Policy Note 71 | May 2026

# *What UK university AI policies actually do:* A study of 96 institutions

Sam Illingworth, Edinburgh Napier University

Universities claim to help students use artificial intelligence. I tested that claim across 96 institutions. The answer is uncomfortable.

Most UK university AI policies use the language of education while operating as compliance instruments. They promise critical thinking but deliver audit trails. They name support yet deliver surveillance. The gap between what these policies say and what they structurally do is the central finding of this study, and it matters for everyone in the education sector who believes AI policy should serve students.

This study examines what university AI policies actually do when you read them closely. Where I identify structural failings, the purpose is to identify patterns that the sector can learn from. All 96 policies as they existed at the time of the research are available in an open data repository for anyone who wishes to examine the evidence directly.[1] Where institutions are named in this report, it is either because their policies offer examples of effective practice worth learning from or because the specific language of a policy illustrates a structural pattern that would be less clear in the abstract. The aim throughout is to help the sector improve.

The computational analysis of vocabulary tells one story. The qualitative reading of the same documents tells another. That discrepancy, between the words institutions choose and the work those words perform, is the heart of what follows.

## Method

I started with 163 UK higher education institutions that hold degree-awarding powers. I tried to find their AI policies, scraping them programmatically on one day: 16 February 2026. The scraping tested what a student, prospective applicant, parent, regulator or member of the public could find through a search engine without logging in.

The technical work was done using Claude Code, an AI coding agent made by Anthropic. I designed the study, chose the research questions, built the analytical framework, selected the keyword vocabularies, decided which policies to read in full and made every interpretive judgement. Claude Code wrote the scraping scripts, ran the keyword counts, generated the spreadsheet data files and built the online data repository. It did in hours what would have taken weeks by hand.

Of 163 institutions, 96 had publicly accessible AI policies. The other 67 either had no discoverable policy or one locked behind authentication walls (requiring a university login to access), broken URLs or content that loads dynamically (meaning the policy text is assembled by the browser after the page loads, making it invisible to search engines and automated tools). That means 41% have no public AI policy that fitted within the scope of this study.

For the 96 policies I could access, I did two things.

First, I counted. I built two keyword vocabularies: 34 terms associated with detection and surveillance (such as 'plagiarism', 'misconduct', 'penalty', 'Turnitin' and 'zero tolerance') and 43 terms associated with education and support (such as 'critical thinking', 'literacy', 'agency', 'assessment design' and 'student voice'). The vocabularies were constructed inductively from a pilot reading of 15 policies. Each term was included only where it could be reliably classified as serving one framing function. Boundary cases (terms like 'integrity' or 'responsibility' that function differently depending on their context) were resolved by testing them against the pilot sample: if a term appeared predominantly in enforcement contexts across the pilot, it was classified as detection-oriented and *vice versa*. The full vocabularies with rationale for each inclusion decision are in the data repository.

For each policy, I calculated a 'framing ratio': the proportion of matched keywords belonging to the detection vocabulary. A score of zero meant exclusively education-oriented language. A score of one meant exclusively detection-oriented language.

Secondly, I read. I selected 19 policies for qualitative coding, stratified across all four institution types (see the next section), which were deliberately chosen to include the full range of framing ratios: high, low and mid-range. The aim was diversity of approach, not statistical representativeness. I excluded my own institution, Edinburgh Napier University, from the qualitative sample to avoid a conflict of interest; its policy remains in the computational analysis.

The coding framework I used has three layers, working from the abstract to the operational.

Four principles provide the theoretical orientation: student-centredness; trust; relevance; and agency. These principles were developed with Rachel Forsyth and are set out in full in our book *GenAI in Higher Education*.[2] They inform the analysis but are not applied directly to the policy text.

Six thematic categories developed for this study organise the analysis at the level of the document as a whole: student-centredness; trust; relevance; agency; deficit-model indicators; and structural / institutional framing. The first four categories carry the four principles into the operational framework. The last two emerged from the pilot reading as patterns the four principles alone did not capture.

Thirty-one codes sit within the six categories and do the tagging work. Each code captures a specific feature a policy might or might not contain (for example, an explicit trust statement, a record of student involvement in policy development and whether the policy distinguishes between AI use and academic misconduct). The full set of codes and their definitions is in the open data repository.

