

# BUILDING AI-READY UNIVERSITIES

## EVIDENCE-BASED PRIORITIES FOR THE NEXT FIVE YEARS

AI is no longer an experimental question for universities: diffusion has outpaced governance, informal practice is hardening into institutional structure, and the gap between them is where risk accumulates. This white paper sets out the state of the art as of March 2026, drawing on more than 130 sources that include systematic and tertiary reviews, international policy instruments, and sector surveys covering thousands of institutions. Table A summarises the six findings on which this argument rests and shows where each is developed in the sections that follow.

The ABC Framework for AI-Ready Universities was developed through thematic analysis of the parallel sessions of the **AI in Higher Education Summit** (Tascioglu, 2026) held at ESCP Business School, Paris campus, on the 17<sup>th</sup> and 18<sup>th</sup> of March 2026. The ABC pillars draw on ESCP's mission and vision to educate "Accountable, Bold & Creative" leaders capable of driving human-centred transformations in business and beyond. The framework is used to provide further structure to ESCPTech Institute's state-of-the-art report on AI in higher education. Specifically, Table B in the Recommendations for senior leadership section operationalises the framework through three mandates, six institutional priorities for 2026-2030, and a four phase roadmap.

# TABLE A

## Key Findings

FINDING	EVIDENCE	ABC PILLAR
AI DIFFUSION IS STRUCTURAL, BUT GOVERNANCE MATURITY IS NOT	<b>Student AI use reaches 92% in the UK</b> (Freeman, 2025), and 86 per cent globally (Stanford AI Index, 2025). Faculty adoption of AI is at 61 per cent but <b>88% report only minimal to moderate integration</b> (DEC, 2025).	A B
NEAR-TERM BENEFITS ARE CREDIBLE, BUT INSTITUTION-SCALE EFFECTS REMAIN CONDITIONAL	<b>Productivity, accessibility, and formative-support gains are documented across synthesis studies in teaching, research, and administration.</b>	B C
THE SAME SYSTEMS THAT PRODUCE GAINS INTRODUCE PERSISTENT RISKS	Four concurrent risk clusters emerge: assessment validity, epistemic quality, analytics accuracy, privacy, data security and vendor opacity.	A
EQUITY IS AN UPSTREAM DESIGN CONSTRAINT, NOT A DOWNSTREAM FIX	Access gaps appear through infrastructure, language coverage, disability accessibility, and paywalled premium tiers. Surveillance and detection errors disproportionately affect already-marginalised groups.	A C
THE EU AI ACT MAKES INSTITUTIONAL ACTION MANDATORY, NOT OPTIONAL	39 per cent of institutions have an AI-related acceptable-use policy (Stanford AI Index, 2025). 80 per cent of faculty do not find institutional guidelines comprehensive. <b>6% find AI literacy resources comprehensive (DEC, 2025).</b>	A
ASSESSMENT AND INTEGRITY CONSTITUTE THE PRIMARY DESIGN FRONTIER	Bans and detection-dependent regimes cannot preserve credential meaning once AI-mediated workflows become routine.	C

These findings unfold across the sections that follow. The diffusion, benefits, and risks sections build the diagnostic case (Findings 1–3). The equity and EU AI Act sections frame the institutional response (Findings 4–5). The assessment frontier carries the argument into the recommendations and four-phase roadmap that close the paper (Finding 6).

# AI DIFFUSION IS STRUCTURAL, BUT GOVERNANCE MATURITY IS NOT

AI capability evolves through a layered trajectory that runs from analytics to automation and agentic workflow integration, through to generative systems participating in open-ended, language-mediated tasks. The layers coexist and do not replace one another: learning analytics and early-warning systems persist, routine task automation is widespread, and conversational systems diffuse rapidly as tutors, writing partners, and coding assistants. Across teaching, research, and administration, value relocates from classification and process execution into dialogue under uncertainty, where models function as collaborators in drafting, iteration, critique, and explanation. Risk shifts correspondingly from isolated tool errors to the design and governance of workflows that use generated text, code, or recommendations as inputs to consequential decisions.

