

ELITE-AI

Integrating AI Support into Educational Policy and Practice

Manual of Good Practices for Policy Makers

Policy Development · Stakeholder Engagement · Resource Allocation · Evaluation and Monitoring · Professional Development

Informed by the ELITE-AI Stakeholder Workshop for Policy Makers, 17 June 2026

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

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

elite-ai.unitbv.ro

Table of Contents

Table of Contents	2
About This Manual	3
1. Why Policy Makers Need This Manual: The Governance Gap	3
2. Policy Development	5
2.1 What a policy actually needs to cover.....	5
2.2 Learning from national and international models	6
2.3 A phased, living-document approach	6
2.4 Anchoring policy in evidence, not enthusiasm or suspicion	7
3. Stakeholder Engagement	8
3.1 Who the stakeholders are	8
3.2 What ELITE-AI's own workshop revealed about engagement failure	8
3.3 Engagement mechanisms that work	9
4. Resource Allocation	10
4.1 The uneven distribution of AI capability	10
4.2 Procurement and compliance models.....	11
4.3 A resource-allocation checklist.....	11
5. Evaluation and Monitoring	11
5.1 What to monitor	11
5.2 Models from elsewhere	12
5.3 The ELITE-AI observation-sheet model as a monitoring instrument.....	12
6. Professional Development	13
6.1 Why training is the precondition, not an add-on	13
6.2 What the evidence shows about what actually drives adoption.....	13
6.3 A professional development agenda.....	13
7. The ELITE-AI Evidence Base	14
7.1 Four countries, six months, eight teachers.....	14
7.2 The Policy Makers' Workshop, 17 June 2026	15
7.3 The monograph and the tool repository.....	15
8. A Consolidated Policy Makers' Checklist	15
9. Conclusion and Next Steps	16
References	17
International and comparative policy frameworks	17
National and institutional examples.....	18
Research on AI adoption, engagement and professional development.....	18
ELITE-AI project sources	18

About This Manual

This manual closes the set of five good-practice manuals produced by the ELITE-AI project (*Empowering Specialized Language Acquisition with Integrated AI*) an Erasmus+ KA220-HED Cooperation Partnership coordinated by Transilvania University of Braşov (UNITBV, Romania) in partnership with University of Maribor (Slovenia), University of Udine (Italy) and Juraj Dobrila University of Pula (Croatia). While the *Quality Assurance Manual* (ELITE-AI Consortium, 2026c), the *Manual of Good Practices for Teachers and Students* (ELITE-AI Consortium, 2025d), and the *Manual of Good Practices for Researchers* (ELITE-AI Consortium, 2025c) speak to people working inside the classroom or the research design, this manual, alongside the *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d) is addressed to the people who set the conditions those classrooms and research projects operate under: ministry officials, rectors and vice-rectors, quality-assurance and legal offices, and department heads and programme coordinators who approve or write AI-related rules.

Its intention, as set out in the project proposal, is to give policy makers a comprehensive, practical guide to integrating AI support into educational settings, organised around five components: policy development, stakeholder engagement, resource allocation, evaluation and monitoring, and professional development. Each of those five components gets its own section below, and each draws on two bodies of evidence at once – what ELITE-AI itself found across four countries, two years of research, six months of classroom trials, and a dedicated workshop for policy makers, and what comparable institutions, national governments and international bodies are doing elsewhere, verified and referenced so that every claim in this manual can be checked and followed up.

The manual draws heavily on the rest of the ELITE-AI outputs. Thus, where a point is argued in full elsewhere in the project’s manuals – the ethical case for a given provision, for instance, or the detailed classroom findings behind a recommendation – this manual cross-references rather than repeats and gives the exact document and section to consult.

Who should use this manual

- Ministry-level officials responsible for national or regional digital-education strategy.
- Rectors, vice-rectors and institutional leadership teams setting institution-wide AI rules.
- Quality-assurance, legal and compliance offices reviewing AI-related policy for regulatory fit.
- Department heads and programme coordinators who approve AI use within specific curricula.

1. Why Policy Makers Need This Manual: The Governance Gap

Although adoption varies between nations, institutions, and user groups, generative AI usage in higher education has become omnipresent and is growing quickly. According to an OECD evaluation conducted in 2025–2026, several national settings had exceptionally high student use, such as around 95% of undergraduates in the United Kingdom and approximately 90% of students in Germany, with academic staff adopting it more cautiously, but still substantially (OECD, 2026a). However, governance has not kept pace with that curve, as the same review cites a 2025 UNESCO survey in which only 19% of higher education institutions worldwide reported having a formal AI policy in place,

with a further 42% still drafting one – leaving a majority of institutions, and the students and staff within them, to rely on personal judgement and unmanaged consumer tools (OECD, 2026a).

A 2026 report from a UK teaching fellow provides a more nuanced account of the same gap and does so from inside a single classroom rather than from a cross-national survey. More than half of students, 53%, report using generative AI for non-academic purposes – drafting emails to employers, decoding accommodation contracts, preparing for job interviews – while fewer than one in five feel their university has given them any ethical guidance on using it at all. On the staff side of the same gap, 41% of UK and European academics report never having used generative AI in any form, and 29% have actively blocked access to it within their own virtual learning environments (Kamel, 2026). The consequence, documented in a separate study Kamel cites, is a climate of mutual suspicion rather than open practice: 62% of students felt their lecturers viewed them with automatic suspicion around AI use, and 57% admitted hiding AI use they believed was legitimate, simply to avoid that suspicion. Students and staff, in other words, are not on opposite sides of a values disagreement – a comparative study she cites found both groups share the same underlying worry, that AI might erode authentic learning, with students in fact more cautious about over-reliance than staff assumed. The gap is one of visibility and trust, not of values, which is precisely why Section 3 below treats stakeholder engagement as a distinct governance problem rather than a subheading under policy development.

