

Responsible AI in Higher Education: A Practical Framework for Institutional Readiness



Purpose

Advancing Human Potential Through Responsible AI

Universities have a long-standing mission to expand knowledge, develop talent, advance research, and contribute to societal progress. Artificial intelligence should be viewed as a strategic enabler of that mission—not as an end in itself. Responsible adoption begins with establishing a shared institutional purpose for AI. This purpose can guide governance decisions, technology investments, workforce development, education strategies, implementation priorities, and the institution's approach to managing evolving legal, regulatory, and societal expectations. This paper lays the groundwork for responsible AI adoption in higher education, offering a practical starting point for institutions as they navigate the opportunities, challenges, and evolving landscape of AI.

Intended Audience and How to Use This Framework

This framework is designed for university leaders responsible for establishing governance structures, guiding AI adoption, and preparing their institutions for an AI-enabled future. Intended audiences include:

- University Presidents responsible for institutional strategy, mission alignment, and stakeholder trust.
- Provosts and Chief Academic Officers responsible for academic innovation, teaching and learning, research, and faculty engagement.
- Chief Information Officers responsible for technology strategy, AI governance, risk management, and institutional AI capabilities.
- Deans and Academic Leaders responsible for implementing AI-related priorities within schools, colleges, and academic programs.
- Cross-functional AI Governance Committees responsible for overseeing AI adoption and coordinating institutional decision-making.

How to Use This Framework

Every institution's AI journey is different. Rather than prescribing a single governance model, this framework is intended to help leadership teams assess their current state, identify gaps, and develop an approach that aligns with institutional priorities and culture.



Executive Leadership Perspective: Responsible AI as an Institutional Imperative

Artificial intelligence is reshaping nearly every aspect of higher education, from teaching and learning to research, student services, workforce development, and institutional operations. As universities seek to capture AI's benefits, executive leaders face a critical challenge: ensuring AI is deployed in ways that advance institutional mission, uphold stakeholder trust, and support long-term resilience.

Responsible AI governance is no longer a technical issue delegated solely to information technology teams or AI specialists. It is a leadership responsibility that requires executive sponsorship, organizational authority, sustained investment, and institution-wide coordination. Universities that treat AI governance as an unfunded administrative assignment risk fragmented oversight, inconsistent practices, and diminished trust among students, faculty, staff, regulators, and other stakeholders.

At the same time, institutions should avoid viewing governance solely through the lens of risk management. Effective governance enables innovation by creating the confidence, accountability, and structures necessary to scale AI responsibly. Across industries, organizations are grappling with questions of accountability, oversight, procurement, transparency, and risk management. Higher education has an opportunity to lead by demonstrating how innovation and responsibility can advance together.

Questions for Executive Leaders

As AI adoption accelerates, university presidents, provosts, deans, CIOs and CAOs should consider:

- How does AI support our institutional mission, strategic priorities, and student success objectives?
- Who is accountable for AI governance, and do they have the authority, resources, and executive sponsorship necessary to succeed?
- Do we have visibility into the AI systems being developed, procured, and used across the institution?
- How are we evaluating AI systems throughout their lifecycle, from procurement and implementation to monitoring and retirement?
- Are faculty, staff, students, and researchers equipped with the knowledge and skills needed to use AI responsibly?
- How are we adapting to evolving legal, regulatory, ethical, and societal expectations?

Governance Requires More Than Accountability

Successful AI governance depends on more than assigning ownership. Institutions should establish governance structures supported by executive authority, dedicated funding, cross-functional participation, and clear decision-making processes. Governance programs should provide visibility into institutional AI activities, establish mechanisms for risk assessment and oversight, and ensure accountability for outcomes.

Without the necessary authority and resources, governance teams may struggle to influence institutional practices, resulting in inconsistent adoption and ineffective oversight.

Preparing Future Leaders for an AI-Enabled World

Higher education plays a unique role in shaping how future leaders understand and govern AI. The objective is not simply to teach students how to use AI tools. Institutions have an opportunity to prepare graduates to design, evaluate, govern, procure, and oversee AI systems across business, government, healthcare, nonprofit, and academic settings.

This responsibility extends beyond computer science and engineering programs. Responsible AI requires perspectives from law, public policy, journalism, business, health sciences, humanities, and other disciplines. Universities that integrate AI literacy, governance, ethics, risk management, and accountability across the curriculum will be better positioned to prepare students for the realities of an AI-enabled economy and society.

A Responsible AI Journey for Higher Education

Stage 1: Build Awareness and AI Literacy

Goal:
Establish a foundational understanding of AI across leadership, faculty, staff, and students.

