

# ELITE-AI

## Data Privacy, Fairness, Transparency and Responsible AI Use in ESP and Foreign Language Instruction

---

### Ethical Provisions Manual

---

Partner institutions: UNITBV (Romania) · UM (Slovenia) · UNIUD (Italy) · UNIPU (Croatia)

Erasmus+ Grant No. 2024-1-RO01-KA220-HED-000248845

[elite-ai.unitbv.ro](http://elite-ai.unitbv.ro)

## Table of Contents

|                                                                                              |           |
|----------------------------------------------------------------------------------------------|-----------|
| <b>1. Purpose and Scope of This Manual.....</b>                                              | <b>3</b>  |
| <b>2. International and European Ethical Foundations .....</b>                               | <b>3</b>  |
| 2.1 UNESCO: a human-centred, rights-based starting point.....                                | 4         |
| 2.2 The European Union's regulatory and ethical architecture .....                           | 4         |
| 2.3 Research ethics: the BERA lens .....                                                     | 4         |
| <b>3. Data Privacy, Security and Informed Consent .....</b>                                  | <b>5</b>  |
| <b>4. Bias, Fairness and Equity of Access .....</b>                                          | <b>6</b>  |
| <b>5. Transparency, Disclosure and Accountability .....</b>                                  | <b>7</b>  |
| <b>6. Pedagogical Integrity and Human Agency .....</b>                                       | <b>8</b>  |
| <b>7. Ethics in the WP3 Classroom: What the Consortium Observed.....</b>                     | <b>9</b>  |
| <b>8. The Policy Makers' Workshop, 17 June 2026 .....</b>                                    | <b>9</b>  |
| 8.1 Who should govern AI in education? .....                                                 | 10        |
| 8.2 Which uses of AI are most ethically concerning? .....                                    | 10        |
| 8.3 What stands between AI and educators? .....                                              | 10        |
| 8.4 Should ethics training be mandatory? .....                                               | 10        |
| 8.5 The closing provocation .....                                                            | 11        |
| 8.6 What the floor discussion added .....                                                    | 11        |
| <b>9. International Practice: How Universities Abroad Govern AI.....</b>                     | <b>11</b> |
| 9.1 Stanford University – default-restrictive, instructor-delegated .....                    | 12        |
| 9.2 University of Oxford – separated teaching and research tracks, free licensed access..... | 12        |
| 9.3 University of Cambridge – decentralised, discipline-specific framework .....             | 12        |
| 9.4 KU Leuven – responsible use over prohibition.....                                        | 12        |
| 9.5 A converging pattern .....                                                               | 13        |
| <b>10. What Participants Told Us: Evaluation Questionnaire Findings.....</b>                 | <b>13</b> |
| <b>11. Continuous Evaluation and Professional Development .....</b>                          | <b>14</b> |
| <b>12. A Consolidated Ethical Provisions Checklist.....</b>                                  | <b>14</b> |
| <b>13. Conclusion and Next Steps .....</b>                                                   | <b>16</b> |
| <b>References .....</b>                                                                      | <b>16</b> |
| International and European frameworks.....                                                   | 16        |
| University AI governance practice .....                                                      | 17        |
| Research on AI disclosure and use .....                                                      | 18        |
| ELITE-AI project sources.....                                                                | 18        |

## 1. Purpose and Scope of This Manual

This manual is the output of Activity A2 of Work Package 4 of ELITE-AI (Enhancing Language Instruction Through Ethical AI Integration in ESP), an Erasmus+ KA220-HED Cooperation Partnership bringing together Transilvania University of Braşov (UNITBV, Romania), University of Maribor (Slovenia), University of Udine (Italy) and Juraj Dobrila University of Pula (Croatia). Its brief, as set out in the project proposal, is to gather in one place the principles that should govern data privacy and security, informed consent, bias and fairness, transparency and accountability, the ethical use of student data, responsible AI development, pedagogical integrity, equity and inclusivity, continuous evaluation, and professional development. These principles apply wherever AI tools are introduced into English for Specific Purposes (ESP) and foreign language (FL) teaching.

Two bodies of evidence sit behind it. The first is external: the ethical frameworks published by UNESCO and the European Union, the research-ethics guidance issued by the British Educational Research Association, and the AI governance choices already made by universities well beyond the consortium. The second is internal: six months of classroom trials across four countries under Work Package 3, and the stakeholder workshop for policy makers held online on 17 June 2026 under Activity A4, in which around 95 participants from at least a dozen institutions and countries contributed live polling responses, open-text comments and floor discussion.

The manual is not a legal text and does not replace institutional data-protection or research-ethics review. It is a practical companion to the three documents already published by the consortium – *The Quality Assurance Manual* (WP4) (ELITE-AI Consortium, 2026a), *The Manual of Good Practices for Researchers* (WP2) (ELITE-AI Consortium, 2025a), and *The Manual of Good Practices for Teachers and Students* (WP3) (ELITE-AI Consortium, 2025b) – and it is written for the audience Activity A2 specifies: educators, administrators and the policy makers who will shortly receive its companion volume, *The Manual of Good Practices for Policy Makers*, due at the close of the project in August 2026.

### How to read this manual

**Sections 2–6** set out the ethical foundations – international, European and disciplinary – and translate each into classroom-level and institution-level guidance.

**Sections 7–9** report what the consortium’s own trials, the policy makers’ workshop and comparable universities abroad are actually doing.

**Sections 10–12** close the loop with questionnaire evidence, a professional-development agenda and a consolidated checklist that institutions can adapt directly.

## 2. International and European Ethical Foundations

Three bodies of guidance anchor this manual: UNESCO’s global, rights-based framework; the European Union’s regulatory and ethical architecture, which is binding rather than advisory for the four partner institutions; and the research-ethics tradition represented here by the British Educational Research Association. Each addresses a different layer of the problem, and together they cover the full span of Activity A2’s brief.