**Adoption is now broad: student use reaches 92 per cent in the United Kingdom (Freeman, 2025) and faculty adoption stands at 61 per cent globally, yet 88 per cent of adopting faculty report only minimal to moderate integration (DEC, 2025).**

Informal practice therefore becomes structural before pedagogy, assessment, data governance, and accountability are institutionalised. Meanwhile the operative environment is shaped by bottom-up toolchains, embedded platform features, peer learning, and inconsistent local norms, so students encounter contradictory rules between courses and campuses. UK 2025 evidence captures the resulting student picture: 80 per cent report a clear institutional policy on AI in assessment, and **76% believe their institution can detect AI-generated work**. This is up from 65 per cent the previous year (Freeman, 2025), even as detector reliability remains structurally unstable for multilingual learners and students with disabilities, those most exposed to false-positive risk.

*Near-term benefits are credible, but institution-scale effects remain conditional.*

Generative AI and adjacent tools deliver credible short-run benefits around productivity, accessibility, and individualised support, yet institution-scale effects remain contingent on governance and design conditions that are still forming.

In teaching and learning, defensible near-term gains concentrate in formative support, personalised pacing and explanation, and accessibility for multilingual and disabled learners. Durable learning effects depend on whether AI outputs are treated as inputs to human critique and verification.

In research, benefits concentrate in workflow support (literature triage, synthesis, outlining, drafting, editing, translation, presentation refinement), while leaving core scholarly judgement as the governing constraint. Bottlenecks shift toward verification, provenance, and responsibility boundaries. Effects on research quality and novelty remain uncertain, with productivity gains better understood as reallocations of effort toward checking, verification, and judgement.

In administration and student services, benefits cluster around process support and first-line conversational interfaces: drafting communications, handling routine queries, service navigation, and automation of structured tasks where data infrastructures exist.

These gains remain inseparable from their enabling conditions. Integrated data systems, clear escalation pathways, and purpose limitation matter once conversational support and analytics expand monitoring scope. Legal exposure and legitimacy risks intensify in high-impact contexts.

**The same systems that produce gains introduce persistent risks**

**Across the evidence base, gains and risks appear as joint outcomes that intensify as stakes rise. Four concurrent risk clusters operate in parallel.**

- **The first cluster concerns assessment validity.** Generative AI undermines the evidential status of many unsupervised outputs, and detector-centred integrity regimes are structurally unstable once any permitted assistance exists. Errors and false accusations are unevenly distributed, disproportionately affecting multilingual learners and students with disabilities. Integrity shifts from misconduct detection toward the redesign of validity conditions, supported by transparent norms for acceptable use, disclosure, and documentation that preserve due process and legitimacy.
- **The second cluster concerns epistemic quality and skill formation.** Generative outputs can be fluent yet inaccurate, biased, or misaligned with local rubrics and disciplinary standards. This fluency induces over-trust and automation bias. The issue is not an isolated error but the gradual weakening of verification practices and foundational skills if calibrated delegation and disciplined checking are not taught and assessed, with beginners particularly exposed due to limited epistemic vigilance and domain priors.
- **The third cluster concerns analytics and predictive decision support.** Accuracy in early warning, engagement detection, and classroom analytics can be weak, sometimes near chance. Here the governing threshold is consequentiality: low-stakes recommendations may tolerate error under human oversight, but high-stakes interventions and eligibility-related decisions require very high performance, auditable reasoning, and contestability. This logic extends to student services triage, progression interventions, admissions-related decision support, and staff-facing systems.
- **The fourth cluster concerns privacy, data security, and vendor opacity.** Educational AI relies on sensitive data integrated through cloud services and vendor networks, where logging, retention, and secondary use can be opaque. Safety depends on institutional clarity regarding data use, access, and storage of prompts and outputs, as well as on technical and contractual safeguards. These vulnerabilities traverse functions, making cross-domain safety requirements necessary. Otherwise, opacity and weak data governance produce correlated vulnerabilities across teaching tools, research platforms, and administrative systems that domain-specific guidance cannot contain.

*Equity is an upstream design constraint, not a downstream fix.*

Language coverage, disability accessibility, uneven model performance, and paywalled premium tiers produce access gaps that stratify capability within the same programme. In multilingual institutions, governance choices about language coverage and local validation determine whether AI tools reduce or widen learning divides.