Examples:

- AI literacy programs
- Faculty development
- Student AI fluency
- Executive education

Stage 2: Experiment and Pilot

Goal:
Test AI applications in low-risk environments.

Examples:

- Classroom pilots
- Student support pilots
- Research pilots
- Administrative efficiency pilots

Stage 3: Determine which technologies should be scaled

Goal:
Establish a foundational understanding of AI across leadership, faculty, staff, and students.

Examples:

- Vendor assessments
- Procurement reviews
- Data and privacy evaluations
- Risk assessments

Stage 4: Establish Governance and Oversight

Goal:
Create structures that enable responsible growth.

Examples:

- Executive AI Steering Committee
- Responsible AI Council
- Risk-based review processes
- AI inventories

Stage 5: Scale and Institutionalize

Goal:
Move from isolated pilots to enterprise-wide adoption.

Examples:

- AI integrated into teaching and learning
- Research support
- Student services
- Operational processes
- Continuous monitoring

Stage 6: Continuous Improvement and Innovation

Goal:
Continuously evaluate, govern, and improve AI usage.

Examples:

- Ongoing monitoring
- Updating governance
- Regulatory alignment
- Innovation pipelines

Human-Centered

Establish a clear institutional position that AI should augment human capabilities, with people retaining responsibility and accountability for its use. This means:

- Students remain responsible for their learning, critical thinking, academic integrity, and validation of AI-assisted work.
- Faculty retain ownership and accountability for teaching, assessment, mentorship, and curriculum decisions.
- Researchers remain responsible for interpreting findings, validating outputs, and ensuring research integrity.
- Academic advisors, student affairs professionals, and student support staff maintain responsibility for advising, intervention, and high-impact decisions affecting learner success.
- Institutional leaders retain accountability for strategic decisions that influence resource allocation, policy, and institutional direction.

AI systems should be designed to augment human capabilities, reduce administrative burden, and create more time for high-value human capabilities. AI investments are most effective when explicitly connected to strategic priorities.

Institutions can evaluate every AI initiative against questions such as:

- What educational, research, operational, or societal outcome is the initiative intended to advance?
- Who is expected to benefit, and how will the institution determine whether those benefits are realized?
- What potential tradeoffs or unintended consequences could affect students, faculty, staff, or other stakeholders?
- How will the initiative strengthen, rather than diminish, human judgment, interaction, and accountability?
- How does the initiative align with institutional mission, values, and strategic priorities?
- What evidence, measures, or feedback will inform decisions to continue, modify, scale, or discontinue the initiative?

Universities can reduce risk by avoiding technology adoption driven solely by market pressure or competitive trends. Every AI initiative should have a clearly defined educational, operational, or societal outcome.



Governance

Building Institutional Trust Through Oversight

Effective governance and automated monitoring promote innovation by providing clarity, accountability, and risk management. Recent IBM Institute for Business Value research reinforces this point: 1 in 4 failed AI initiatives can be traced to weak governance, 27% of AI efficiency gains stem from strong governance, and organizations with higher governance maturity have 18% faster AI adoption rates.¹ Universities that establish governance early are better positioned to scale AI safely and confidently. IBM's Responsible Technology approach demonstrates that governance should act as an enabler rather than a barrier to innovation.

Executive AI Steering Committee

The Steering Committee provides institutional direction and executive accountability for AI adoption.

Responsibilities include:

- Establishing a university-wide vision and strategy for responsible AI.
- Approving major AI investments and enterprise-wide initiatives.
- Defining institutional risk tolerance.
- Aligning AI priorities with strategic planning, accreditation, and compliance requirements.
- Ensuring coordination across academic, administrative, technology, risk, privacy, security, legal, compliance, and governance functions.
- Reviewing progress against institutional AI goals and performance indicators.
- Serving as the final decision-making authority for high-impact AI initiatives.

Responsible AI Council

The Responsible AI Council serves as the primary governance body responsible for translating institutional principles into operational practices.

Responsibilities include:

- Developing responsible AI principles and implementation standards.
- Reviewing high-risk AI use cases before implementation.
- Establishing criteria for transparency, fairness, and accountability.
- Reviewing incidents and emerging AI risks.
- Monitoring evolving regulatory requirements and industry standards.
- Maintaining a university-wide inventory of AI assets and applications.
- Recommending improvements to policies, controls, and governance processes.
- Providing guidance to departments pursuing AI innovation initiatives.
- Establishing monitoring requirements for the use of AI.
- Developing and deploying university-wide training on above.