## 2.1 UNESCO: a human-centred, rights-based starting point

UNESCO's 2021 'Recommendation on the Ethics of Artificial Intelligence' and its 2023 'Guidance for Generative AI in Education and Research' ([unesdoc.unesco.org/ark:/48223/pf0000381137](https://unesdoc.unesco.org/ark:/48223/pf0000381137) and [unesdoc.unesco.org/ark:/48223/pf0000386693](https://unesdoc.unesco.org/ark:/48223/pf0000386693)) set the tone that the rest of this manual follows: AI should support human agency and human oversight rather than substitute for either. The 2023 guidance is specific about consequences – it asks governments and institutions to validate generative AI tools for age-appropriateness and pedagogical soundness before they reach a classroom, to protect the data of minors and young adults with particular care, and to build AI literacy into both teacher training and the curriculum rather than leaving competency to chance. UNESCO's 2024 'AI Competency Frameworks for Students' ([doi.org/10.54675/JKJB9835](https://doi.org/10.54675/JKJB9835)) and 'AI Competency Frameworks for Teachers' ([doi.org/10.54675/ZJTE2084](https://doi.org/10.54675/ZJTE2084)), and its 2025 follow-up report 'AI and the Future of Education: Disruptions, Dilemmas and Directions' ([unesco.org/en/articles/ai-and-future-education-disruptions-dilemmas-and-directions](https://unesco.org/en/articles/ai-and-future-education-disruptions-dilemmas-and-directions)), extend the same logic: literacy and critical judgement are treated as prerequisites for safe use, not optional extras.

## 2.2 The European Union's regulatory and ethical architecture

For the consortium's four institutions, EU law is not background reading; it is the compliance floor. In the 'Regulation' (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (<https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>), laying down harmonised rules on artificial intelligence implementing the proposal commonly referred to as the 'Artificial Intelligence Act' first published in 2021, which classifies systems used to evaluate learning outcomes, to determine admission, or to monitor students during assessment as high-risk (Annex III.3), thus bringing obligations for fundamental rights impact assessment under Article 27 and, from 2 August 2026, the full Annex III regime. Article 50 requires that AI-generated or AI-manipulated text, audio, image or video content be marked as such wherever it is deployed – a rule with direct relevance to any AI-generated teaching material. 'The General Data Protection Regulation' (<https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>) applies in parallel: institutions act as data controllers whenever they decide how learner data is processed, they must communicate that processing in plain language under Articles 12–15, and a data protection impact assessment is required under Article 35 before any high-risk system, AI included, is put into use.

'The European Commission's Ethical Guidelines on the Use of Artificial Intelligence and Data in Teaching and Learning for Educators' ([doi.org/10.2766/7967834](https://doi.org/10.2766/7967834)), first published in 2022 and substantially revised in 2026, translate this legal frame into seven working dimensions that this manual borrows throughout: human agency and oversight, transparency and explainability, diversity and inclusion, fairness and non-discrimination, societal and environmental wellbeing, privacy and data governance, and accountability. The revised guidelines are blunt about a common misconception worth repeating here – that AI tools are neutral because they are mathematical. In reality, they are trained on human-generated data and can reproduce the assumptions and gaps in that data, which is precisely why the fairness questions in Section 4 below matter as much for a grammar-checking tool as for an admissions algorithm.

## 2.3 Research ethics: the BERA lens

'The British Educational Research Association's Ethical Guidelines for Educational Research' (<https://www.bera.ac.uk/wp-content/uploads/2024/04/Ethical-Guidelines-for-Educational-Research-5th-edition.pdf>), now in a fifth edition, were written for a wider purpose than AI, but two of its principles map directly onto ELITE-AI's own practice. The first concerns consent specificity: where

a study uses digital or AI-based tools, BERA expects consent forms to be explicit about what is collected, how it is stored, and for how long – a standard the consortium had already adopted in its own observation-sheet protocol, and one echoed almost word for word in the *Manual of Good Practices for Researchers* (ELITE-AI Consortium, 2025a). The second concerns jurisdiction: BERA reminds researchers that data protection obligations are set nationally (in the UK, the Data Protection Act 2018 alongside the GDPR) and that researchers gathering data online should track how those obligations evolve rather than assume a single one-off compliance check is sufficient. Every WP3 classroom trial obtained ethics committee approval at its home institution before data collection began, a step consistent with this guidance and observed by each partner.

| Framework                                  | Primary emphasis                                  | Where it binds this manual |
|--------------------------------------------|---------------------------------------------------|----------------------------|
| UNESCO (2021, 2023, 2024)                  | Human oversight, age-appropriateness, AI literacy | Sections 3, 6, 11          |
| EU AI Act – GDPR – Ethical Guidelines 2026 | Legal compliance, seven ethical dimensions        | Sections 3, 4, 5, 6        |
| BERA Ethical Guidelines, 5th ed.           | Consent details, jurisdictional compliance        | Section 3                  |

### 3. Data Privacy, Security and Informed Consent

Every AI tool used in an ESP classroom collects something – a writing sample, a pronunciation recording, a log of which prompts a student tried and discarded – and much of that data leaves the institution’s own servers the moment the tool is called. *The Manual of Good Practices for Researchers* (ELITE-AI Consortium, 2025a), produced earlier in this project, already states the working rule for research contexts: informed consent forms for AI-integrated studies should name each tool, describe what it collects, say where that data is stored, for how long, and whether it may be used to train future models. This manual extends the same rule to everyday teaching, where no ethics committee reviews the activity but the data at stake is identical.

The GDPR gives this rule legal weight. Articles 12 to 15 require that institutions, acting as data controllers, explain their processing of learner data in language a student can actually follow, not in a linked terms-of-service document written for lawyers. Article 35 requires a data protection impact assessment before any system that poses a meaningful risk to learners’ rights is deployed – a threshold that most generative AI tools used for grading, feedback or profiling will meet. Where the AI Act’s Article 27 fundamental rights impact assessment also applies, the two assessments should be run together rather than treated as separate paperwork exercises, since they draw on much of the same information about what the tool does and who it affects.