I applied the codes paragraph by paragraph across the 19 policies, reading every policy in full.

One limitation of scraping should be noted. Programmatic scraping captures the full HTML source of a webpage, which is not always identical to what a reader sees on screen. In at least one case, institutional text present in the HTML (an acknowledgement of student involvement in policy development) was not visibly rendered on the page. The text existed in the source code but was invisible to a student reading the guidance. This study uses scraped text. Readers checking quotes against live webpages may find minor discrepancies where content is present in HTML but hidden by a page's design. This is itself an instructive finding: even the acknowledgement of student voice can be structurally invisible.

## What the numbers show

The computational analysis classified 83 of 96 policies (86%) as education-dominant. Nine (9%) fell into a balanced category between the two. Only four policies (4%) were classified as detection-dominant. The mean framing ratio across the documents was 0.215, well inside the education range.

The most frequent education keywords were 'learning' (690 occurrences), 'guidance' (555), 'support' (484) and 'skills' (289). The most frequent detection keywords were 'misconduct' (570), 'academic integrity' (349), 'plagiarism' (131) and 'penalty' (124).

Institutions were split into four categories: Russell Group members, other pre-1992 universities (that is, universities founded before 1992 that are not Russell Group members), post-1992 universities and specialist institutions. These categories carry real weight in UK higher education. They shape funding, reputation, student demographics and institutional culture. They also, as this study shows, shape how institutions talk to students about AI.

Pre-1992 institutions outside the Russell Group had the lowest mean ratio (0.179). Every single pre-1992 policy looked at was classified as education-dominant. Russell Group institutions averaged 0.207, post-1992 institutions 0.223 and specialist institutions 0.312 (though that last figure is drawn from only six policies with a range from 0.022 to 0.791, making the mean less reliable).

The headline looks reassuring. So does the accessibility picture, on the surface: Russell Group institutions had the highest rate of publicly accessible AI policies at 75%, followed by post-1992 institutions at 69%, other pre-1992 institutions at 68% and specialist institutions at just 18%. The specialist sector, comprising conservatoires, arts universities and smaller focused providers, is nearly absent from the public AI policy conversation.

Taken at face value, these results paint a reassuring picture. The sector appears overwhelmingly committed to education and support. Only a handful of institutions have policies dominated by punishment language.

However, there is a problem with taking these results at face value.

## What the words actually do

The qualitative coding of 19 policies reveals a systematic gap between what policies say and what they do. The framing ratio measures vocabulary. It does not measure function. A policy can use the word 'learning' while meaning 'compliance'. It can discuss 'skills' while building an audit trail.

As a result, human reading is essential. The computational analysis is useful as a mapping tool, revealing what the sector's vocabulary looks like at scale, and it identifies outliers worth investigating. The most important findings in this study come entirely from reading the policies closely and asking what the language is doing.

Take the University of Southampton.[3] Its framing ratio is 0.081, placing it firmly in the education category. Southampton's opening aspires to produce 'critically digitally literate' graduates, but read beyond that and you find a binary list of acceptable and unacceptable uses anchored by an early threat: 'We will take disciplinary action that may result in penalties on your marks.' The opening and the body of the policy read as if written by different logics. The critical-literacy aspiration does no work beyond the first paragraph. The document says 'learning' but means 'compliance'.

Loughborough University is the most revealing case.[4] Its framing ratio is 0.364, computationally education-dominant. Language like 'demonstrate honesty', 'take responsibility' and 'good scholarship'

pushes the education score up. But the document's operational function is an audit trail. Its four-step structure (save outputs, write acknowledgement, describe use and reference AI) has nothing to do with learning. It is an evidence-retention protocol for potential misconduct investigations. The vocabulary of learning is pressed into the service of surveillance.