ELITE-AI's own stakeholder workshop for policy makers (ELITE-AI Consortium, 2026e), held online on 17 June 2026, found the same gap at ground level and in participants' own words. Asked what stood between AI and educators at their institution, the single most-endorsed open-text answer was a lack of AI literacy; one participant described their institution as having 'no regulations at all', with teachers 'unsupported and abandoned, left to our own devices'; another reported that a policy did exist but had never been disseminated beyond the committee that wrote it. The full detail of that workshop is set out in the companion *Ethical Provisions Manual*, Section 8 (ELITE-AI Consortium, 2026d); this manual treats its findings as the starting evidence for what follows.

This manual does not propose a single model policy – no such thing would transfer cleanly across a ministry, a comprehensive university and a language department, and every comparative example examined for this manual and for the *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d) before it was explicitly provisional, revised on a cycle measured in months rather than years. What it offers instead is a structured way to close the governance gap in a given institution's own context, organised around the five components the project brief specifies.

The five components of this manual

Policy development – what a working AI policy needs to cover, and how to build one without starting from a blank page (Section 2).

Stakeholder engagement – who needs a voice in the policy and why consultation failures are the most common governance failure ELITE-AI observed (Section 3).

Resource allocation – funding tool access and training equitably across disciplines, not just at a central level (Section 4).

Evaluation and monitoring – treating policy as a living instrument with a fixed review cycle, not a one-off release (Section 5).

Professional development – training as the precondition for enforcement, not an afterthought to it (Section 6).

2. Policy Development

2.1 What a policy actually needs to cover

A recent roadmap for higher-education AI policy, developed by a multi-institutional European team and published in 2025, distils the design task into four priorities that this manual adopts as its own starting checklist: a clear regulatory framework tying institutional rules to binding law; guidance specific to each stakeholder group rather than a single undifferentiated statement; an anchoring in what the research on AI in education actually shows, rather than intuition or vendor claims; and a sustained commitment to AI literacy as the condition that makes the rest of the policy usable (Stracke et al., 2025). ELITE-AI's own *Ethical Provisions Manual*, in its Section 12 (ELITE-AI Consortium, 2026d) arrives at a compatible answer from a different direction, working from workshop evidence and comparative practice to a nine-part checklist covering data privacy, informed consent, bias and fairness, transparency and accountability, ethical use of student data, responsible AI development, pedagogical integrity, equity and inclusivity, and continuous evaluation. A policy that addresses both – the four structural priorities and the nine substantive areas – covers the ground this manual and its companion volume between them treat as necessary.

What a policy should not attempt is comprehensiveness at the expense of usability, the clearest single lesson from the comparative practice reviewed for this manual and for the *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d), in Section 9, being that authority decentralises well and rigidity does not: Stanford requires course-level policies from individual instructors rather than a single institutional rule; Cambridge deliberately declined to write one university-wide policy on the grounds that a mathematics degree subject programme and a modern languages course need different answers; KU Leuven's rule is a two-line requirement – instructors state whether GenAI is permitted for an assignment, students disclose when they use it – rather than an exhaustive rulebook. A policy document that tries to anticipate every tool and every classroom scenario centrally is a policy that will be obsolete before it is approved.

Stanford's own CS336 course, *Language Modeling from Scratch*, offers a genuinely controlled example of what a course-level policy looks like once an instructor writes one properly, rather than as a single line in a syllabus. Its policy does not simply say whether AI is allowed; it defines a role – AI tools are positioned explicitly as a teaching aid, not a solution generator – and then operationalises that role with matched lists of what an AI agent should and should not do. It should explain concepts, point to course materials, review code a student has already written, and ask guiding questions; it should not write code, complete assignment sections, or convert a requirement directly into a working solution. The policy even scripts the boundary with worked dialogue examples, showing an acceptable exchange (the agent asks what the student has already tried and offers a specific, non-paste-able diagnostic step) against an unacceptable one (the agent simply supplies the finished code). This is precisely the stakeholder-specific, task-specific guidance that Stracke et al. (2025) identify as a design priority in the abstract; CS336 shows what it looks like once someone actually writes it (Stanford University, CS336 course staff, 2026).

A functional example: Stanford CS336's AI-use policy

Role definition – AI agents are framed explicitly as a teaching aid, not a solution generator – stated once, at the top of the document, before any rule.

Matched should/should-not lists – each permitted use (explain, point to materials, review, ask guiding questions) is paired with an adjacent prohibited use (write code, complete sections, convert requirements into solutions).

Operational dialogue examples – the policy shows, not just tells – a sample acceptable exchange next to a sample unacceptable one, removing ambiguity about where the line sits.

A stated escalation route – ‘when in doubt, refer the student to course staff or office hours’ – the policy does not try to resolve every edge case itself.

2.2 Learning from national and international models

An OECD policy review published in July 2026 surveyed member and partner countries and grouped the responses it found into five recurring types, drawn from real, named national and institutional examples rather than abstracted principle (OECD, 2026b). Governments are issuing system-level guidance on responsible use, as Ireland, Australia and the United Kingdom have each done. They are developing common approaches to compliance and procurement, most visibly where Italy negotiated a national agreement with OpenAI and Australia established a national partnership with Anthropic, converting what would otherwise be dozens of separate institutional contracts into a single negotiated, auditable arrangement. They are building competence at national scale through funded training programmes, as Germany, Switzerland and South Korea have each done for their higher-education staff. They are supporting evidence collection and pilot programmes before wider adoption and broader implementation, as in *Australia’s GenAI Knowledge Hub* and the *Netherlands’ eduGenAI pilots*, while a smaller number are funding the development of education-specific GenAI tools rather than relying solely on general-purpose consumer products, as ETH Zurich has done with its *Ethel teaching assistant* (OECD, 2026b).