Consent in an AI-integrated classroom has one complication that a standard consent form does not: the tool itself may change its data practices after a course begins, particularly where an institution moves from a free tier to a licensed one, or where a vendor updates its own policy. WP3’s own experience – eight teachers across four countries, twenty-nine institutions represented at the June workshop, and a repository of one hundred catalogued tools – argues for treating tool selection as an ongoing responsibility rather than a one-off approval, and for keeping a simple, living record of which tools are in use in which course, updated each semester rather than each accreditation cycle.

### Consent and privacy: minimum provisions

1. Name every AI tool used in a course, in plain language, before the course begins.
2. State what each tool collects, where it is stored, and whether it trains on submitted data.
3. Treat free-tier tools with the same scrutiny as licensed ones – free tools often monetise data instead of a subscription fee.
4. Run GDPR Article 35 and AI Act Article 27 assessments together where a tool touches assessment or profiling.
5. Review the tool list every semester, not only at the point of initial approval.

## 4. Bias, Fairness and Equity of Access

Fairness in an ESP classroom has two distinct faces and combining them is a common mistake. The first is algorithmic: an AI tool trained mainly on general-purpose English may perform less reliably on the specialised registers this project actually researches – criminal-justice terminology, tourism correspondence, early-years pedagogy, students/teachers mobility practices – and may contain cultural or linguistic assumptions that disadvantage some learners more than others. The second is structural: access to AI tools is not evenly distributed. The consortium's own slide on digital equity puts this plainly – the ESP fields observed across the four countries, from law to sport science and tourism, vary enormously in the specialist AI tools available to them, and resource-allocation decisions have to account for that variation rather than assume a single institutional licence solves the problem for every discipline.

The revised EU ethical guidelines frame both faces under the same two headings, fairness and non-discrimination, and diversity and inclusion, and the working questions they propose translate directly into ESP practice: has the tool been checked for bias across different learner groups before adoption?; does it take a learner's age, comprehension level and disciplinary vocabulary into account?; and, where a student lacks reliable access outside class, is there an offline or lower-bandwidth alternative that does not put them at a disadvantage?. These are not rhetorical questions. When policy makers at the 17 June workshop were asked, in an open-text poll, what the most significant barrier was between AI and educators at their institution, 'lack of funding for paid, reliable AI tools' and 'lack of AI literacy' were among the most-repeated and most-endorsed answers, the latter drawing the single highest number of agreement votes (ten) of any response in that poll.

Equity of access depends on more than institutional provision alone. A tool that is licensed institution-wide but poorly understood by staff produces the same disadvantage as no licence at all – a point several workshop participants made independently and one this manual returns to in Section 11 on professional development.

### Fairness and equity: questions worth asking before adoption

**For the course:** Has this tool been tested for bias across the learner groups actually enrolled, not a generic population?

**For the discipline:** Does the tool handle the specialised register of this ESP field – legal, medical, tourism, technical – or only general English?

**For access:** Do all students have equivalent access outside class, and if not, what is the offline or low/no-cost alternative?

**For redress:** Is there a named, quick route for a student or teacher to flag a biased or unfair output once the tool is in use?

## 5. Transparency, Disclosure and Accountability

Academic publishing offers this manual a ready-made case study in graduated disclosure, because four major publishers have already had to answer, in binding policy, the exact question ESP instructors now face in the classroom: when AI assists with language, does that assistance need to be named? The comparison was presented at the consortium's own 7 May 2026 workshop and is reproduced here because the underlying logic transfers directly to teaching.

| Publisher                  | Position                    | Core requirement                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Elsevier                   | Permissive, with disclosure | A named declaration statement at submission; AI may never substitute for human critical judgement, and every AI-suggested reference must be independently verified ( <a href="https://elsevier.com/about/policies-and-standards/the-use-of-generative-ai-and-ai-assisted-technologies-in-writing-for-elsevier">elsevier.com/about/policies-and-standards/the-use-of-generative-ai-and-ai-assisted-technologies-in-writing-for-elsevier</a> ). |
| Oxford University Press    | Strictest: prior approval   | Written permission from the editor is required before any AI-generated text, data analysis, graph or image is submitted ( <a href="https://academic.oup.com/pages/for-authors/books/author-use-of-artificial-intelligence">academic.oup.com/pages/for-authors/books/author-use-of-artificial-intelligence</a> ).                                                                                                                              |
| Cambridge University Press | Most transparent            | AI cannot be listed as author; uniquely, Cambridge contacted 20,000 authors in October 2024 to seek consent before licensing their work for model training ( <a href="https://cambridge.org/news-and-insights/cambridge-launches-ai-research-ethics-policy">cambridge.org/news-and-insights/cambridge-launches-ai-research-ethics-policy</a> ).                                                                                               |
| Peter Lang                 | Strictest on peer review    | AI is permitted only for language checking; reviewers are explicitly banned from uploading any manuscript into an AI system ( <a href="https://peterlang.com/ai-statement">peterlang.com/ai-statement</a> ).                                                                                                                                                                                                                                  |

Three principles recur across all four policies, and this manual adopts them as the working definition of accountability for the consortium's own classrooms. Human accountability cannot be delegated: if an AI tool fabricates a reference, a fact or a grammatical judgement and nobody checked it, the human who submitted the work is responsible, not the tool. Disclosure is not optional, though what it must cover varies by context and risk – a named-tool-and-purpose statement, as Elsevier requires, is the minimum that should apply to any piece of ESP coursework produced with functional AI assistance. And confidentiality in peer review, or its classroom equivalent – a classmate's draft, a colleague's exam paper – is treated as absolute by every publisher examined; no such material should enter an AI system without the author's explicit permission, a rule several workshop participants flagged as the least well known of all.

The evidence on how well disclosure is actually working is mixed, and worth stating precisely rather than approximately. A *Nature* survey of five thousand researchers, conducted by Diana Kwon and published in May 2025, found deep and persistent division over whether AI use in manuscript preparation should be declared at all, and if so, in how much detail (Kwon, 2025: 574–578). A separate excess-vocabulary analysis of biomedical abstracts by Dmitry Kobak and colleagues estimated that

roughly one in eight 2024 PubMed abstracts showed detectable signs of large-language-model assistance (Kobak, D. et al., 2025: eadt3813) – a figure close to the fourteen per cent cited at the consortium’s own Ethics Talk, and one that illustrates the same point made throughout this section: written policy and lived practice do not yet match, in research or in the classroom, and closing that gap is exactly the task this manual sets out to support.