The University of Glasgow's opening statement is genuinely well-crafted: 'Rather than seek to prohibit your use of these tools, we want to support you in learning how to use them effectively, ethically, critically, and transparently.'[5]

The promise is not entirely kept. The body of the statement is a long list of permitted and prohibited behaviours, anchored in academic integrity. The policy also warns that AI research tools 'will in practice break academic integrity rules', contextualising this within a broader statement about the limitations of current AI tools and the risks of uncritical reliance on them. That framing positions the warning as practical advice, but its effect is deterrence: it tells students that using AI for research is likely to result in an integrity violation, regardless of intent.

Other policies in the sample absorb AI into existing penalty matrices alongside ghost-writing services and essay mills, with penalties extending in some cases to post-graduation revocation of awards. These approaches treat AI as a threat to be contained rather than a technology to be understood.

Of the 19 policies I coded qualitatively, the computational classification was misleading for 12: the eight coded as detection-oriented (where education-dominant scores masked compliance instruments) and the four coded as performatively educational (where education vocabulary functioned as a veneer over compliance work). For two further policies, the education-dominant classification was technically correct but obscured significant detection-oriented elements that a reader would encounter as the dominant experience of the document. The misclassification followed a consistent pattern: policies that used education vocabulary for non-educational purposes were systematically classified as education-dominant. I found three mechanisms through which this occurred.

**Performative education framing:** Policies adopt the language of learning, support and development while structurally functioning as compliance instruments. This operates at every level: individual sentences ('assessment is an important part of learning and students who do not complete assessments appropriately risk not only wasting their own time at university')[6], document structure (aspirational openings followed by regulatory bodies) and institutional location (AI guidance embedded within academic misconduct frameworks). The performative gap may well be unintentional. Institutions may genuinely wish to support student learning. It is a structural consequence of producing educational guidance within regulatory architectures. The response to AI has been absorbed into existing governance structures, and many of those structures were already organised around detection and punishment. AI did not create new deficit architectures; it inherited them.

**Conditional trust:** No policy in my sample of 19 explicitly states 'we trust students'. Some embody trust structurally, but the dominant model across the documents is conditional: students are trusted to use AI appropriately, provided they declare their use, retain evidence of their processes, acknowledge AI contributions and submit to verification if requested. A student saving their drafts because a policy tells them to demonstrate the process they undertook is building a defence file.

The University of Liverpool makes the logic explicit: 'A student that refuses to declare how they have used the technology … may be attempting to hide the fact that the work is not their own.'[6] Non-declaration is read as concealment. Silence is evidence of guilt. Students are required to be

transparent about their AI use. Institutions are not required to be transparent about their detection practices.

This matters because if a policy is going to ask students to make good decisions about AI, it must begin by treating them as capable of doing so. A policy that starts from the assumption that students will cheat unless watched, and structures its guidance accordingly, produces exactly the adversarial dynamic it claims to prevent. Students who feel surveilled are less likely to engage honestly with AI and more likely to conceal its use. Research on academic integrity consistently shows trust-based approaches produce better outcomes than surveillance-based ones. Research synthesising a decade of work on honour codes shows institutions with trust-based systems have significantly lower rates of self-reported cheating than those relying on surveillance and deterrence.[7] A large-scale Australian study found that students who reported dissatisfaction with their teaching and learning environment were significantly more likely to engage in contract cheating.[8] Explicitly stating 'we trust students' would be a structural commitment that shapes every subsequent paragraph.

**Structural location:** Where a policy sits within an institution's website is a stronger predictor of its framing than the language it uses. Of the seven policies in my qualitative sample located primarily within academic misconduct frameworks, six were coded as detection-oriented with strong deficit indicators. The exception was Durham University, which explicitly acknowledges its location within a misconduct policy and then works to exceed that frame.[9] A policy placed in a misconduct framework inherits that framework's assumptions about students, regardless of the vocabulary it employs. Conversely, policies located within teaching and learning or study skills architectures, as at – for example – the University of Stirling and Canterbury Christ Church University, were more likely to be genuinely educational.[10] The institutional decision about where to house AI policy is more consequential than the language used within it. This is not surprising: a document's context shapes how it is read and what it is understood to authorise. One might reasonably ask whether the causality runs the other way, with institutions selecting detection-oriented vocabulary first and then locating the policy accordingly. In practice, the misconduct frameworks predated the AI guidance in every case I examined. The AI policy was written into an existing regulatory architecture rather than the architecture being chosen to fit the policy's tone. The vocabulary inherits the location, not the reverse.