*Deploying tools without parallel investment in interpretive capacity transfers advantage instead of removing it, and equity belongs inside procurement criteria and assessment design.*

**Three inclusive design principles translate this stance into institutional practice.**

- **First**, accessibility by design is the binding decision test: each AI deployment is a distributional intervention, and the central judgement is whether deployment narrows or widens divides, with language coverage, cultural pluralism, and disability accessibility operating as admissibility conditions. Enforcement that disregards assistive contexts chills participation and increases false accusations for multilingual and disabled students.
- **Second**, participatory governance functions as a safety infrastructure: student support, accessibility services, and affected user groups hold defined roles and influence in decision-making. Otherwise, harms remain systematically invisible and governance becomes performative, and high-stakes uses remain disfavoured until differential error and harm are empirically bounded. Third, literacy and capability building close the loop: expanding provision without parallel competence development reproduces the stratification it intends to reduce. Hence, institutional provisioning is coupled with universal AI literacy and structured onboarding tied to policy and pedagogy, while licensing eligibility, default tool availability, and training requirements determine whether access translates into equitable capability or amplifies prior advantage.

**Two specifics sharpen this constraint for multi-campus institutions:**

- Generative systems perform best in English, with reliability declining materially in other languages, so normalising AI use without multilingual validation transfers advantage instead of removing it. (UNESCO, 2023c)
- Cultural classification of misconduct also varies systematically: **respondents in high uncertainty-avoidance contexts are 3.67 times more likely to classify generative AI use as cheating**, and respondents in long-term-orientation contexts 2.87 times more likely (Yusuf et al., 2024). A single integrity rule applied identically across campuses results in predictably uneven enforcement and student exposure to disciplinary risk.

## **THE EU AI ACT MAKES INSTITUTIONAL ACTION MANDATORY, NOT OPTIONAL**

Institutional action is now warranted on regulatory grounds. The costs of fragmented adoption extend beyond performance, and the binding issue becomes responsibility under rights-based expectations and regulation. Universities are educators and major employers, placing them directly inside the EU AI Act's high-risk framing for education and employment-related uses. This includes corresponding obligations for documentation, oversight, accountability, and AI literacy measures for persons operating or using these systems.

*Under these conditions, AI is not a discretionary innovation agenda delegated to local pilots. It is a governance, compliance, equity, and legitimacy agenda that determines whether adoption yields coherent educational value or correlates with institutional risk.*

Annex III of the EU AI Act designates admissions, assignment to institutions, learning-outcome assessment, level-setting, and proctoring as presumptively high-risk education uses, alongside recruitment, promotion, termination, task allocation, monitoring, and evaluation as high-risk employment uses (Article 6, Annex III). Article 5 of the EU AI Act sets bright-line prohibitions.

Three of these directly affect universities:

### EMOTIONAL INFERENCE

in workplaces and educational institutions

(with narrow exceptions for medical or safety purposes)

### SOCIAL SCORING

that affects rights or access

### SYSTEMS

that exploit age, disability, or socioeconomic vulnerabilities

Engagement surveillance, well-being inference, and composite cross-context scoring proposals therefore become categorically impermissible regardless of pedagogical claim. Commission guidelines further specifying the Annex III classification boundaries were due in February 2026, so institutions must treat classification as a living control process with documentation that is resilient to rule refinement (European Parliament and Council, 2024).

High-risk classification can also catch student-services automation that universities do not initially read as regulated. Additional high-risk provisions apply where AI is used to determine eligibility for scholarships, bursaries, or subsidised services. It also affects situations where deployments link to creditworthiness or insurance-related offerings on behalf of public functions. Efficiency initiatives in admissions, finance, and student support can therefore become compliance-intensive.

Current evidence supports risk-based sequencing, institutional provisioning and literacy, assessment redesign, and continuous monitoring, while remaining insufficient to justify high-stakes deployment without local validation and lifecycle governance. Attention therefore shifts from tool debates to implementation architecture. This entails rights-based foundations, risk classification, procurement and data governance, equity by design, maturity building, and change management that aligns norms and accountability across heterogeneous disciplines and campuses.