At the classroom level, the consortium’s own WP3 trials already show a version of this gap closing itself, ahead of any formal policy. ELITE-AI’s summary of what the observation sheets recorded notes that ethical boundaries, including disclosure norms, emerged spontaneously from the very first sessions, before any institution had issued formal guidance – a bottom-up pressure for governance that this manual treats as evidence in its own right, not merely as an aspiration.

### A minimum disclosure statement

‘[Tool name, version] was used to [generate / edit / translate / check] [specific part of the work], and the output was reviewed and substantially revised by [author].’

**Why it works:** it names the tool, states the purpose, and asserts human oversight – the three elements every publisher policy in this section requires in some form.

## 6. Pedagogical Integrity and Human Agency

The single finding that reappears most often across this project, in the WP3 observation sheets, in ELITE-AI’s workshop presentations (Costantini, 2026; Nechifor, 2026a, 2026b) and in the floor discussion on 17 June, is a one-line rule: AI should function as support, not as a replacement. It is easy to state and harder to operationalise, and this section sets out what operationalising it actually involves.

The clearest practical consequence, raised repeatedly during the policy makers’ workshop, is a shift in what assessment rewards. Speaking to a participant’s account of an institution with no AI-specific regulation – only outdated plagiarism rules – consortium partners argued that the more defensible response is to weight the drafting process, not only the finished product: a student’s evolving drafts, the decisions they made along the way, and their ability to explain and defend those decisions, all become harder for an AI tool to fabricate than a polished final paragraph. Punishing a student for using AI to polish the English of a piece they wrote and researched themselves, it was noted, sits oddly next to the fact that journal editors have always routinely asked published researchers to do exactly that, i.e. to polish and proofread their final manuscripts for language and style errors, whether by themselves or through commissioned work by language specialists.

A second consequence concerns detection. The consortium’s own testing found that AI-detection software returns false positives even on text written well before generative AI existed – passages from nineteenth-century authors were flagged as AI-generated in informal tests described during the workshop discussion. Relying on detection scores as evidence of misconduct is therefore not a safe default; a policy built primarily around catching AI use with a detector is a policy built on a measurement tool known to fail in both directions. The more durable approach, consistent with the EU guidelines’ human agency and oversight dimension, keeps a teacher’s informed judgement, exercised through process-based assessment and direct conversation with the student, ahead of any automated signal.

A third consequence is definitional rather than procedural: several contributors, from a university director in Italy to a lecturer in Slovenia, argued independently that institutions need to define what counts as an ethical violation before they can enforce one. Training that helps staff understand how a given AI tool actually works has to come before disciplinary policy, not after it – otherwise the policy ends up penalising unfamiliarity rather than misconduct.

### Human agency: guiding questions for the classroom

**Before the task:** Can I explain to students why AI is or is not appropriate for this specific activity, and not only that it is or is not allowed?

**During the task:** Do I know what a good AI-assisted answer looks like well enough to recognise a poor one?

**After the task:** Am I assessing the process a student went through, and not only the artefact they submitted?

**Institutionally:** Have staff received tool-specific training before being asked to enforce a tool-specific policy?

## 7. Ethics in the WP3 Classroom: What the Consortium Observed

Work Package 3 put these principles into practice before this manual existed to state them. Eight teachers – Andreea Nechifor and Cristina Dimulescu at UNITBV, Mojca Kompara Lukančič and Tilen Smajla at the University of Maribor, Eleonora Gallitelli, Emanuela Li Destri and Francesco Costantini at the University of Udine, and Jelena Gugić at the Juraj Dobrila University of Pula – ran AI-integrated sessions across pragmatics, applied grammar, ESP, criminal justice terminology, tourism English and early-years pedagogy, and recorded what happened on a shared, self-designed observation sheet, generating a dataset the project describes as statistically substantial in its own right.

Four findings from that dataset carry direct ethical weight. The shift from a teacher-centred to a shared, student-initiated use of AI tools happened gradually and was itself something the observation sheets could track, not merely assume. No single AI tool dominated the sessions – tool variety mattered, with different tools serving distinct purposes even within a single fifty-minute class, which argues against any policy written around one named product. Ethical boundaries, including disclosure norms, emerged from the very first classes, ahead of formal institutional policy in every partner country – the bottom-up pressure for governance already noted in Section 5. And AI access itself is not uniform: *The Manual of Good Practices for Researchers* (ELITE-AI Consortium, 2025a), drafted earlier in the project, warns explicitly that a study run at a well-resourced European university with institutional tool licences is not automatically transferable to a context where any of those conditions are absent, and that low engagement with a tool may reflect cost or connectivity barriers rather than learner attitude.

These four findings are the empirical backbone of Sections 3 to 6 above, where they are not stated as abstract principle but as something the consortium watched happen, across four countries, over six months of trials.

## 8. The Policy Makers' Workshop, 17 June 2026

Activity A4 organised an online workshop for policy makers, educational leaders and institutional decision-makers on 17 June 2026, 11:00–12:00 EET. Registration closed the previous day; the session drew participants from the Erasmus+ partner institutions plus visiting policy makers and academics from Spain, Ghana, China-based programmes and several Italian and Romanian universities, among others. Partner-researchers presented the project's results and its account of innovation in AI-assisted language teaching and then ran five live Slido polls before the floor was opened for open discussion. What follows is drawn directly from the poll results (ELITE-AI Consortium, 2026c) and the workshop discussion, and it is the evidentiary core of this manual, since it is the one section built entirely from what policy makers themselves said rather than from published guidance.