These three mechanisms share a common logic. The deficit model (a framing that defines students by what they are assumed to lack) does not require punitive language to reproduce itself. It requires only the assumption that students cannot be trusted to think. That assumption can be carried by aspirational vocabulary, by conditional frameworks and by the decision about which webpage to put the policy on. The surface can say 'education' while the structure says 'control'.

The University of East Anglia provides the clearest example of how these mechanisms interact.[11] Its policy uses a 'green light / red light' framework for both staff and students. But the register shifts depending on who is being addressed. The student section is directive and rule-bound. The staff guidance, written by senior academics, is rich in pedagogical reasoning, assessment redesign strategies and disciplinary examples. Both are publicly accessible in the same document. The content is not hidden. But the difference in depth reveals whom the institution designed the policy for. Staff are addressed as professionals trusted with judgement. Students are given instructions. That framing is a choice. It echoes a pattern HEPI identified in its work on the impact of menstruation in higher education, where institutional policies were aimed primarily at staff.[12]

The revised picture shows it is no longer right to claim that 86% of institutional AI policies are educational.

Of the 19 policies I coded, five were genuinely educational, four were performatively educational, two were balanced and eight were detection-oriented. This sample was designed for diversity, not representativeness, so I will not project specific percentages onto the full higher education sector. But the qualitative findings suggest the 86% figure overstates the sector's educational orientation. The HEPI / Kortext *Student Generative AI Survey 2026* corroborates this: 68% of students say AI skills are essential to thrive, but only 48% feel their teaching staff are helping them develop those skills. Students were evenly split on whether their institution encourages AI use (37% agree, 36% disagree), while 65% said assessment has changed significantly in response to AI. Qualitative responses revealed students felt anxious about being falsely accused of cheating for writing well. That survey measures student perception rather than policy content, so it is corroboration rather than confirmation. But students can see what keyword analysis cannot.[13]

## The universities getting it right

Four policies demonstrate what AI guidance looks like when it starts from a different assumption: those from Durham University; the University of Stirling; Canterbury Christ Church University; and Arts University Plymouth.

Durham University's policy is over 9,000 words long. Some will reasonably ask whether any student will read a policy of that length. The answer is probably not, cover to cover. But Durham uses that length to think out loud. Its 'Background ideas' sections constitutes a genuine critical literacy resource, addressing environmental impact, labour exploitation, copyright and algorithmic bias. It includes a 17-row guidance table. Most remarkably, it distinguishes explicitly between academic misconduct and 'unwise or unethical' behaviour, and refuses to use the misconduct process to police the latter. It tells students:

> ***Judgments about this will differ, but given the significant ethical concerns surrounding generative AI your teachers should respect your decisions about this. In particular, your teachers should not demand that you use AI tools, or create an environment where a high volume of AI usage is unavoidable …***

The length serves a purpose. A short policy is necessarily a list of rules. A longer policy can explain why the rules exist, what the institution's reasoning is and where uncertainty remains. What matters is whether the policy models the kind of critical thinking it asks students to develop. Whether every student reads every word is a secondary question.

Canterbury Christ Church University explicitly names its approach as 'educative, rather than punitive.' The document addresses sustainability, ethics and equity of access. It asks students to 'consider whether you really need to use a GenAI tool for your task' before listing permitted and prohibited uses.

Arts University Plymouth encourages its students to:

> ***take a proactive approach to understanding generative AI and how it could augment or disrupt their own creative processes.***[14]

The University of Stirling's AI in Assessment Scale provides a five-level framework from prohibition to creative integration, giving both staff and students a shared language for decisions.

None of these policies is perfect. Durham University's is very long and located within a misconduct framework (which, as noted, it explicitly works to exceed). Canterbury Christ Church University still prohibits AI use in assessed work. Stirling is pedagogically sound but sometimes thin in detail. But they demonstrate that it is possible to produce AI policy that does not treat students as a problem to be solved.

What these four share is student agency, institutional honesty, critical literacy over tool proficiency and student voice in development. They span all four institution types: one Russell Group, one pre-1992, one post-1992 and one specialist. How an institution frames AI is a choice, and it can be made differently.