¹ IBV "Go further, faster with AI"


AI Risk Review Process

Every institutional AI initiative benefits from a structured review process that supports consistent decision-making before, during, and after deployment.

Key components include:

- Identification of the intended outcomes and affected stakeholders.
- Assessment of potential impacts on students, faculty, staff, and external communities.
- Evaluation of privacy, security, and data management considerations.
- Assessment of whether existing cybersecurity policies, controls, and governance processes are sufficient to support the proposed AI application, drawing on established frameworks such as [National Institute of Standards and Technology's \(NIST\) Cyber Security Framework \(CSF\)](#) where appropriate.
- Creation of inventory, documentation, and ownership records to support ongoing visibility, traceability, and accountability across the AI lifecycle.
- Identification of bias and fairness risks.
- Determination of required human oversight and review checkpoints.
- Determination of where automated management and governance of IT assets is needed.
- Definition of monitoring and performance requirements during implementation and deployment.
- Establishment of escalation procedures when concerns arise.

This process helps create consistency while allowing innovation to proceed at scale.



People

Building an AI-Ready Institution

The long-term success of AI depends less on technology and more on institutional capability. Universities must prepare students, faculty, leaders, and staff to engage with AI confidently, critically, and responsibly. Recent global research on higher education identified AI literacy³ and workforce readiness as among the most important priorities for institutions navigating AI adoption³. As AI becomes embedded in everyday work, AI literacy is increasingly the starting point, while AI fluency is becoming a foundational capability. AI literacy is understanding what AI is, how it works, and its capabilities and limitations. AI fluency goes further. It is the ability to apply AI to meaningful work, evaluate outputs critically, and exercise sound judgment when answers are uncertain or incomplete. As AI takes on more routine tasks, capabilities such as critical thinking, problem solving, collaboration, adaptability, and ethical decision-making become even more important. Universities can support the development of both AI literacy and AI fluency through curriculum, experiential learning, professional development, and accessible learning resources such as [IBM SkillsBuild](#), which provides free learning opportunities to help students, educators, and lifelong learners build AI knowledge and practical skills for an AI-enabled future.

Student Capabilities

Institutions can offer foundational AI and data literacy as a general education requirement available to all students, regardless of major. Institutions can create pathways for students across all disciplines, not only technical majors, to participate in building AI systems that reflect the values, standards, and edge cases of their own fields.

As AI becomes increasingly accessible, students are no longer solely consumers of technology. Across disciplines, they have growing opportunities to become builders, creators, and innovators who use AI to solve problems, conduct research, develop new forms of content, and create solutions for real-world challenges. AI fluency extends beyond understanding how AI works. It includes the ability to apply AI effectively, critically evaluate outputs, exercise sound judgment, and responsibly create AI-enabled solutions. Employers increasingly expect graduates to enter the workforce with practical experience applying AI in authentic settings, combining academic knowledge with hands-on experience that was traditionally developed after entering the workplace.

Students should be able to:

- Understand the basic concepts, capabilities, and limitations of AI systems.
- Evaluate AI-generated outputs for reliability, accuracy, and bias.
- Apply critical thinking when using AI tools for research and learning.
- Understand the ethical implications of AI in professional and societal contexts.
- Demonstrate AI fluency by responsibly designing, building, or applying AI-enabled solutions to address academic, professional, or societal challenges.
- Utilize AI responsibly within academic environments.
- Develop interdisciplinary perspectives on AI applications and impacts.
- Demonstrate AI literacy and AI fluency in ways that enhance career readiness and lifelong learning.
- Navigate emerging workplace expectations regarding AI proficiency.

² https://www.ibm.com/think/insights/ai-literacy?ngo-id=0302&utm_source=bett-framework&utm_campaign=bett-framework&utm_medium=other

³ <https://www.digitaleducationcouncil.com/resource-library-items/ai-in-higher-education-global-survey-2026>

In Practice:

Responsible AI and Emerging Technology Education Fordham University (United States)

Fordham University's Advanced Certificate in Ethics and Emerging Technologies provides an interdisciplinary approach to preparing learners for the ethical, governance, and societal challenges associated with emerging technologies. The program illustrates how institutions can help students develop the judgment and decision-making skills needed to evaluate the broader impacts of AI.

[Learn more: Fordham Ethics & Emerging Technologies](#)

In Practice:

AI Workforce Readiness and Experiential Learning AICTE and IBM SkillsBuild (India)

The collaboration between the All India Council for Technical Education (AICTE) and IBM SkillsBuild provides students access to AI-focused learning pathways, digital credentials, virtual internships, and project-based experiences in areas such as artificial intelligence, generative AI, data analytics, and cybersecurity. The initiative illustrates how institutions can complement classroom learning with scalable, industry-aligned experiences that build both technical and responsible AI capabilities.