### 8.1 Who should govern AI in education?

The opening poll asked participants, in free text, who should be responsible for creating and enforcing AI governance policy in education. The resulting word cloud, drawn from thirty-nine responses, converged on a small set of terms used repeatedly: government, university, ministry, teachers, students, institutional policy makers, rectoral team and AI experts. The live synthesis of the responses, mirroring the partners' opinion and approach, described the emerging shape as a deliberately layered, top-down structure – ministry, then university, then faculty, then individual instructor – with the explicit caution that the appropriate degree of central control shrinks at each layer, because subject matter varies so widely even among linguists teaching in the same department, let alone between arts and applied science faculties in the same institution.

### 8.2 Which uses of AI are most ethically concerning?

The second poll asked participants to select, from six options, which uses of AI in education they found most ethically concerning; multiple selections were permitted. One answer dominated from the first responses onward and finished at 98% of all votes: students submitting AI-generated text without acknowledging it. Teachers using AI to grade student essays followed at roughly 43–46%, researchers using AI to design research plans at around 19%, teachers using AI to design assessments at about 10%, and students using AI for translation and teachers using AI to create educational content tied for last place at around 2% each. The gap between the top answer and everything else is the single clearest signal to emerge from the entire workshop: undisclosed generation, not AI use itself, is where policy makers locate the ethical risk.

### 8.3 What stands between AI and educators?

An open-text poll asked what the most significant barrier was between AI and educators at participants' own institutions. Answers clustered around five recurring themes, several of which participants would 'like' to register agreement with an existing answer rather than repeat it: a lack of AI literacy among staff (the single most-endorsed answer, with ten agreement votes); the absence of any institutional regulation at all ('there are no regulations at all at my institution – we as teachers are unsupported and abandoned, left to our own devices,' one participant wrote); insufficient teacher training and workshops; unresolved copyright questions; and, less frequently but pointedly, 'prejudice against technology' among some colleagues. 'Lacking good practices for meaningful integration of AI into courses' drew the second-highest number of agreement votes, at eight.

### 8.4 Should ethics training be mandatory?

Asked directly whether AI ethics training should be mandatory for all staff who use AI tools, 86% of respondents answered yes, 3% no, and 11% 'it depends.' This is the most decisive result of any poll in

the workshop and is treated in Section 11 below as a direct mandate for the professional-development provisions of this manual.

### 8.5 The closing provocation

The final poll asked participants to finish the sentence: ‘The one thing my institution still hasn’t figured out about AI is...’ Answers ranged from the procedural – ‘how to make AI policies in the first place,’ ‘no clear guidelines’ – to the more reflective: one participant wrote that their institution had not worked out ‘how to balance the raw speed of technological evolution with the deeply human pace of genuine learning and ethical reflection,’ and another, more bluntly, that AI risked ‘killing the critical thinking at students’ creativity.’ Several answers converged on a single proposed remedy – a mandatory course or module in AI literacy for both students and staff – that this manual adopts as a concrete recommendation in Section 12.

### 8.6 What the floor discussion added

Four contributions from the open floor add texture the polls cannot capture on their own. One participant, from the University of Santiago de Compostela in Spain, described their institution’s regulations as addressing plagiarism in detail while saying nothing about AI at all, leaving teachers to treat AI use as a matter of individual judgement and, in their words, to frame it as something to be fought rather than managed – ‘people are abandoned ... it’s like the enemy.’ By contrast, another attendee, speaking on behalf of the University of Udine as the Rector’s delegate for AI, described an institution that has already developed AI policies and is actively working on their implementation across teaching and research. Yet his intervention also highlighted how demanding this process is: universities must reconcile European, national and regional regulations, accommodate very different uses of AI across disciplines, and continuously update guidance as technologies evolve. He stressed that reliable AI-detection tools do not exist and argued that assessment should focus instead on whether students can explain and critically engage with their work.

An instructor from the University of Ghana reported a different failure mode entirely: her institution has an approved AI policy covering students, researchers and teachers, but it has not been disseminated – ‘it’s just in the shell’ – so staff outside the committee that drafted it are largely unaware of its content and fall back on individual discretion regardless. A professor who teaches Slovene at a Chinese university affiliated with the University of Ljubljana, offered the sharpest contrast of the day: where European institutions are still debating whether to permit AI in the classroom, her Chinese institution has instructed staff to use it, and the university is restructuring its language curriculum partly in response to AI-driven competitive pressure for students. Her own conclusion, delivered to the room, was that training has to come before regulation – ‘we need to do something, but we are not skilled, actually’ – a sentiment project partners echoed directly in closing that segment of the discussion.

## 9. International Practice: How Universities Abroad Govern AI

The workshop discussion referenced Stanford, Oxford and Cambridge as early movers on institutional AI policy, and speakers noted that the consortium had already begun researching comparable universities’ approaches. This section extends that comparison to five institutions with published, verifiable policies, chosen to represent a spread of regulatory philosophies rather than a single model.

## 9.1 Stanford University – default-restrictive, instructor-delegated

Stanford's Generative AI Policy Guidance, issued by its Office of Community Standards, sets a default rule rather than a blanket ban: absent an explicit statement from the course instructor, using or consulting generative AI is treated exactly as assistance from another person would be under the Honor Code – which means substantial undisclosed AI use in an assessed exam or assignment is, by default, a violation. The weight of decision-making sits with individual instructors and departments, who are expected to publish their own course-level AI policy in the syllabus; Stanford's Teaching Commons provides templates for exactly this purpose. An Academic Integrity Working Group, formed in 2024, continues to revise guidance each academic year ([communitystandards.stanford.edu/generative-ai-policy-guidance](https://communitystandards.stanford.edu/generative-ai-policy-guidance); [teachingcommons.stanford.edu/policy-guidance](https://teachingcommons.stanford.edu/policy-guidance)).