These five response types map closely onto the five components this manual is organised around, and the correspondence is not an accident: guidance on responsible use and purchasing are both policy-development questions (this section); competence-building is professional development (Section 6); pilots and evidence collection are evaluation and monitoring (Section 5); and the underlying tension in every example – who is penalised and for what, and who has the capacity to decide – runs through resource allocation and stakeholder engagement (Sections 3 and 4). A policy maker drafting a first institutional policy does not need to invent a framework from nothing, considering that the template of one already exists across a dozen national systems and can be adapted rather than authored from scratch.

2.3 A phased, living-document approach

Every policy examined for this manual and its companion volume is explicitly provisional. Cambridge revised its *AI Policy Framework* twice within a year of first publishing it. Italy’s national coordination of universities, as noted by a university leader at an ELITE-AI workshop, has highlighted the need for shared policies to remain responsive to the rapid pace of technological change and evolving national, regional, and EU regulatory developments, as reported in the *Ethical Provisions Manual*, Section 8 (ELITE-AI Consortium, 2026d). Jisc, the UK’s tertiary-education digital agency, published its own *Strategic Framework for AI in Colleges and Universities* in August 2025, built around three pillars – skills and knowledge, technology, and governance – paired with a self-assessment instrument, its *AI Education Maturity Model* (2025), that lets an institution benchmark its own readiness against the sector rather than assume compliance from a single policy document (Jisc, 2025). The UK’s

Department for Education maintains a separate, narrower policy paper on generative AI in education, first published in 2023 and last updated in August 2025 (UK Department for Education, 2025).

A maturity-model approach of this kind is worth adopting directly rather than reinventing: it treats policy development as a staged process – awareness, then guidance, then integration, then optimisation – with each stage assessed against concrete institutional evidence rather than against the mere existence of a document. An institution with an approved but undisseminated policy, the failure mode reported at ELITE-AI’s own workshop, would score honestly low on such an instrument despite technically having ‘a policy’, which is precisely the gap this manual is written to close.

2.4 Anchoring policy in evidence, not enthusiasm or suspicion

UNESCO’s own foundational guidance for this audience – *AI and Education: Guidance for Policy-Makers*, published in 2021 and still the starting reference point five years on – opens from a position this manual shares: that policy adjusts to what AI actually does in a classroom, not to what a vendor or a headline claims it does (Miao et al., 2021). That distinction has only sharpened since, with multiple 2025 and 2026 meta-analyses claiming that generative AI enhances academic performance, while a critical review of those same meta-analyses found that each failed basic standards of study selection, methodological rigour, or statistical integrity. The evidence therefore warrants a much more cautious conclusion: rather than demonstrating a positive effect on learning, it currently provides no reliable basis for claiming one (Holmes, 2026). Two further studies push the caution further: the first one is a 2024 trial in Türkiye which found that student performance improved only while ChatGPT access was available and fell below the no-access control group once that access was withdrawn – a pattern its authors describe as a cognitive support rather than a learning gain; the second one is a 2025 MIT study using EEG to record participants’ brain activity, which found generative-AI users performed worse than a non-AI control group at every level measured, a pattern reported as ‘cognitive atrophy’ (Holmes, 2026).

None of this argues against AI in education; it argues against treating adoption itself as the policy goal. Philosopher Gert Biesta’s distinction, cited in the same UNESCO Courier analysis, is useful here: education serves three functions – qualification (acquiring knowledge and skills), socialisation (finding one’s place in a social and civic world) and subjectification (becoming a person capable of independent judgement) – and generative AI, however capable, addresses only the first (Holmes, 2026). A policy that measures its own success purely by adoption and satisfaction metrics is measuring qualification alone while silently dropping the other two. Section 5 below returns to this point directly: what a policy monitors has to include evidence of learning, not only evidence of use.

What ‘anchored in evidence’ means in practice

- Treat vendor and headline claims about learning gains as marketing, not evidence, until an independently reviewed study says otherwise.
- Ask what function of education a given AI use serves – qualification, socialisation, or subjectification – and do not expect it to serve the latter two.
- Watch for the ‘cognitive support’ pattern directly: does performance hold up once AI access is removed, or only while it is present?
- Consult UNESCO’s *Guidance for Policy-Makers* (2021) and its 2023 *Guidance for Generative AI* alongside this manual, as both were written for exactly this audience.

Policy development: essentials

1. Anchor the policy to mandatory law first – the AI Act, GDPR, and national data-protection statute – then add institution-specific provisions (*Ethical Provisions Manual*, Section 2).
2. Write stakeholder-specific guidance rather than one undifferentiated statement: students, teaching staff, researchers and administrators face different questions.
3. Decentralise course- and discipline-level decisions to instructors within a common institutional structure, rather than legislating every scenario centrally.
4. Set a fixed review interval – no longer than one academic term – and assess policy maturity against evidence, not against the document's mere existence.
5. Consult the comparative examples in Section 2.2 and in the *Ethical Provisions Manual*, Section 9, before drafting from scratch.
6. Anchor claims about learning benefit in independently reviewed evidence, not vendor or headline claims (Section 2.4).

3. Stakeholder Engagement

3.1 Who the stakeholders are

When ELITE-AI's own workshop asked policy makers, in an open word-cloud poll, who should be responsible for creating and enforcing AI governance policy in education, the answers converged on a layered structure rather than a single owner: government and ministry, university and rectoral team, institutional policy makers, teachers, students, and AI experts all appeared repeatedly among the 39 responses, mentioned by the *Ethical Provisions Manual*, in Section 8.1 (ELITE-AI Consortium, 2026d). That result is itself the argument for this section: no single stakeholder group can or should own AI policy alone, and a policy drafted by one group in isolation – a committee, an IT department, a rector's office – without input from the others is a policy at high risk of the dissemination and legitimacy failures described below.