# ASSESSMENT AND INTEGRITY CONSTITUTE THE PRIMARY DESIGN FRONTIER

Benyayer & Tascioglu (2026) applied the ABC Framework to synthesise and introduce the parallel session contributions of the proceedings. The framework translates into three pillar-level commitments that persist through cycles of technological change. Each pillar below sets out its mandate and the two 2026-2030 priorities that carry it. The staged roadmap is presented separately in the Executive roadmap section.

Three benchmarks calibrate where most institutions sit when this paper goes to press. Echoing Table A, in 2025 only 39 per cent of higher education institutions held an AI acceptable use policy, with a strong size gradient favouring larger universities (Stanford AI Index, 2025). Faculty preference is conditional permission: 57 per cent favour permitting AI with disclosure and specific instructions, 38 per cent with disclosure alone, 23 per cent support bans, and 11 per cent would mandate use (DEC, 2025).

*Yet 83 per cent of faculty report concern that students cannot evaluate AI outputs (DEC, 2025), so the literacy mandate sits upstream of the assessment-reform mandate, not parallel to it.*

# TABLE B

## Operational architecture of the ABC Framework

PILLAR	MANDATE	PRIORITIES (2026-2030)	ROADMAP CONNECTION
<p><b>A - ACCOUNTABLE</b></p>	<p>Make coherence and accountability non-negotiable: institution-wide baselines for acceptable use, disclosure, and data handling. Accountability allocation and complaint routes for sensitive domains. Governance as continuous revision with protected deliberation spaces.</p>	<ul style="list-style-type: none"> <li>Participatory AI governance balancing academic freedom with responsibility allocation and risk-based regulatory alignment.</li> <li>Ongoing evaluation, technology scanning, and wellbeing-centred design as operating cycles.</li> </ul>	<p><b>Phase 0</b> prerequisites (governance foundations, risk mapping).</p> <p><b>Phase 1</b> interim guidance and baselines.</p> <p><b>Phase 2-3</b> lifecycle-based governance.</p>
<p><b>B - BOLD</b></p>	<p>Treat institutional access and compliance architecture as the primary control surface: supported tools as governance and equity instruments. Risk classification embedded into procurement and oversight. Live AI inventory of systems, data flows, vendors.</p>	<ul style="list-style-type: none"> <li>Institution-wide AI strategy and operating-model consolidation.</li> <li>Equitable access through secure institutional provision.</li> <li>Compulsory AI and data literacy (shared with Creative).</li> </ul>	<p><b>Phase 0</b> maturity diagnosis.</p> <p><b>Phase 1</b> secure access pilots.</p> <p><b>Phase 2</b> scaling and integration.</p> <p><b>Phase 3</b> ubiquitous access with continuous procurement governance.</p>
<p><b>C - CREATIVE</b></p>	<p>Sponsor capability and assessment reform as the integrity strategy: structured staff development as institutional requirement. AI and ethics literacy staged by role. Assessment portfolios reviewed for AI vulnerability with a shift from detection to transparent norms and task redesign.</p>	<ul style="list-style-type: none"> <li>Compulsory AI and data literacy staged and task-tied.</li> <li>Assessment and integrity redesign for AI-saturated conditions, with research integrity standards.</li> </ul>	<p><b>Phase 1</b> baseline literacy and assessment pilots.</p> <p><b>Phase 1</b> assessment redesign pilots scaled to programme level.</p> <p><b>Phase 2-3</b> curriculum redesigned for ubiquitous access procurement governance.</p>

The three pillar commitments below detail the framework's mandates, followed by a four-phase implementation roadmap that runs from prerequisites to AI-by-design integration.

## **ACCOUNTABLE.**

---

The mandate (Table B) holds coherence and accountability as non-negotiable. A small set of institution-wide baselines for acceptable use, disclosure, and data handling anchors coherence, while course and unit rules remain locally contextual within those shared baselines. Rights-based and human-centred principles translate into operational requirements. This entails human oversight and informed disclosure for consequential decisions, accountability allocation across roles, complaint and redress routes in sensitive domains, and protected deliberation spaces that allow governance to revise itself as evidence accumulates.

## **BOLD.**

---

Institutional access and compliance architecture is the primary control surface (Table B). Institutionally supported tools function as both governance and equity instruments wherever AI is integrated into learning or work. Risk classification becomes a design constraint for the operating model, embedded into procurement, deployment, and oversight, so that prohibited practices can be excluded and high-risk domains can be governed through heightened controls. Meanwhile, a live AI inventory of systems, data flows, and vendors provides meaningful oversight at scale.