## 9.2 University of Oxford – separated teaching and research tracks, free licensed access

Oxford maintains two distinct policies rather than one: a Guidance on Safe and Responsible Use of GenAI for students and teaching, and a separate Policy for Using Generative AI in Research overseen by its Research Practice Sub-Committee, which requires that substantive AI use be disclosed and that researchers protect intellectual property and sensitive data. Since September 2025, Oxford has provided every member of staff and student with free, single-sign-on access to ChatGPT Edu, paired with mandatory information-security training that now includes an AI component – a deliberate attempt to close the literacy gap this manual's own workshop evidence identifies as the largest single barrier ([ox.ac.uk/gen-ai](https://ox.ac.uk/gen-ai); [ox.ac.uk/research/support-researchers/research-practice/policy-generative-ai-research](https://ox.ac.uk/research/support-researchers/research-practice/policy-generative-ai-research); [ox.ac.uk/students/life/it/genai-tools/guidance-on-safe-and-responsible-use-of-genai](https://ox.ac.uk/students/life/it/genai-tools/guidance-on-safe-and-responsible-use-of-genai)).

## 9.3 University of Cambridge – decentralised, discipline-specific framework

Cambridge deliberately declined to write a single university-wide rule, on the grounds that appropriate AI use differs too much between a mathematics tripos and a modern languages one. Its AI Policy Framework instead gives faculties, departments and colleges a common structure to adapt locally, sitting alongside a university-wide statement on AI and assessment and a firm, absolute rule that examiners may never upload student work into an AI system. First published in February 2025 and already revised twice, the guidance is scheduled for another review in summer 2026 – an explicit acknowledgement that a policy written once and left alone will not keep pace with the technology ([blendedlearning.cam.ac.uk/artificial-intelligence-and-education/ai-policy-framework](https://blendedlearning.cam.ac.uk/artificial-intelligence-and-education/ai-policy-framework); [information-compliance.admin.cam.ac.uk/data-protection/guidance/ai-guidance](https://information-compliance.admin.cam.ac.uk/data-protection/guidance/ai-guidance)).

## 9.4 KU Leuven – responsible use over prohibition

KU Leuven, in Belgium, is frequently cited by independent policy trackers as an early European leader in this space for choosing a responsible-use framework over a ban when GenAI first reached classrooms. Its rule for teaching is simple and enforceable: instructors must state clearly whether GenAI is permitted for a given assignment, and students must be transparent about their use of it; undisclosed substantial AI use is treated as an irregularity under the university's examination regulations. The policy explicitly ties classroom practice to the European Code of Conduct for Research Integrity's requirement that researchers report their use of AI and automated tools in a way that allows verification ([kuleuven.be/english/education/student/educational-tools/responsible-use](https://kuleuven.be/english/education/student/educational-tools/responsible-use)).

of-generative-artificial-intelligence; [nieuws.kuleuven.be/en/content/2023/ku-leuven-opts-for-responsible-use-of-generative-artificial-intelligence-in-research-and-education](https://nieuws.kuleuven.be/en/content/2023/ku-leuven-opts-for-responsible-use-of-generative-artificial-intelligence-in-research-and-education)).

## 9.5 A converging pattern

Read together with the floor discussion reported in Section 8 – Italy’s evolving national coordination, Ghana’s dissemination gap, and the mandated, top-down adoption as described in the Chinese institution – three lessons recur across every example in this section regardless of legal system or national culture. Disclosure obligations are now close to universal, but institutions calibrate how much detail they require to the actual risk of the specific use, not to AI use in the abstract. Authority is devolved to the course or discipline level almost everywhere examined, because a single fixed rule cannot survive contact with the range of subjects a comprehensive university teaches – the same point project partners made from the ELITE-AI floor. And every policy reviewed for this manual is explicitly provisional, revised on a cycle measured in months rather than years. None of the five institutions in this section, and none of the consortium’s own four, has published a policy it describes as finished.

## 10. What Participants Told Us: Evaluation Questionnaire Findings

Thirty-two policy makers completed the post-workshop evaluation questionnaire (ELITE-AI Consortium, 2026b) comfortably short of overcounting but consistent with the consortium’s own target of at least forty stakeholders across the two teacher- and policy-facing workshops combined. The results corroborate the poll data reported in Section 8 rather than complicating it.

| Question                                               | Result                                                 |
|--------------------------------------------------------|--------------------------------------------------------|
| Topics were interesting                                | 29 of 32 agreed or strongly agreed (18 strongly agree) |
| Workshop improved knowledge or skills                  | 23 of 32 agreed or strongly agreed                     |
| Gave concrete ideas for institutional AI policy        | 26 of 32 agreed or strongly agreed                     |
| Discussion enabled reflection on governance challenges | 27 of 32 agreed or strongly agreed                     |
| Speakers rated 5 out of 5                              | 26 of 32 respondents                                   |
| Workshop overall rated 5 out of 5                      | 23 of 32 respondents (remaining 9 rated it 4)          |
| Interested in accessing project resources              | 31 of 32 respondents                                   |

The open-text responses add the qualitative register the ratings alone cannot. Several participants described the session as reinforcing rather than changing an existing view – ‘it did not change my perspective, I would say it reinforced it’ – which this manual reads as a sign of convergence rather than redundancy: independent policy makers arriving at the same conclusions as the consortium’s own trial data is itself evidence the recommendations in Section 12 are on solid ground. Others named a specific gap this manual addresses directly: one respondent wrote that the workshop ‘made me think of the need of an institutional approach on using AI in education in general, not only in language teaching,’ and another, writing in Romanian, put the same point as a matter of principle – that AI tools are now part of the contemporary educational landscape and must be used in teaching and

research, but with respect for academic ethics rules that remain, at this point, largely implicit rather than written down. A third theme, raised independently by several respondents, was a direct request for more of this kind of training: ‘I would like to attend more AI related workshops,’ one wrote; another asked for sessions that go into ‘even more detail.’

## 11. Continuous Evaluation and Professional Development

Every strand of evidence gathered for this manual points to the same conclusion: the largest ethical risk in AI-integrated language teaching is not malicious misuse but simple unfamiliarity, among staff who have not been trained, in institutions that have not disseminated the policy they already wrote, applying rules that were current a year ago but are not current now. Eighty-six per cent of workshop respondents said ethics training should be mandatory for all staff who use AI tools; the single most-endorsed answer to the barriers question was a lack of AI literacy; and one floor comment – ‘we need to do something, but we are not skilled’ – was echoed, in substance, by policy makers from at least three other countries during the same session.