The University of Edinburgh's own guidance illustrates one workable model of this multi-stakeholder principle in practice. Its published student guidance on working with generative AI, first issued in October 2024 with a fixed review date attached, requires a specific declaration wherever GenAI is used in assessed work – naming the tool, its publisher, the URL, and a one-sentence description of how it was used – while leaving the substantive judgement about what counts as acceptable use to individual course organisers, who are expected to state their own position clearly rather than default to a single university-wide rule (University of Edinburgh, 2024). The structure matters as much as the specific declaration format: a policy that is centrally defined but locally applied avoids both extremes this manual has warned against – a rule disconnected from classroom reality, and a vacuum where individual staff are left to guess.

3.2 What ELITE-AI's own workshop revealed about engagement failure

Three distinct engagement failures surfaced independently from policy makers at ELITE-AI's 17 June workshop, all mentioned by the *Ethical Provisions Manual*, in Section 8.6 (ELITE-AI Consortium, 2026d) and each is instructive precisely because it is a different failure, not a repetition of the same one. One institution had no AI-specific policy at all, leaving staff to treat the question as one of individual judgement, with one participant describing the prevailing attitude toward AI as adversarial – 'it's like the enemy' – rather than governed. A second institution had an approved policy covering

students, researchers and teachers that had simply never been disseminated beyond the committee that drafted it, so staff outside that committee remained unaware of its content and continued to exercise individual discretion regardless of what the document said. A third context, described by a participant teaching within a system where institutions have moved to mandate rather than debate AI use, illustrated the opposite failure mode: adoption imposed without the staff training to support it, prompting the closing observation from the floor that ‘we need to do something, but we are not skilled’.

None of these three failures is solved by writing a better-worded policy document. The first needs a policy; the second needs a dissemination plan as deliberate as the drafting process itself; the third needs training sequenced before, not after, implementation. Stakeholder engagement, in other words, is not a single step in policy development but a distinct and recurring discipline that has to be revisited at drafting, at execution, and at every subsequent revision.

Hence, ELITE-AI’s own workshop is not an isolated data point here. A first-person account from a UK teaching fellow, published the same year, describes arriving at the identical failure from the opposite side of the lecture theatre: having announced a blanket ban on AI use, she was asked by a student why that meant deleting the CV she had just rewritten with ChatGPT the night before for a job application – a moment the author describes as revealing that staff and students were, in her words, ‘on different planets’ (Kamel, 2026). The stakes of getting engagement wrong are not abstract: in the same account, most students who used AI legitimately still hid that use from staff, out of a reasonable fear of being treated with suspicion regardless of what they actually did. A policy that is not co-developed with the people it governs does not just fail to be followed, but it actively teaches the people it governs to conceal their practice from it, which makes the policy blind to exactly the behaviour it is aimed to shape.

3.3 Engagement mechanisms that work

The OECD’s own synthesis of effective GenAI use in education establishes a direct dialogue with the people affected by a policy, explicitly naming teachers, students, families and unions among its core recommendations for policy makers, alongside ensuring human oversight remains available as an alternative to any AI-mediated process (OECD, 2026a). ELITE-AI’s own classroom evidence mirrored in Section 7 of the *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d) offers a complementary, bottom-up argument for the same conclusion: across all four partner institutions, ethical boundaries and disclosure norms emerged from teachers’ and students’ own classroom practice before any formal institutional policy existed to require them. A policy that consults frontline staff and students before drafting is not just simply more democratic, but it is also more likely to describe practice that already works, rather than legislating against it.

One concrete mechanism is worth naming because it is inexpensive and directly reported to work: co-writing the rules with the people who will follow them. In the same UK account cited above, an instructor handed her class a blank shared document and asked them to draft their own three rules for AI use in the course; the group produced sensible, specific norms – cite any AI-generated text, do not work with any unrevised copy-pasted text, ask the instructor when unsure – and then enforced those rules on each other as peer expectation rather than imposed policy (Kamel, 2026). This is the classroom-level version of the standing-committee principle above: authorship by the people governed by a rule is not a courtesy extended to them, but what makes the rule durable once the drafting meeting ends.

Stakeholder engagement: essentials

1. Establish a standing, cross-functional committee – academic leadership, IT, legal, research ethics, and faculty – rather than a single office owning policy alone.
2. Consult frontline teaching staff and students before drafting; classroom practice often already contains the norms a policy needs to formalise.
3. Treat dissemination as a distinct task from drafting, with its own plan and its own success measure – an unread policy governs nothing.
4. Where adoption is commanded rather than optional, sequence staff training before the implementation takes effect, not after.

4. Resource Allocation

4.1 The uneven distribution of AI capability

Resource allocation in AI policy is usually discussed in terms of funding for AI licences. ELITE-AI's classroom evidence suggests, however, that the more consequential inequality concerns not simply who has access to AI tools, but which disciplines and specialised language registers those tools are actually equipped to support. The consortium's ESP fields – law, tourism, criminal justice terminology, Erasmus+ mobility, education pedagogy – varied significantly in the specialist AI tools available to them even within the same ESP domain, let alone institution, and a single institution-wide licence for a general-purpose tool does not close that gap, according to Section 7 of the *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d) because general-purpose tools are precisely the ones that perform least reliably on specialised professional registers. Resource allocation decisions, in other words, have to be made at the level of the discipline and its actual register, not only at the level of the institution's overall AI budget.

The same structural pattern reappears at a larger scale internationally. The phenomenon of 'globalisation' Nechifor put forward in 2025 addresses exactly these issues, i.e., both the positives and the negatives of a digital educational environment in a globalised world. Research on AI adoption in under-resourced higher-education systems, including recent work on the AI divide across sub-Saharan African institutions, documents how funding, connectivity and licensing constraints deepen and reinforce existing inequality rather than simply lagging better-resourced systems on the same trajectory (Malatji, 2026). The general principle transfers even where the specific context does not: a resource-allocation policy that assumes uniform access, whether across disciplines within one university or across institutions within one country, will systematically disadvantage whichever group it did not have in mind when it was written.