Faculty Capabilities

As AI becomes increasingly integrated into education, faculty play a critical role in maintaining academic integrity by designing learning experiences, assessments, and classroom expectations that promote authentic learning, responsible AI use, and student accountability without reducing the educational experience to monitoring or enforcement. Institutions benefit from establishing clear guidance, governance, and support structures that promote consistent expectations for AI use across teaching, assessment, research, and student engagement. Faculty development programs should equip educators to apply these institutional principles within their disciplines while adapting to the unique needs of their learners. Faculty require support not only in using AI, but also in leading responsible adoption, modeling effective practices, and helping students navigate an increasingly AI-enabled society and workforce.

Faculty development programs should help educators:

- Redesign assessments for an AI-enabled learning environment.
- Integrate AI effectively into teaching and learning activities.
- Evaluate the pedagogical value and risks of AI systems.
- Foster a culture of academic integrity and responsible AI use that emphasizes learning, accountability, and trust rather than enforcement alone.
- Conduct AI-enabled research responsibly.
- Understand ethical considerations associated with AI applications.
- Guide students in the critical and responsible use of AI.
- Participate in institutional AI governance and policy discussions.

In Practice:

Advancing Interdisciplinary AI Education Cornell University (United States)

Cornell's AI education initiatives provide students and educators opportunities to engage with AI concepts and applications across disciplines. The approach illustrates how institutions can advance AI literacy and critical thinking beyond traditional computer science and engineering programs.

[Learn more: Cornell AI Education](#)

In Practice:

Preparing Faculty for AI-Enabled Teaching and Learning

Andhra Pradesh Faculty Development Initiative (India)

As part of a broader higher education collaboration, faculty development programs are helping educators build knowledge of AI, generative AI, responsible AI practices, and project-based learning methodologies. The initiative demonstrates how institutions can prepare faculty to integrate AI into teaching, assessment, and student mentoring while supporting broader institutional adoption.



Leadership Capabilities

University leaders need a distinct set of competencies.

Senior administrators should be prepared to:

- Understand emerging AI opportunities and risks.
- Make informed decisions about AI investments.
- Lead institutional transformation efforts.
- Establish governance and accountability structures.
- Navigate regulatory and compliance challenges.
- Communicate effectively with stakeholders about AI initiatives.
- Build a culture that balances innovation with responsibility.

In Practice:

Establishing an Institutional AI Strategy Marist University (United States)

Marist University has developed a university-wide AI strategy focused on guiding responsible AI adoption, supporting innovation, and preparing the institution for the opportunities and challenges associated with AI. The initiative illustrates how executive leadership can align AI efforts with broader institutional priorities while creating a foundation for governance and long-term capability building.

[Learn more: Marist AI Strategy](#)

In Practice:

Establishing System-Wide Guidance for Responsible AI State University of New York (SUNY)

SUNY has developed system-wide guidance to support responsible AI adoption across its institutions. The initiative demonstrates how higher education systems can establish common principles, expectations, and governance considerations while allowing individual campuses flexibility to implement AI in ways that align with their local priorities, governance structures, and educational missions.

[Learn more: SUNY Responsible AI Guidance](#)

In Practice:

Establishing an AI Advisory Committee Arizona State University (United States)

Arizona State University established an AI Advisory Committee to provide strategic guidance on the institution's AI activities and opportunities. Composed of faculty and staff from across the university, the committee advises senior leadership on responsible AI adoption, emerging technologies, and long-term institutional priorities. The model illustrates how higher education institutions can create cross-functional governance structures that bring together diverse expertise to guide AI strategy and oversight.

[Learn more: ASU AI Advisory Committee](#)

In Practice:

Leading Institutional Change in Response to AI Massachusetts Institute of Technology (United States)

MIT convened an Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training to assess how AI is reshaping education and to develop recommendations for the Institute. The committee brought together faculty, students, and staff to examine questions related to teaching, assessment, learning outcomes, institutional values, and governance. The resulting report demonstrates how university leadership can create structured processes to evaluate AI's impact and guide institution-wide decision-making (edited)

[Learn more: MIT AI and Education](#)

Risk-Based Review and Oversight

Not all AI applications present the same level of risk. Rather than evaluating AI systems solely based on their functional purpose, a risk-based approach helps universities consider the potential impact of a technology on individuals, communities, institutional operations, and society.

