.org/?p=188529

Malone, E. (2025) Students suspected of AI misuse to face in-person interviews under new recommendations. The Irish Times, 31 December 2025. https://www.irishtimes.com/ireland/education/2025/12/31/students-suspected-of-ai-misuse-to-face-in-person-interviews-under-new-recommendations/

Media & Learning (2025) PISA 2029 to assess students' media and AI literacy. https://media-and-learning.eu/subject/artificial-intelligence/pisa-2029-to-assess-students-media-and-ai-literacy/

Nemecek, A., Chaudhary, V. and Ayday, E. (2026) Watermarks without verification: AI text watermarking after the EU AI Act. arXiv:2609.09604. https://arxiv.org/abs/2609.09604

Pedersen, M. (2026) Norway wants students to learn AI responsibly, but not use it to cheat. NordiskPost, 19 January 2026. https://www.nordiskpost.com/2026/01/19/norway-ai-cheating-students-doubled-in-2025/

Petersen, L. B. (2025) Artificial intelligence forces several degree programmes at Aarhus University to change exams. Omnibus, 1 December 2025. https://omnibus.au.dk/en/archive/show/artikel/kunstig-intelligens-tvinger-flere-uddannelser-paa-aarhus-universitet-til-at-aendre-eksamener

PublicTechnology (2026) Exams watchdog warns over high-tech cheating (Ofqual malpractice figures, summer 2025), 5 June 2026. https://www.publictechnology.net/2026/06/05/education-and-skills/exams-watchdog-warns-over-high-tech-cheating/

Ren, K., Raghavan, M. and Garg, N. (2026) Hitting a moving target: test-time adaptation for AI text detection under continual distribution shift. arXiv:2606.25152. https://arxiv.org/abs/2606.25152

Ross, J. (2026) 'Checkmate' looms for traditional tests amid AI glasses epidemic. Times Higher Education, 5 September 2026. https://www.timeshighereducation.com/node/744796

Ross, J. (2026) 'No defence' against wearable AI in exams, researchers warn. Times Higher Education, 30 April 2026. https://www.timeshighereducation.com/node/742661

Rowsell, J. (2026) Are universities returning to in-person exams to combat AI cheating? Times Higher Education, 29 June 2026. https://www.timeshighereducation.com/node/743640

SBS (2026) National Taiwan University entrance exam: AI glasses case, June 2026. https://sbsstar.net/article/N1008607806/why-a-hoodie-in-this-weather-why-the-glasses-were-hot-top-university-in-uproar

Stephenson, R. and Armstrong, C. (2026) Student Generative AI Survey 2026. HEPI Report 199, 12 March 2026. https://www.hepi.ac.uk/reports/student-generative-ai-survey-2026/

TEQSA (2025) Enacting assessment reform in a time of artificial intelligence. Tertiary Education Quality and Standards Agency, September 2025. https://www.teqsa.gov.au/sites/default/files/2026-06/enacting-assessment-reform-in-a-time-of-artificial-intelligence-web.pdf

UC San Diego Today (2023) Oral exams improve engineering student performance and motivation, 23 February 2023. https://today.ucsd.edu/story/oral-exams-improve-engineering-student-performance-motivation

University of Sydney (2023) Embracing the future of assessment at the University of Sydney. Teaching@Sydney, December 2023. https://educational-innovation.sydney.edu.au/teaching@sydney/?p=19768

Utvalget om kunstig intelligens i høyere utdanning (2025) Foreløpige vurderinger: notat fra utvalget om kunstig intelligens i høyere utdanning, 16 December 2025. https://files.nettsteder.regjeringen.no/wpuploads01/sites/589/2025/12/Forelopige-vurderinger-notat-fra-utvalget-om-KI-i-hoyere-utdanning.pdf

Van Vlasselaer, M., Van Droogenbroeck, F. and Spruyt, B. (2026) Who wrote this? Evaluating the reliability of AI detection tools in higher education. International Journal for Educational Integrity 22, 16. https://doi.org/10.1007/s40979-026-00226-w

Wagner, M. G. (2026) Students performing human. The Augmented Educator, 8 May 2026. https://www.theaugmentededucator.com/p/students-performing-human

Webb, M. (2025) Opinion: writing under watch – reflections on creativity, surveillance and assessment. Jisc National Centre for AI, 28 March 2025. https://nationalcentreforai.jiscinvolve.org/wp/?p=5103

Webb, M. (2026) AI detection and assessment: an update for 2026. Jisc National Centre for AI, 15 October 2026. https://nationalcentreforai.jiscinvolve.org/wp/2026/10/15/ai-detection-and-assessment-an-update-for-2026/

Williams, T. (2026) Fear of being flagged by AI detectors drives stress among students (YouGov/Studiosity survey of 2,373 UK students). Times Higher Education, 23 February 2026. https://www.timeshighereducation.com/news/fear-being-flagged-ai-detectors-drives-stress-among-students

World Economic Forum (2025) The Future of Jobs Report 2025. https://www.weforum.org/publications/the-future-of-jobs-report-2025/digest/

UNIwise (2025) Beyond the buzzword of authentic assessment, 25 November 2025. https://uniwise.eu/resources/blog/beyond-the-buzzword-of-authentic-assessment

UNIwise (2026) An oral exam is not a question type. It is a process. https://uniwise.eu/resources/blog/an-oral-exam-is-not-a-question-type-it-is-a-process

UNIwise (2026) Forward rather than back – why the exam hall is not the only answer to AI. https://uniwise.eu/resources/blog/forward-rather-than-back

UNIwise (2026) The closed-book exam comeback: why securing exams against AI needs layers, not a single setting. https://uniwise.eu/resources/blog/the-closed-book-exam-comeback-why-securing-exams-against-ai-needs-layers-not-a-single-setting

UNIwise (2026) The EU AI Act and assessment: December 2027 is not a snooze button. https://uniwise.eu/resources/blog/the-eu-ai-act-and-assessment-december-2027-is-not-a-snooze-button

UNIwise (2026) There is no quick fix for AI in assessment – and AI detection isn't one. https://uniwise.eu/resources/blog/there-is-no-quick-fix-for-ai-in-assessment

UNIwise (2026) Training AI on students' work: what European institutions need to consider. https://uniwise.eu/resources/blog/training-ai-on-students-work