There is a sharper version of this same inequality worth naming directly, sometimes called the 'flipped digital divide': students from better-resourced backgrounds continue to have access to human teachers and small-group instruction, while less advantaged students are increasingly offered AI tools as the substitute (Holmes, 2026). A resource-allocation policy that treats an AI licence as an adequate answer to a staffing or funding deficit risks establishing exactly this pattern rather than closing it. AI access is not a substitute for the human contact a well-resourced student still receives; budgeted correctly, it should add to institutional capacity, not quietly replace it for the students who can least afford the substitution.

4.2 Procurement and compliance models

Institutions and governments reviewed for this manual sit along a spectrum of purchasing models, from fully centralised to fully devolved, and each has different resource-allocation consequences. At the national end, Italy's agreement with OpenAI and Australia's partnership with Anthropic convert what would otherwise be dozens of separate institutional negotiations into a single audited national arrangement, with the efficiency benefits and the loss of institution-specific flexibility that centralisation typically brings (OECD, 2026b). At the institutional end, as referred to by the *Ethical Provisions Manual* in Section 9.2 (ELITE-AI Consortium, 2026d), Oxford's provision of free, single-sign-on access to ChatGPT Edu for every member of staff and student, paired with mandatory information-security training, exemplifies a middle path: institution-wide licensing that still leaves room for departmental judgement about which specialised tools to add on top. Neither end of the spectrum is inherently correct; the choice has to reflect the size, disciplinary spread and existing IT capacity of the institution or system making it.

4.3 A resource-allocation checklist

Whatever the acquisition model, ELITE-AI's own evidence and the comparative literature converge on the same non-negotiable approach: funding for training has to be treated as seriously as funding for licences. A tool that is licensed institution-wide but poorly understood by the staff meant to use it produces the same practical disadvantage as no licence at all, a point several policy makers made independently at ELITE-AI's own workshop, as captured by the *Ethical Provisions Manual*, in Section 4 (ELITE-AI Consortium, 2026d). Any resource-allocation line that funds tool access without a matching line for the training needed to use that access well has, in practice, only allocated half the resource the policy intended.

Resource allocation: essentials

1. Budget by discipline and register, not only by institution – a general-purpose licence does not close a specialised-tool gap.
2. Choose a procurement model – national, institutional, or departmental – that matches actual IT capacity and disciplinary spread, not the model that is simplest to announce.
3. Fund AI-literacy training in the same budget line as tool licences, not as a residual or optional add-on.
4. Provide an offline or low-cost alternative wherever reliable access cannot be guaranteed equally to every student or member of staff.

5. Evaluation and Monitoring

5.1 What to monitor

Monitoring an AI policy is often reduced, in practice, to tracking tool-adoption numbers: how many staff logged into the licensed platform, how many students used it. Nevertheless, ELITE-AI's own evidence argues for a wider brief, as rendered by the *Ethical Provisions Manual*, in Sections 3-6 (ELITE-AI Consortium, 2026d), considering that what mattered across the consortium's own classroom trials was not usage volume but four qualitative signals: whether disclosure norms were actually being followed, whether assessment practice had adapted to keep pace with AI-assisted work, whether access remained equitable across disciplines and student groups, and whether a workable

route existed for a student or teacher to flag a biased or unfair AI output once a tool was in use. A monitoring framework built only around usage statistics will miss all four, and, per Section 2.4 above, it should also monitor for evidence of learning itself, not only for evidence of use – the two are not the same thing and can move in opposite directions.

One low-cost assessment practice illustrates what monitoring the second signal can look like at the level of a single course, i.e., whether assessment has actually adapted to the requirements of today's digital educational landscape. One instructor now asks students to submit their full AI conversation history alongside a final piece of written work, not as a disciplinary trap but as a window onto how the output was actually produced and revised (Kamel, 2026). The submitted history becomes evidence of process rather than only of product, which is exactly the shift Section 6 of this manual's companion *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d) argues assessment needs to make, and it costs an institution nothing to require beyond the willingness to read a transcript instead of, or alongside, a polished final draft.

5.2 Models from elsewhere

Two distinct monitoring traditions are worth drawing on, at different levels of formality. At the more formal end, the United States' National Institute of Standards and Technology publishes an *AI Risk Management Framework*, with a dedicated profile for generative AI, organised around four functions – govern, map, measure and manage – that together provide a rigorous, auditable structure for tracking AI risk over the lifetime of a deployment rather than only at the point of initial approval (NIST, 2023; NIST, 2024). At a lighter and more sector-specific level, Jisc's own *AI Education Maturity Model* allows an institution to self-assess its readiness across the same broad domains – skills, technology, governance – named in its companion *Strategic Framework for AI in Colleges and Universities*, without requiring the full apparatus of a formal risk-management programme (Jisc, 2025).

Evidence-collection infrastructure is a third, complementary model. Australia's *GenAI Knowledge Hub* and a set of pilot programmes run under the name *eduGenAI* in the Netherlands both function as centralised, ongoing evidence-gathering mechanisms that feed back into national policy revision rather than unique compliance audits (OECD, 2026b). The common thread across all three models, formal and informal, is the same one this manual has returned to throughout: monitoring is treated as continuous, not as a single certification event, and its findings are expected to change the policy rather than simply confirm it.

5.3 The ELITE-AI observation-sheet model as a monitoring instrument

ELITE-AI's own methodology offers policy makers a monitoring instrument that is already built, already piloted, and considerably lighter than a formal risk-management programme. The consortium's observation-sheet protocol – structured, cross-institutional and repeatable – generated 1,870 individual responses across four countries during Work Package 3: 1,522 in-class student self-assessments, 221 at-home student self-assessments, and 127 teacher observation sheets (ELITE-AI Consortium, 2025d, 2026a). That dataset was collected without disrupting normal teaching, using instruments teachers themselves helped design, which is precisely the kind of low-burden, high-fidelity monitoring a resource-constrained institution can realistically sustain over multiple semesters rather than abandoning after a single pilot cycle.