A risk-based framework enables institutions to scale innovation while applying additional governance, oversight, and safeguards to applications that may present greater ethical, operational, legal, reputational, or societal risks.

Establishing Risk Categories

A consistent methodology helps universities assess AI use cases across multiple dimensions, including:

Impact on Stakeholders

Institutions can evaluate whether an AI application may significantly influence decisions, experiences, opportunities, or outcomes for individuals and groups that interact with the institution.

Considerations may include:

- Magnitude and duration of potential impacts.
- Degree of influence on decision-making processes.
- Ability for stakeholders to understand, challenge, or appeal decisions.
- Potential consequences of inaccurate or inappropriate outputs.
- Implications for institutional trust and reputation.

Data Sensitivity

Institutions can assess the nature of the data being used, processed, stored, or generated by AI systems.

Considerations may include:

- Use of personal or sensitive information.
- Student, employee, or research data.
- Privacy obligations and consent requirements.
- Data governance and stewardship requirements.

Fairness and Institutional Impact

Institutions can evaluate whether an AI application could create unintended outcomes, inconsistencies, or barriers that may affect stakeholders, institutional objectives, or trust in decision-making processes.

Considerations may include:

- Potential for unintended bias or skewed outcomes.
- Reliability and consistency of outputs across different populations and use cases.
- Alignment with institutional values, policies, and mission.
- Accessibility and usability considerations.
- Potential effects on stakeholder trust, participation, and engagement.
- Mechanisms for monitoring, remediation, and continuous improvement.

Transparency and Explainability

Institutions should assess whether stakeholders can understand when and how AI is being used.

Considerations may include:

- Clarity of AI-generated outputs.
- Availability of explanations and supporting information.
- Stakeholder awareness of AI involvement.
- Documentation and traceability.

Human Oversight and Accountability

Clear identification of accountability helps clarify who remains responsible for decisions and outcomes involving AI.

Considerations may include:

- Human review requirements.
- Escalation pathways.
- Governance ownership.
- Ongoing monitoring and accountability mechanisms.

Institutional and Societal Impact

Institutions should evaluate broader implications beyond immediate operational outcomes.

Considerations may include:

- Reputational risk.
- Impact on public trust.
- Legal and regulatory considerations.
- Impact on academic values and institutional mission.
- Potential societal consequences.

Applying Enhanced Review

Universities should establish governance processes that apply additional review, monitoring, documentation, and oversight when AI systems present elevated risk across one or more categories. Automation tools can be utilized to track valuable AI assets and their owners, metrics and how those assets are used across the institution.

Depending on an institution's governance model, enhanced review may involve:

- Cross-functional governance committees.
- Legal, privacy, security, and compliance reviews.
- Additional testing and validation.
- Independent bias assessments.
- Expanded transparency requirements.
- Ongoing monitoring and periodic reassessment.

The objective is not to limit innovation, but to ensure that governance efforts are proportionate to risk and aligned with institutional values. Operating this way allows governance to function as an enabler of responsible scale, helping universities innovate confidently while maintaining trust, accountability, and responsible stewardship.

With governance and automation capabilities in place, universities can move faster: they can evaluate new technologies consistently, scale successful AI assets across academic and administrative environments, respond to regulatory or security changes with greater confidence, and build the trust needed for students, faculty, staff, and external partners to adopt AI responsibly. In this way, governance becomes the operating foundation for sustained innovation, helping institutions adopt emerging technologies quickly while preserving human judgment, institutional accountability, and public trust.

Continue the Conversation

As universities translate responsible AI principles into practice, leadership alignment and workforce readiness become critical success factors. IBM's AI for Impact workshops help institutions convene key stakeholders, build AI literacy and AI fluency, explore governance considerations, and develop a shared understanding of ethical AI adoption.

Learn more about IBM AI for Impact workshops and resources: [IBM AI for Impact](#)



About the collaborators



About Bett

At Bett, we bring together educators from across the entire education landscape with flagship events in London, Brasil, USA and Asia. By bringing together educators, leaders, students, policymakers and 600+ innovative EdTech providers, we are the global community for education. Our mission is simple, to drive Better Education, Globally. Let's work together to change the game for education worldwide!

About IBM

IBM is a global leader in hybrid cloud, AI, and consulting, serving clients in more than 175 countries. Built on a foundation of trust, transparency, and innovation, IBM develops technologies that help organizations solve complex challenges. Through its commitment to skilling 30 million people by 2030, IBM partners with nonprofits, higher-ed institutions, and workforce organizations to expand access to free digital skills training.

