

Responsible AI Governance Policy

Model for institutional adoption / Security update 25 September 2026

Complete the adoption record and confirm institutional authority, applicable requirements and implementation arrangements. This model expresses HumanSkills recommendations for adaptation and institutional approval.

Institutional adoption record

Institution: [Institution name]

Approval authority: [Governing body or authorized officer]

Policy owner: [Accountable executive]

Effective date: [Date] Review date: [Date]

Related institutional policies and contact directory: [Policy directory URL]

1 Purpose and commitment

[Institution name] supports the responsible use of artificial intelligence to improve learning, strengthen professional capability, expand access to services and improve the quality of institutional work. AI use must advance the institution's mission and respect the people affected by it.

The institution enables useful AI with safeguards proportionate to the information involved, the people affected and the consequences of error. Routine, low-risk uses may proceed under published standards. Uses involving sensitive information, consequential decisions or actions within institutional systems require the review specified in this policy.

People remain responsible for the work and decisions they undertake with AI. The institution remains responsible for the systems it provides, the conditions under which people use them and the support needed to use them well. Success is demonstrated through better learning, service or work outcomes. Usage volume alone does not establish value.

The policy owner and institutional leaders will maintain a mission-linked plan for responsible AI development. At least annually, they will examine opportunities to improve learning, service and institutional work, the consequences of adoption and delay, and the capability and access needed to pursue worthwhile uses. The plan must identify priorities, accountable owners, bounded trials where appropriate and evidence for continuation. Adoption for its own sake is not an institutional outcome.

2 Scope and definitions

This policy applies to faculty, staff, students, trustees, contractors and others acting for the institution when developing, acquiring, configuring or using AI for institutional activities. It covers credit and noncredit instruction, workforce education, research, student services and administrative operations, regardless of device or account ownership.

AI includes systems that generate content, classify or predict, recommend or support decisions, or carry out actions. It includes AI features within existing software. An AI agent is a system that can use connected tools or carry out a sequence of actions within assigned permissions.

A consequential use materially affects a person's rights, safety, educational standing, employment, financial support or access to essential services. Protected information means information subject to institutional, contractual or legal restrictions on access or use, including student and employee records, health information, credentials and confidential partner information.

This policy works with existing academic freedom, shared governance, integrity, accessibility, privacy, security, procurement, records, employment and conduct requirements. System and district requirements remain applicable. Program standards may add safeguards within their authority. Conflicts must be resolved by the policy owner and the appropriate academic or administrative authority before the affected use proceeds.

The policy owner maintains practical scope guidance. Ordinary software functions that do not introduce a material change in data access, decision influence or automated action may be excluded from separate AI review. An exclusion does not remove existing obligations.

Accessibility accommodations, employment agreements, research approvals and program obligations must be reconciled through the responsible offices. No individual approval under this policy changes another office's reserved authority. The policy owner records unresolved conflicts, identifies the deciding authority and communicates the interim permitted scope.

3 Governance and decision rights

The policy owner is accountable for implementation, resourcing, coordination and annual reporting to the approval authority. The institution may assign these duties to existing governance bodies and staff. The size and form of the process must fit institutional capacity while meeting this policy's requirements.

The policy owner designates a coordinating group or existing body with appropriate academic, administrative, technical and student participation. It must have access to privacy, security, accessibility, procurement, legal and relevant program expertise. It maintains standards, resolves cross-unit questions and recommends improvements. Recommendations and approvals must be distinguished in its terms of reference.

Academic governance retains its established authority over curriculum, teaching and assessment. Faculty determine course and assignment uses within applicable program and institutional requirements. Relevant employee and student representatives participate in decisions that materially affect their work, learning or access to services.

Each institutional AI service, integration or consequential deployment must have a named owner accountable for its purpose, outcomes, affected people and continuing operation. Technical responsibilities must also be assigned. One person may hold several roles when qualified and adequately supported. A vendor or committee does not replace a named institutional owner.

Reviewers must have the expertise, information, time and authority needed to carry out their duties. Users must follow permitted purposes, protect information, check outputs appropriately and report concerns. The institution will publish a current directory of decision, support, complaint and incident contacts.

4 Permission and proportionate review

Routine use is permitted without individual case approval when it uses public or otherwise authorized information, follows published tool and account standards, remains under the user's control, and does not trigger the review conditions below. Examples include exploring ideas, drafting material for review, generating practice examples and summarizing public information. Course and assignment rules continue to apply.

The institution will publish approved tools and standard uses, allowed data, account requirements, limits and support contacts. It will provide a route to propose additional tools or uses. Standard permission authorizes only the stated conditions; approval of a product does not authorize every possible use. Normal procurement and software installation requirements remain in effect.

Prior review is required when a proposed use:

- Processes protected information beyond an existing authorization.
- Materially influences a consequential decision or creates substantial safety, fairness or accessibility concerns.
- Connects to institutional systems, expands access to records or takes actions on behalf of people.
- Creates an institutional AI service or substantially changes an existing service's purpose, users, data or permissions.
- Introduces a new contract, vendor or material AI feature requiring procurement or security review.

Review must document the intended benefit, affected people, alternatives, information and rights, likely failures, safeguards, testing, responsible owners and monitoring. The decision may authorize a standard use, approve a limited pilot, require changes, approve deployment or decline the use. Pilot authorization does not automatically permit general deployment.