Evaluation and monitoring: essentials

1. Monitor disclosure, assessment adaptation, equity of access, and pathways for reporting and resolving problems, not only usage statistics.
2. Adopt or adapt an existing framework – NIST’s govern/map/measure/manage functions for formal risk management, or Jisc’s maturity model for a lighter self-assessment – rather than building one from scratch.
3. Treat monitoring as continuous, with findings expected to revise the policy, not merely confirm it at a single audit point.
4. Consider the ELITE-AI observation-sheet protocol as a directly adaptable, low-burden model for either domestic or even cross-institutional monitoring.

6. Professional Development

6.1 Why training is the precondition, not an add-on

Every section of this manual arrives, from a different direction, at the same conclusion: policy fails first at the point where the people expected to follow it have not been shown how. At ELITE-AI’s own workshop, 86 per cent of policy makers said AI ethics training should be mandatory for all staff who use AI tools, and the single most-endorsed answer to a separate question about barriers was a lack of AI literacy – not a lack of rules, a lack of skill to follow the rules that already existed, according to the *Ethical Provisions Manual*, Sections 8.3–8.4 (ELITE-AI Consortium, 2026d). One participant working under a system that had moved to mandate AI use rather than debate it put the same point unambiguously from the other direction: ‘we need to do something, but we are not skilled.’ Recent research on AI literacy among educators reaches an aligned conclusion from outside the project: literacy has to be built deliberately and early, because it does not develop as a side effect of exposure to the tools themselves (Pei et al., 2026).

6.2 What the evidence shows about what actually drives adoption

A 2023 study of AI adoption factors among teachers, co-authored by one of the researchers who later helped draft UNESCO’s own AI competency framework for teachers, found that engagement depended less on mandates than on teachers’ perceived usefulness of a given tool for their own specific teaching context (Cukurova et al., 2023). A 2026 meta-analysis of faculty acceptance studies across the Technology Acceptance Model (TAM) and Unified Theory of Acceptance and Use of Technology (UTAUT) literatures, pooling data from over three thousand faculty respondents, reached a compatible finding: ‘perceived usefulness’ and ‘ease of use’, according to TAM, or ‘performance expectancy’ and ‘effort expectancy’, according to UTAUT (Venkatesh et al., 2003), were consistently the strongest predictors of both attitude toward AI tools and intention to adopt them, more so than institutional command or social pressure from colleagues (Alotaibi, 2026). Read together, the practical implication for policy makers is that professional development framed around demonstrated classroom usefulness – showing staff a tool solving a problem they already have – will outperform development framed only around compliance with a new rule.

6.3 A professional development agenda

The national cases discussed in Section 2.2 demonstrate how competence-building can be funded and implemented systematically at scale: Germany, Switzerland and South Korea have each run

structured training programmes for higher-education staff as a distinct policy instrument in its own right, not folded into a general IT-onboarding process (OECD, 2026b). ELITE-AI's own recommendation, developed from its classroom trials and repeated across this manual and its companion *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d), is more modest in scale but identical in sequencing logic: training has to come before enforcement, and dissemination of what staff are trained on has to be treated as seriously as the training content itself.

Not every institution can fund a national-scale programme, and the evidence suggests it may not need to. The UK teaching fellow's account of closing her own institution's staff-student trust gap begins with what she calls admitting to being a beginner: a recurring, informal monthly gathering where colleagues bring their AI failures rather than their successes, run without the ceremony or expense of a formal programme, which she reports did more to change actual classroom practice than any written policy (Kamel, 2026). The underlying finding, that staff who report not knowing how a tool works or fearing they will look unskilled using it are the same staff most likely to either block access outright or adopt it without any structure at all, argues for treating psychological safety as a legitimate professional-development design goal, not a soft add-on to the technical curriculum, i.e., permission to be visibly still learning.

Professional development: a minimum agenda

1. Develop a short, mandatory 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. Base training on real teaching and research needs and demonstrate how AI can address specific teaching and research tasks that staff already face, rather than focusing primarily on compliance – perceived usefulness is a stronger driver of genuine adoption than mandate.
3. Assign institutional responsibility for keeping the training current, through a designated contact or committee and a fixed review interval of no more than one academic term.
4. Deliver training before any AI mandate takes effect, rather than after it – and communicate the training programme as deliberately and systematically as the policy itself.
5. Provide low-cost, informal, no-shame opportunities for continued learning – recurring peer meet-ups can complement, rather than replace, funded formal training.

7. The ELITE-AI Evidence Base

This manual's recommendations are not written from policy theory alone. They are grounded in three bodies of evidence the consortium generated directly, briefly summarised here with pointers to where each is documented in full.

7.1 Four countries, six months, eight teachers

Work Package 3 ran AI-integrated teaching trials across pragmatics, applied grammar, ESP, criminal-justice terminology, tourism English and early-years pedagogy, taught by eight teachers across the four partner institutions over six months. As already mentioned above, the trials generated 1,870 structured observation-sheet responses – 1,522 in-class student self-assessments, 221 at-home student self-assessments and 127 teacher observation sheets (ELITE-AI Consortium, 2026a) – and catalogued more than one hundred AI tools in active classroom use across the four language skills (ELITE-AI, 2025a). The full pedagogical findings from this dataset are reported in the *Quality Assurance Manual* (ELITE-AI Consortium, 2026c) and the *Ethical Provisions Manual* (ELITE-AI

Consortium, 2026d), consequently, this manual draws on their ethical and governance implications only.

7.2 The Policy Makers' Workshop, 17 June 2026

Around 95 participants from at least a dozen institutions and countries took part in ELITE-AI's dedicated online workshop for policy makers, contributing live poll responses, open-text comments and floor discussion. Table 1 summarises the headline results, the *Ethical Provisions Manual*, Section 8 (ELITE-AI Consortium, 2026d) giving the full detail, including the floor discussion this manual draws on throughout.