This manual therefore recommends training before enforcement as a sequencing principle, not only a preference. A disciplinary policy applied to staff who have not been shown how a given tool actually works risks penalising unfamiliarity rather than misconduct, a distinction several workshop participants raised without prompting. Cambridge’s decision to revise its guidance twice within a year of first publishing it, and Italy’s expectation of a quarterly review cycle reported by one participant, both model the alternative this manual recommends: treat AI policy as a living instrument with a fixed review date, not a document issued once and left to age.

Monitoring, in this manual’s own terms, does not require new infrastructure. WP3’s observation-sheet methodology – structured, cross-institutional and repeatable, already piloted across four countries and nearly two thousand responses when combined with *The Manual of Good Practices for Teachers and Students* (ELITE-AI Consortium, 2025b) – is itself a model other institutions can adopt for monitoring AI integration without adding significant burden to teaching staff.

### Professional development: a minimum agenda

1. A mandatory, short AI-literacy module for all staff who use AI tools in teaching, assessment or research – the direct answer to the workshop’s 86% vote.
2. A named institutional contact or committee responsible for keeping the AI policy current, with a fixed review interval no longer than one academic term.
3. Active dissemination of any approved policy – an approved policy that staff have never seen, as reported from Ghana, is functionally no policy at all.
4. Peer observation and shared reflection, on the WP3 model, as the primary monitoring instrument rather than detection software.

## 12. A Consolidated Ethical Provisions Checklist

This section gathers the recommendations made throughout the manual under the nine headings Activity A2 specifies, so that an institution or policy maker can work through it directly rather than search back through the preceding sections.

### 1. Data privacy and security

- Name every AI tool in use, in plain language, before a course begins.
- Run GDPR Article 35 and, where relevant, AI Act Article 27 assessments together.
- Review the tool list every semester.

### 2. Informed consent

- State what each tool collects, where it is stored, and whether it trains on submitted data.
- Extend research-grade consent practice to everyday teaching, not only formal studies.

### 3. Bias and fairness

- Check tools for bias across the learner groups actually enrolled.
- Verify a tool's performance on the specialised ESP register it will actually be used for.
- Provide an offline or low/no-cost alternative wherever access is uneven.

### 4. Transparency and accountability

- Require a minimum disclosure statement – tool, purpose, extent of human review – for functional AI-assisted work.
- Treat peer material – a classmate's or colleague's draft – as confidential by default; never upload it to an AI system without permission.
- Do not rely on AI-detection scores as evidence of misconduct.

### 5. Ethical use of student data

- Treat learner data collected by third-party AI tools with the same care as institutional records.
- Communicate data practices in language appropriate to the learner's age and level.

### 6. Responsible AI development and procurement

- Prefer institutionally inspected or licensed tools over unreviewed free tiers where sensitive data is involved.
- Pilot a tool with a small group before full-course adoption, on the WP3 model.

### 7. Pedagogical integrity

- Weight process – drafts, decisions, defensible reasoning – alongside the finished product.
- Keep AI as support for a learning aim the teacher defines, never as a substitute for it.

### 8. Equity and inclusivity

- Treat disciplinary variation in AI tool access as a resource-allocation problem, not a discipline problem.
- Consider AI literacy training as seriously as tool licences – a licensed tool nobody understands helps no one.

### 9. Continuous evaluation and professional development

- Make AI ethics training mandatory for staff who use AI tools – the position 86% of workshop respondents endorsed.
- Set a fixed policy review interval, no longer than one academic term.
- Actively disseminate any approved policy; an unread policy governs nothing.

## 13. Conclusion and Next Steps

This manual closes the loop described throughout the 17 June workshop: the ideas raised that day by the participants' contributions, together with six months of WP3 classroom evidence, the ethical frameworks already governing some of the consortium's four institutions, as well as already existing provisions, manuals and instructions across well-known educational establishments, have been assembled into 'The Ethical Provisions Manual', which constitutes Activity A2's output. Its recommendations are deliberately practical rather than aspirational, each one tracing back to a specific piece of evidence reported somewhere in this document, whether a GDPR article, a workshop poll result, or an observation the consortium made in its own classrooms.

The consortium's closing message to the 17 June workshop applies equally here: the aim was never to arrive at a finished, permanent policy, but at a coherent and evidence-based starting point that individual institutions can adapt to their own context – and one built, unusually for a document of this kind, as much from what policy makers said in an hour of live discussion as from what any framework document prescribes in advance.

## References

### International and European frameworks

European Commission. (2026). Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators (revised ed.). Publications Office of the European Union.

European Commission. (2021). Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) (COM/2021/206 final). [eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206)

European Commission. (2020). Digital Education Action Plan (2021–2027). [education.ec.europa.eu/focus-topics/digital-education/action-plan](https://education.ec.europa.eu/focus-topics/digital-education/action-plan)

UNESCO. (2021). Recommendation on the ethics of artificial intelligence. [unesdoc.unesco.org/ark:/48223/pf0000381137](https://unesdoc.unesco.org/ark:/48223/pf0000381137)

Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. [doi.org/10.54675/EWZM9535](https://doi.org/10.54675/EWZM9535)

- Miao, F., & Shiohira, K. (2024). AI competency framework for students. UNESCO. doi.org/10.54675/JKJB9835 — unesdoc.unesco.org/ark:/48223/pf0000391105
- Miao, F., & Cukurova, M. (2024). AI competency framework for teachers. UNESCO. doi.org/10.54675/ZJTE2084 — unesdoc.unesco.org/ark:/48223/pf0000391104
- UNESCO. (2023). Guidance for generative AI in education and research. unesdoc.unesco.org/ark:/48223/pf0000386693
- UNESCO. (2025). AI and the future of education: disruptions, dilemmas and directions. unesdoc.unesco.org/ark:/48223/pf0000395236
- British Educational Research Association. (2024). Ethical guidelines for educational research (5th ed.). <https://www.bera.ac.uk/wp-content/uploads/2024/04/Ethical-Guidelines-for-Educational-Research-5th-edition.pdf>
- Council of Europe. (2024). Chounta, I.-A. et al., The state of artificial intelligence and education across Europe. [rm.coe.int/the-state-of-artificial-intelligence-and-education-across-europe-the-1680b29929](https://rm.coe.int/the-state-of-artificial-intelligence-and-education-across-europe-the-1680b29929)
- OECD. (2026). OECD Digital Education Outlook 2026: Exploring effective uses of generative AI in education. doi.org/10.1787/062a7394-en
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, OJ L, 2024/1689, 12.7.2024. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- European Commission, Directorate-General for Education, Youth, Sport and Culture. (2026). *Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators*. Publications Office of the European Union. doi.org/10.2766/7967834 (catalogue no. NC-01-25-186-EN-N)
- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, OJ L 119, 4.5.2016, p. 1–88. <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>