## The accessibility gap

For 39 institutions, the URLs I identified did not return AI policy content. In some cases the pages had been moved or removed; in others, the pages existed but had been replaced with content unrelated to AI. For a further 28 institutions, no public AI policy could be discovered through systematic web search. Taken together, the pattern suggests that AI policies are often treated as ephemeral documents rather than stable institutional commitments, and that institutions vary widely in how findable they make their guidance. Students searching for their own university's AI policy via Google may not find it.

Some of these institutions may have policies accessible to enrolled students through authenticated portals (intranets that require a university login). If a policy is intended solely for current students, an intranet may be a reasonable location. But this framing underestimates who else needs to be able to find a university's AI policy, such as prospective students and their parents making choices about where to study. In a sector where AI is reshaping how students learn and how they are assessed, a university's AI policy is relevant information for applicant decision-making. If it exists only behind a login wall, it is invisible at the point when it could be most useful.

Poor accessibility also hampers other institutions looking for models, examples and sector norms. The reason universities are each finding their own way (a point I return to below) is partly that they cannot easily see what other institutions are doing. A publicly accessible policy contributes to a shared sector conversation. A hidden one does not.

The Office for Students, QAA and professional bodies all have an interest in how institutions are responding to AI. If a policy is not publicly accessible, external scrutiny depends on formal review processes rather than routine visibility. Given the pace at which AI is changing teaching and assessment, that lag matters.

This study can only analyse publicly accessible policies. The 67 institutions where policies are inaccessible are absent from the analysis. If their policies are more enforcement-oriented (which is plausible, given that less visible documents face less external scrutiny), then the picture this study paints is more optimistic than reality. Employers and professional bodies who need to understand what AI competencies graduates have been expected to develop are also left to guess.

## The sector coordination gap

One of the most striking findings is how much variation exists across a sector that has no shortage of bodies designed to bring universities together.

Universities UK, GuildHE, Jisc, UCISA and AHUA: these organisations exist precisely to coordinate sector-wide responses to shared challenges. AI is among the most significant shared challenges UK

higher education has faced. Yet each university appears to be developing its AI policy in isolation, producing documents that range from genuine critical literacy resources to one-paragraph additions to misconduct procedures.

The Russell Group published principles on the use of generative AI tools in education in 2023, and several institutions in this study reference them.[15] But these are high-level commitments ('support students and staff to become AI-literate', 'adapt teaching and assessment'), not operational guidance. They do not address the structural questions this study raises: where a policy should sit; how trust should be framed; what student involvement looks like; or how to avoid the deficit model. There is no model AI policy, no agreed minimum standard, no shared vocabulary for the levels of AI use that might be permitted or encouraged and no sector-wide framework that institutions could adapt to their own context.

This is not an argument for uniformity. Institutional context matters. A conservatoire's relationship with AI-generated content differs fundamentally from an engineering faculty's. A research-intensive university faces different questions from one focused on professional practice. Any model policy would need to be adapted rather than adopted wholesale.

But the current situation, where 163 institutions with degree-awarding powers (and a much wider provider base behind them – the Office for Students' Register lists over 400 higher education providers in England alone) are independently developing responses to the same technology, produces unnecessary duplication, wide variation in quality and a sector-wide picture that students, staff, regulators and the public cannot see. The four exemplary policies identified in this study (from Durham, Stirling, Canterbury Christ Church and Arts University Plymouth) could serve as starting points for a sector conversation about what good AI policy looks like. Jisc, with its role in digital infrastructure and its existing work on AI in education, is one obvious convener.

Variation has consequences for students. A student transferring between institutions, or comparing their university's approach with a friend's, will find policies that differ in fundamental orientation, not only in detail. Some institutions treat AI as a tool to be learned. Others treat it as a threat to be contained. Students moving between institutions have no way of knowing whether their own university's approach is typical, progressive or unusually restrictive.

## What a student-centred AI policy would look like

A model policy should start with five principles.

**1. Location determines framing.** A policy embedded in a misconduct framework is likely to reflect the assumptions of that framework. AI guidance should sit within teaching and learning architectures. If institutional governance requires a misconduct reference, it should be a cross-reference rather than the primary home.