Poll	Headline finding
Who should govern AI in education?	A layered structure: ministry, university, faculty, individual instructor – not a single owner.
Most ethically concerning use of AI	Undisclosed AI-generated student text, at 98% of respondents – far ahead of any other option.
Barriers between AI and educators	Lack of AI literacy was the single most-endorsed answer, ahead of lack of regulation or funding.
Should ethics training be mandatory?	86% yes, 3% no, 11% 'it depends'.
Post-workshop questionnaire (ELITE-AI Consortium, 2026f). (n=32)	31 of 32 respondents want continued access to project resources; 26 of 32 said the workshop gave concrete policy ideas.

7.3 The monograph and the tool repository

The consortium's classroom and literature-review work is also published in full-length academic form: Blažević, Costantini, Dimulescu, Gallitelli, Gugić, Kompara Lukančič, Li Destri, Nechifor and Smajla, *Empowering Language Acquisition through Artificial Intelligence: A Comprehensive Monograph* (Cambridge Scholars Publishing, 2026), which situates the project's classroom findings within the wider international literature on AI in language education and catalogues the full tool repository referenced throughout this manual (ELITE-AI Consortium, 2025b) and its companion volumes. Individual project findings have also been published separately, the complete set of project outputs, including further published articles (ELITE-AI Consortium, 2026b), being maintained at elite-ai.unitbv.ro (ELITE-AI Consortium, 2025a).

8. A Consolidated Policy Makers' Checklist

This section gathers the essentials from each of the five preceding sections into a single working checklist, organised in the order a policy maker is likely to need them: first develop the policy, then engage the people it affects, then fund it properly, then monitor it, then keep training current.

1. Policy development

- Anchor the policy to binding law first, then add institution-specific provisions.
- Write stakeholder-specific guidance; decentralise course-level decisions to instructors.

- Set a fixed review interval of no more than one academic term.

2. Stakeholder engagement

- Establish a standing, cross-functional committee spanning academic leadership, IT, legal and faculty.
- Consult frontline staff and students before drafting; classroom practice often already contains the needed norms.
- Treat dissemination as a distinct, resourced task, separate from drafting.

3. Resource allocation

- Budget by discipline and specialised register, not only by institution.
- Fund AI-literacy training in the same line as tool licences.
- Guarantee an offline or low-cost alternative wherever access cannot be equal.

4. Evaluation and monitoring

- Monitor disclosure, assessment adaptation, equity and compensation – not only usage statistics.
- Adopt an existing framework (NIST's functions, or a lighter maturity model) rather than building one from scratch.
- Treat findings as grounds to revise the policy, not merely to confirm it.

5. Professional development

- Make AI-literacy training mandatory, and sequence it before any mandate takes effect.
- Frame training around demonstrated usefulness for real tasks, not compliance alone.
- Name an institutional owner for keeping training current and disseminated.

9. Conclusion and Next Steps

This manual completes the set of four good-practice manuals ELITE-AI set out to produce: for researchers, for teachers and students, and – across this volume and its companions the *Ethical Provisions Manual* (ELITE-AI Consortium, 2026d) and the *Quality Assurance Manual* (ELITE-AI Consortium, 2026c) – for policy makers. Its recommendations are deliberately practical rather than aspirational, and deliberately plural rather than prescriptive: no single national or institutional model reviewed here was treated as the correct one, because none of the institutions or governments whose practice this manual draws on treats its own policy as finished either. What transfers across every example, from Stanford to KU Leuven to Italy's national university coordination, is not a specific rule but a shared discipline – decentralise what should be decentralised, consult before mandating, fund training as seriously as access, monitor continuously, and revise on a fixed and honest schedule.

ELITE-AI's own contribution to that discipline is the evidence assembled across this manual: six months of classroom trials in four countries, a workshop that put policy makers from a dozen institutions in direct conversation with the consortium's own findings, and two monographs, an interactive lesson plan, and tool repository that will remain available after the project itself closes. The project concludes formally in August 2026, with all outputs – the five manuals, the monographs, the observation-sheet datasets and the published articles – maintained at elite-ai.unitbv.ro for institutions that wish to adapt this work to their own context.

References

International and comparative policy frameworks

OECD. (2026a). OECD Digital Education Outlook 2026: Exploring effective uses of generative AI in education. *OECD Publishing*. doi.org/10.1787/062a7394-en

OECD. (2026b). Policies supporting responsible and systematic GenAI adoption in higher education. *OECD Education Spotlights*, No. 23. OECD Publishing. doi.org/10.1787/c4e5621f-en

Stracke, C. M., Nahar, N., Punzo, V., Massaro, S., Pappa, D., Di Grassi, A., Bećirović, S., Hollins, P., Ziouvelou, X., Skenduli, M. P., & Burgos, D. (2025). Artificial intelligence policies for higher education: Manifesto for critical considerations and a roadmap. *MAP Education and Humanities*, 6, 61–73. doi.org/10.53880/2744-2373.2025.6.61

National Institute of Standards and Technology (NIST). (2023). *Artificial Intelligence Risk Management Framework* (AI RMF 1.0). NIST AI 100-1. doi.org/10.6028/NIST.AI.100-1

National Institute of Standards and Technology (NIST). (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. doi.org/10.6028/NIST.AI.600-1

UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. unesdoc.unesco.org/ark:/48223/pf0000381137_eng

Miao, F., Holmes, W., Huang, R., & Zhang, H. (2021). *AI and education: Guidance for policy-makers*. UNESCO. unesco.org/en/articles/ai-and-education-guidance-policy-makers – unesdoc.unesco.org/ark:/48223/pf0000376709

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

Miao, F., & Shiohira, K. (2024). *AI competency framework for students*. UNESCO. doi.org/10.54675/JKJB9835