## University AI governance practice

- Stanford University, Office of Community Standards. Generative AI Policy Guidance. [communitystandards.stanford.edu/generative-ai-policy-guidance](https://communitystandards.stanford.edu/generative-ai-policy-guidance)
- Stanford University, Teaching Commons. Policy & Guidance. [teachingcommons.stanford.edu/policy-guidance](https://teachingcommons.stanford.edu/policy-guidance)
- University of Oxford. Generative AI at Oxford. [ox.ac.uk/gen-ai](https://ox.ac.uk/gen-ai)
- University of Oxford. Policy for using generative AI in research. [ox.ac.uk/research/support-researchers/research-practice/policy-generative-ai-research](https://ox.ac.uk/research/support-researchers/research-practice/policy-generative-ai-research)
- University of Cambridge, Blended Learning Service. AI Policy Framework. [blendedlearning.cam.ac.uk/artificial-intelligence-and-education/ai-policy-framework](https://blendedlearning.cam.ac.uk/artificial-intelligence-and-education/ai-policy-framework)
- KU Leuven. Responsible use of generative Artificial Intelligence. [kuleuven.be/english/education/student/educational-tools/responsible-use-of-generative-artificial-intelligence](https://kuleuven.be/english/education/student/educational-tools/responsible-use-of-generative-artificial-intelligence)
- Elsevier. The use of AI and AI-assisted technologies in writing for Elsevier. [elsevier.com/about/policies-and-standards/the-use-of-generative-ai-and-ai-assisted-technologies-in-writing-for-elsevier](https://elsevier.com/about/policies-and-standards/the-use-of-generative-ai-and-ai-assisted-technologies-in-writing-for-elsevier)
- Oxford University Press. Author use of Artificial Intelligence. [academic.oup.com/pages/for-authors/books/author-use-of-artificial-intelligence](https://academic.oup.com/pages/for-authors/books/author-use-of-artificial-intelligence)
- Cambridge University Press & Assessment. Cambridge launches AI research ethics policy. [cambridge.org/news-and-insights/cambridge-launches-ai-research-ethics-policy](https://cambridge.org/news-and-insights/cambridge-launches-ai-research-ethics-policy)

Cambridge University Press & Assessment. Authorship and contributorship – research publishing ethics guidelines. [cambridge.org/core/services/authors/publishing-ethics/research-publishing-ethics-guidelines-for-journals/authorship-and-contributorship](https://www.cambridge.org/core/services/authors/publishing-ethics/research-publishing-ethics-guidelines-for-journals/authorship-and-contributorship)

Peter Lang. AI Statement. [peterlang.com/ai-statement](https://peterlang.com/ai-statement)

## Research on AI disclosure and use

Kwon, D. (2025). Is it OK for AI to write science papers? *Nature* survey shows researchers are split. *Nature*, 641(8063), 574–578. doi.org/10.1038/d41586-025-01463-8

Kobak, D., González-Márquez, R., Horvát, E.Á., & Lause, J. (2025). Delving into ChatGPT usage in academic writing through excess vocabulary. *Science Advances*, 11, eadt3813.

## ELITE-AI project sources

Nechifor, A. (2026a). ELITE-AI: Overview of the Project. Presentation, Stakeholder Workshop for Policy Makers, 17 June 2026.

Nechifor, A. (2026b). Ethics in AI: What Journals, Classrooms and Observation Sheets Are Telling Us. Presentation, ELITE-AI Workshop, 7 May 2026.

Costantini, F. (2026). Innovation in Language Acquisition through AI. Presentation, Stakeholder Workshop for Policy Makers, 17 June 2026.

ELITE-AI Consortium. (2025a). *Manual of Good Practices for Researchers*, Work Package 2. Activity 4. [https://elite-ai.unitbv.ro/wp-content/uploads/2026/04/ELITE\\_AI\\_Manual\\_Good\\_Practices.pdf](https://elite-ai.unitbv.ro/wp-content/uploads/2026/04/ELITE_AI_Manual_Good_Practices.pdf)

ELITE-AI Consortium. (2025b). *Integrating AI Support in ESP Classes: A Manual of Good Practices for Teachers and Students*. Work Package 3. Activity 4. [https://elite-ai.unitbv.ro/wp-content/uploads/2026/06/ELITE-AI\\_Manual\\_Good\\_Practices\\_Teachers\\_Students.pdf](https://elite-ai.unitbv.ro/wp-content/uploads/2026/06/ELITE-AI_Manual_Good_Practices_Teachers_Students.pdf)

ELITE-AI Consortium. (2026a). *Quality Assurance Manual*, Work Package 4, Activity 1. [https://elite-ai.unitbv.ro/wp-content/uploads/2026/06/ELITE\\_AI\\_Quality-Assurance\\_MANUAL\\_WP4\\_A1.pdf](https://elite-ai.unitbv.ro/wp-content/uploads/2026/06/ELITE_AI_Quality-Assurance_MANUAL_WP4_A1.pdf)

ELITE-AI Consortium. (2026b). *Stakeholder Workshop for Policy Makers Evaluation Questionnaire Responses*, 17 June 2026

ELITE-AI Consortium. (2026c). *Stakeholder Workshop for Policy Makers Slido Poll Results*, 17 June 2026.