**2. Trust should be the default position, stated explicitly.** Conditional trust, where students must earn confidence through declarations and evidence retention, reproduces the same assumption with better vocabulary. Explicitly stating 'we trust students to make good decisions about AI' is a structural commitment that can shape every subsequent paragraph. It is also, based on the evidence from academic integrity research, more effective than surveillance at producing honest engagement.

Some will object that students themselves may want tighter controls, that not all students want to be trusted with AI when their peers might use it to cut corners. That concern is real. It also conflates two different things: trust in students and fairness in assessment. A policy can trust students to make good decisions about AI while simultaneously designing assessments that hold up against

misuse. Assessment design is the institution's responsibility. Detection is a response to assessment design that has not adapted to AI. If an assessment can be completed by pasting the brief into a chatbot, the assessment needs redesigning.

The intuitive alternative, pervasive surveillance and automated AI detection, fails on three counts. First, the technology is unreliable; its false positives become accusations the institution must investigate and the student must defend against. Secondly, detection produces an arms race. Each advance in detection drives a corresponding advance in evasion, so institutions chasing detection accuracy are chasing a moving target while diverting resources from teaching. Thirdly, the cost of being wrong falls hardest on students who can least afford it: a false accusation of academic misconduct can derail a degree, a visa or a career. The case for trust rests on the recognition that the alternative is structurally unfair, technically unreliable and operationally expensive.

**3. Student voice should shape the policy.** Arts University Plymouth included its Student Union President in its staff-led AI focus group and hosted a student-led symposium 'AI and the horizon of creative practice'. Most institutions showed no evidence of student involvement in policy development. A policy about students, written without students, will reproduce assumptions by default.

**4. Critical literacy should replace tool proficiency as the organising principle.** Teaching students how to use ChatGPT correctly is compliance training. Asking students to interrogate the assumptions, biases and power structures embedded in AI systems is education. The difference between these approaches is the difference between producing compliant users and producing critical thinkers.

The depth and form of critical AI literacy will vary by discipline. A Mathematics or Engineering student may spend more time on the algorithmic, statistical and computational properties of AI tools, while a Media Studies or History student may focus more on questions of bias, representation and epistemic authority. The underlying principle, that students should interrogate AI rather than simply operate it, applies across the sector. What each discipline does with that principle in practice should differ.

**5. Policies should be publicly accessible.** A policy that exists only behind an authentication wall serves an internal compliance function. Making it public is a commitment to transparency, a contribution to the sector conversation and a signal that the institution considers its AI position sound enough to withstand external scrutiny.

## Limitations of this analysis

**Selection bias:** The 96 accessible policies considered in this paper are the more visible and student-facing documents. Policies behind authentication walls may include more enforcement-oriented material. The selection may overrepresent education-framed policies, which would mean the real picture is worse than I found. Equally, hidden policies may be more student-friendly in ways this study cannot assess. The direction of the bias is unknown.

**Single coder:** I coded all 19 policies myself. The coding framework, decision rules and illustrative quotations provide transparency, but inter-coder reliability testing would strengthen the analysis. The framework is publicly available for anyone who wishes to code a sample and compare.

**The computational method identifies vocabulary, and vocabulary alone:** The framing ratio cannot detect when education vocabulary is being used for non-educational purposes. That gap between vocabulary and function is what makes the qualitative reading essential. Computational analysis can map the territory but it is the qualitative analysis that tells you what is actually there.

**Policies change:** This is a snapshot from mid-February 2026. Policies are live documents. Some will have changed since I scraped them. Some URLs in the data repository may no longer point to the current version of a policy, and some institutions may have additional AI guidance available through other channels. The scraped pages may represent one element of a broader institutional approach to AI that includes staff development, curriculum design and internal guidance not visible to the public.

**The data:** Everything needed to reproduce, challenge or extend this study is in the open data repository: all 96 policies with full text; the coding framework (31 codes with definitions); the keyword vocabularies (34 detection terms and 43 education terms, with the rationale for each); the framing ratios for every institution; the qualitative coding of 19 policies; and supplementary material with full coding framework detail. The method is transferable. Researchers in other countries could adapt the keyword vocabularies for local terminology and apply the same analytical framework. The hardest part is finding the policy URLs, which in itself reveals something important about the state of the field.