Miao, F., & Cukurova, M. (2024). *AI competency framework for teachers*. UNESCO. doi.org/10.54675/ZJTE2084

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. eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation). *Official Journal of the European Union*, OJ L 119, 4.5.2016, p. 1–88. eur-lex.europa.eu/eli/reg/2016/679/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)

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478

National and institutional examples

UK Department for Education. (2025). *Generative artificial intelligence (AI) in education* [Policy paper]. First published 29 March 2023, last updated 12 August 2025. gov.uk/government/publications/generative-artificial-intelligence-in-education

Jisc, National Centre for AI. (2025). *A Strategic Framework for AI in Colleges and Universities, and the accompanying AI Education Maturity Model*. nationalcentreforai.jiscinvolve.org/wp/2025/08/28/a-strategic-framework-for-ai-in-colleges-and-universities/

University of Edinburgh, Registry Services. (2024). *Guidance for working with Generative AI ('GenAI') in your studies*. registryservices.ed.ac.uk/sites/default/files/2024-10/Generative%20AI%20Guidance%20for%20Students%20October%202024.pdf

Stanford University, Office of Community Standards. *Generative AI Policy Guidance*. communitystandards.stanford.edu/generative-ai-policy-guidance

Stanford University, CS336 course staff. (2026). *CS336: Language Modeling from Scratch – course policies and AI agent guidelines*. cs336.stanford.edu – github.com/stanford-cs336/assignment1-basics/blob/main/CLAUDE.md

University of Oxford. *Generative AI at Oxford*. ox.ac.uk/gen-ai

University of Cambridge. *Blended Learning Service. AI Policy Framework*. 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

Research on AI adoption, engagement and professional development

Holmes, W. (2026). Learning to think in the AI era. *The UNESCO Courier*, 2 April 2026 (updated 7 April 2026). courier.unesco.org/en/articles/learning-think-ai-era

Kamel, D. (2026). When it comes to AI, faculty are from Mars and students are from Venus. *Times Higher Education*, Campus, 1 June 2026. timeshighereducation.com/campus/when-it-comes-ai-faculty-are-mars-and-students-are-venus

Cukurova, M., Miao, X., & Brooker, R. (2023). Adoption of artificial intelligence in schools: Unveiling factors influencing teachers' engagement. In N. Wang, G. Rebolledo-Mendez, N. Matsuda, O. C. Santos, & V. Dimitrova (Eds.), *Artificial Intelligence in Education (AIED 2023), Lecture Notes in Computer Science*, vol. 13916 (pp. 151–163). Springer. doi.org/10.1007/978-3-031-36272-9_13

Alotaibi, N. (2026). Faculty acceptance of generative AI in higher education: A meta-analysis of TAM and UTAUT studies (2021–2025). *International Journal of Higher Education*, 15(1). sciedupress.com/journal/index.php/ijhe/article/view/28869

Pei, B., Lu, J., Zhang, Z., Tuffour, P., & Park, S. (2026). Enhancing AI literacy for educators: Where to start and to what end? *Computers and Education: Artificial Intelligence*, 10, 100550. doi.org/10.1016/j.caeai.2026.100550

Ateeq, A., Alzoraiki, M., Milhem, M., & Ateeq, R. A. (2024). Artificial intelligence in education: Implications for academic integrity and the shift toward holistic assessment. *Frontiers in Education*, 9, 1470979. doi.org/10.3389/feduc.2024.1470979

Malatji, M. (2026). Bridging the AI divide in sub-Saharan Africa: Challenges and opportunities for inclusivity. arxiv.org/abs/2601.06145

ELITE-AI project sources

Blažević, I., Costantini, F., Dimulescu, C., Gallitelli, E., Gugić, J., Kompara Lukančič, M., Li Destri, E., Nechifor, A., & Smajla, T. (2026). *Empowering language acquisition through artificial intelligence: A comprehensive monograph*. Cambridge Scholars Publishing.

Nechifor, A. (2025a). The 'globalised' discourse of education: Positives and negatives of globalising language acquisition through digitalisation. *Philobiblon. Transylvanian Journal of Multidisciplinary Research in Humanities*, 30(1), 229–247. doi.org/10.26424/philobib.2025.30.1.15

Nechifor, A. (2025b). AI-enhanced pragmatic language acquisition in ESP contexts: Connecting grammar development and communicative competence – A comprehensive framework. *Analele Universității „Ovidius” Constanța. Seria Filologie*, 36(2), 1269–1299.

ELITE-AI Consortium (2025a). Project homepage and results overview. elite-ai.unitbv.ro

ELITE-AI Consortium. (2025b). *AI Tools Repository*. elite-ai.unitbv.ro/wp-content/uploads/2025/03/AI-Tools-Repository_30.03.pdf

ELITE-AI Consortium. (2025c). *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

ELITE-AI Consortium. (2025d). *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

ELITE-AI Consortium. (2026a). *Observation Sheets Database*. Work Package 3, Activities 1, 3 and 4. <https://elite-ai.unitbv.ro/index.php/elite-ai-observation-sheets/>

ELITE-AI Consortium. (2026b). *Research Articles*. Work Package 3. <https://elite-ai.unitbv.ro/index.php/research-articles-elite-ai/>

ELITE-AI Consortium. (2026c). *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

ELITE-AI Consortium. (2026d). *Ethical Provisions Manual*, Work Package 4, Activity 2. https://elite-ai.unitbv.ro/wp-content/uploads/2026/08/ELITE_AI_Ethical_Provisions_Manual.pdf

ELITE-AI Consortium. (2026e). *Stakeholder Workshop for Policy Makers*, Work Package 4, Activity 2. 17 June 2026. https://elite-ai.unitbv.ro/wp-content/uploads/2026/06/Call_for_participation_Workshop_Elite_AI_June_17_Final.pdf

ELITE-AI Consortium. (2026f). *Workshop Evaluation Questionnaire Responses*, 17 June 2026