## **CREATIVE.**

---

Capability and assessment reform are the integrity strategy (Table B). Structured staff development becomes an institutional requirement when responsibility for interpretation and verification spreads beyond specialist roles, and AI and ethics literacy is staged through role-differentiated pathways that integrate technical understanding, critical evaluation, ethical reasoning, and domain practice for faculty, researchers, and professional staff.

*Assessment reform moves from detection-centred regimes toward task redesign, process-visible assessment, transparent norms for permitted use, and oral or interactive components that sustain credential validity under routine AI use.*

Implementation capacity is a staged bundle of enablers. Stalled progress reflects absence of predictable access, training, shared practice, and permission structures, not absence of model capability. **DEC reports that faculty identify five dominant enablers: access to tools and resources (65 per cent), faculty training in AI literacy and skills (64 per cent), best-practice libraries and concrete use cases (60 per cent), clear guidelines on AI in teaching (50 per cent), and environments that encourage innovation while tolerating failure (31 per cent).** The four phases below sequence these enablers over time, with pillar tags showing which pillars engage in each phase.

### **Phase 0. Before you start [A/B].**

Governance and data foundations become first dependencies. The problem is defined, constraints are mapped, use cases are assigned to risk categories (with early exclusion where appropriate), and a maturity diagnosis establishes readiness across rules, access, familiarity, trust, and organisational culture. Integrity and research governance prerequisites are specified before tool normalisation, and early governance choices are tested against socio-economic and cultural constraints.

### **Phase 1. First 12 months [A/B/C].**

Prerequisites become practice through bounded pilots organised around explicit evaluation questions. Secure access to selected tools is provisioned, baseline literacy and ethics training reduce uncertainty, and interim guidance limits contradictory local norms between units. AI Act compliance runs as a parallel workstream for high-impact pilots, and assessment reform opens as the core integrity pilot stream through vulnerability audits, redesign pilots, and explicit disclosure expectations.

## **Phase 2. Years 1 to 3 [A/B/C].**

What survives evaluation is integrated into routine governance, and pilot evidence translates into programme-level outcomes and curriculum governance artefacts. Monitoring expands to learning outcomes and unintended consequences, sociotechnical stewardship roles are created and resourced for oversight and data governance, and evaluation is professionalised as routine institutional capability, not exceptional reporting.

## **Phase 2. Years 1 to 3 [A/B/C].**

AI becomes a standing design parameter across teaching, research, operations, and governance, running continuously throughout the lifecycle. Curriculum and assessment are redesigned for ubiquitous access, centring human judgement, responsible collaboration, transparency, and critical evaluation. Continuous impact monitoring sustains scepticism and protects institutional purpose amid evolving capabilities and regulations.

An AI-ready university is characterised by three capacities: to standardise what requires coherence (the Accountable commitment), to make institutional access and compliance architecture the primary control surface (the Bold commitment), and to redesign what makes credentials and scholarship meaningful under AI saturation (the Creative commitment).



IT ALL STARTS HERE

BERLIN | LONDON | MADRID | PARIS | TURIN | WARSAW

**BERLIN**

Heubnerweg 8 – 10  
14059 Berlin, Deutschland

**LONDON**

527 Finchley Road,  
London NW3 7BG, United Kingdom

**MADRID**

Arroyofresno 1  
28035 Madrid, España  
Maria de Molina 4  
28006 Madrid, España  
Navalmanzano 6  
28035 Madrid, España

**PARIS**

6-8, av. de la Porte de Champerret  
75017 Paris, France  
3, rue Armand Moisant  
75015 Paris, France

**TURIN**

Via Andrea Doria 27,  
10123, Torino, Italia

**WARSAW**

c/o Kozminski University,  
International Relation Office,  
57/59 Jagiellońska St.  
03-301 Warsaw, Poland

**BRANCH CAMPUS DUBAI**

DIFC Academy Dubai International Financial Center, Dubai, UAE

ESCP benefits from the best **international accreditations**.



Each of our European and international campuses benefits from **national recognition** in higher education.

Affiliated to