## Conclusion

Universities are supposed to develop critical thinkers. If an institution's own AI policy cannot model critical thinking about AI, if it resorts to compliance while claiming to educate, then the policy contradicts the mission.

The higher education sector has the bodies, expertise and examples it needs. At least four higher education institutions have already demonstrated AI policy can educate, extend trust and address risk through assessment design rather than detection. The remaining question is whether the sector will learn from them or continue to develop separate responses to the same challenge.

## Endnotes

1 Sam Illingworth, 'UK University AI Policies: Open Data Repository', GitHub, 2026 **https://github.com/sam-illingworth/uk-university-ai-policies**

2 Sam Illingworth and Rachel Forsyth, GenAI in Higher Education: Redefining Teaching and Learning, 2026 **https://www.bloomsbury.com/uk/genai-in-higher-education-9781350535787/**

3 University of Southampton, 'Using Generative AI During Your Studies' **https://www.southampton.ac.uk/about/governance/regulations-policies/policies/using-gen-ai-during-your-studies**

4 Loughborough University, 'Generative AI' **https://www.lboro.ac.uk/students/handbook/assessments/assessment-information/generative-ai/**

5 University of Glasgow, 'AI for Students' **https://www.gla.ac.uk/myglasgow/sld/ai/students/**

6 University of Liverpool, 'Guidance on the Use of Generative AI' **https://www.liverpool.ac.uk/media/livacuk/centre-for-innovation-in-education/digital-education/generative-ai-teach-learn-assess/guidance-on-the-use-of-generative-ai.pdf**

7 Donald L. McCabe, Linda K. Treviño and Kenneth D. Butterfield, 'Cheating in Academic Institutions: A Decade of Research', Ethics and Behavior, Volume 11, Issue 3, 2001, pp.219-232 **https://doi.org/10.1207/S15327019EB1103_2**

8 Tracey Bretag, Rowena Harper, Michael Burton, Cath Ellis, Philip Newton, Sonia Saddiqui, Karen van Haeringen and Pearl Rozenberg, 'Contract cheating: a survey of Australian university students', Studies in Higher Education, Volume 44, Issue 11, 2019, pp.1837-1856 **https://doi.org/10.1080/03075079.2018.1462788**

9 Durham University, 'Academic Misconduct/AI Policy and Guidance' **https://www.durham.ac.uk/departments/academic/common-awards/policies-processes/assessment/academic-misconductai-policy-and-guidance/**

Please note that this source is specific to the Common Awards Partnership at Durham University, a form of validation in collaboration with the Church of England and various partner theological educational institutions. Durham University's other AI guidance is available to staff and students via an internal site.

10 University of Stirling, 'AI Assessment Scale' **https://learnteach.stir.ac.uk/how-to/ai-assessment-scale/**; Canterbury Christ Church University, 'Generative AI Guidance' **https://students.canterbury.ac.uk/study/study-support-and-platforms/ generative-ai-guidance**

11 University of East Anglia, 'Generative AI Policy for Teaching and Learning' **https://assets.uea.ac.uk/f/185167/x/5558a0f812/generative-ai-policy-for-teaching-and-learning.pdf**

12 Rose Stephenson, The Hidden Impact of Menstruation in Higher Education, HEPI Report 191, July 2025 **https://www.hepi.ac.uk/reports/the-hidden-impact-of-menstruation-in-higher-education/**

13 Rose Stephenson and Charlotte Armstrong, Student Generative AI Survey 2026, HEPI Report 199, March 2026 **https://www.hepi. ac.uk/reports/student-generative-ai-survey-2026/**

14 Arts University Plymouth, 'Generative AI' **https://www.aup.ac.uk/reports-legal-privacy/generative-ai**

15 Russell Group, 'Russell Group principles on generative AI in education', 2023 **https://www.russellgroup.ac.uk/sites/default/files/2025-01/Russell%20Group%20principles%20on%20generative%20AI%20in%20education.pdf**