A live pilot or deployment must not proceed while testing demonstrates unauthorized information access or actions, required information-handling or contractual protections are unresolved, no accountable owner can stop the service, or an essential service lacks a workable fallback. Other critical security findings must be resolved or addressed through an authorized, time-limited exception with demonstrated compensating safeguards under section 12. Exceptions cannot waive applicable requirements or permit the failures listed above.

The institution will publish review service standards. The default is acknowledgement within five business days and a decision or explained timetable within fifteen business days. Delays escalate to the policy owner. Silence does not authorize a use requiring review. Existing approval requirements may establish different published timeframes.

Uses that violate applicable requirements, deceive people about identity or authorship, evade security controls, or remove human authority reserved by this policy are prohibited.

AI must not be used to manipulate people through deceptive impersonation or to make unsupported inferences about their emotions, character or sensitive personal attributes for consequential decisions. Monitoring or behavioral profiling requires a documented necessity, an assessment of less intrusive alternatives, appropriate notice and specific review of privacy, fairness and existing rights. Review does not authorize an otherwise prohibited use.

5 Information security and procurement

Information may be provided to AI only when the information and intended use are authorized for the particular tool, account and configuration. Protected information must not be entered into unapproved public or personal services. Use the minimum information needed and follow applicable access and retention rules. Removing names alone does not establish that information cannot identify someone.

Student work, assessment materials, employee records and employer or clinical-partner information must receive the protections required by their ownership, sensitivity and permitted use. Public availability does not itself establish permission to reproduce copyrighted material or use it to train a model.

AI connections must respect existing access permissions. System owners must approve integrations, restrict access to the task and prevent connected services from extending access beyond the approved purpose. Credentials must be managed through authorized security mechanisms.

Institutionally provided AI services and integrations must meet the institution's applicable information-security standards. For services that process protected information or act within institutional systems, require multifactor authentication for human access; managed, task-limited service identities; timely access removal; encryption of protected information in transit and at rest; and controlled credential and encryption-key storage, rotation and revocation. Apply institutional vulnerability-remediation requirements, protected activity logging, and tested recovery or service-fallback arrangements. Required controls must be established before release; any permissible alternative must follow section 12 and applicable requirements.

Treat AI-generated outputs as untrusted input to connected systems. Validate tool arguments, encode displayed output for its context, and use parameterized database operations. Generated code or commands may execute only through an explicitly approved, constrained process with appropriate isolation and testing. A model response must not itself grant authority or bypass normal application security controls.

Procurement and renewal review must address the safeguards relevant to the deployment, including data use and training, security, accessibility, subcontractors, retention and deletion, incident notification, material changes, intellectual property, service continuity and exit arrangements. A subscription tier or vendor assurance alone does not establish that a use meets institutional requirements.

Reviewers must retain sufficient supplier evidence to support the decision, record unresolved limitations and specify conditions of approval. Where essential safeguards cannot be established, the use must be limited, deferred or declined. The service owner must assess material provider changes before continuing affected consequential uses and maintain a workable route to export, transfer or delete institutional information when a service ends.

Use the institution's established vendor-assessment process, such as HECVAT or an equivalent, and supplement it for the proposed AI use. Verify that answers and supporting evidence cover the actual product, configuration, service tier and relevant subcontractors. A completed questionnaire or assurance report is evidence to evaluate, not approval by itself.

Owners of institutionally configured AI services must identify authoritative information sources, permitted reuse, responsible data stewards and update schedules. They must check the relevance, quality and currency of information for the intended task, provide a correction route and prevent

superseded or unauthorized material from continuing to guide the service. Training or evaluation data must have appropriate provenance and permissions.

Restrict who and what can add or change knowledge sources, retrieval indexes, training data and model components. Review their provenance and changes, maintain version records, and test for unauthorized or malicious content before use. Preserve access restrictions through retrieval and provide a tested way to remove compromised material or restore an approved version, including affected indexes and caches.

Institutional or protected information may not be used to train or improve an external provider's models without specific authorization and an appropriate agreement. New AI features within existing software require reassessment when they materially change information handling, decision influence or permitted actions.

Recording, transcription and AI-enabled wearables must comply with institutional recording rules, applicable consent requirements and approved accommodations. Confidential classes, meetings, clinical encounters and workplace interactions must not be recorded or transmitted to AI without the required authorization. The institution will provide a clear route for resolving accommodation and recording requirements together.

6 Human judgment and consequential decisions

A qualified, authorized person must make final determinations materially affecting admission, grades, academic standing, financial support, employment, discipline, access to essential services or comparable rights and interests. AI may assist only within an approved process. Clinical, safety-critical and other regulated activities remain subject to their additional requirements.

Human review must be meaningful. Reviewers must consider relevant evidence, understand the system's limits, identify when its output is unreliable and be able to reject, correct or stop its recommendation. Routine approval without substantive review does not meet this requirement.

Before consequential deployment, the owner must test performance in the intended setting, examine relevant disparities and access barriers, define acceptable limits and establish correction and escalation routes. The owner must assess whether the human review process detects important errors and supports defensible decisions.

Testing must use relevant tasks and affected groups where appropriate information can be obtained lawfully and responsibly. The record must identify uncertainty, known limitations, unacceptable failure conditions and the action required when limits are exceeded. Insufficient evidence of fitness cannot be replaced by an unsupported assurance that a person will check the output.

Affected people must receive understandable notice of AI's material role and a usable route to review by an authorized person. Existing appeal and complaint processes apply, with their published timelines and safeguards. People must be able to raise an AI-related concern without having to understand the technology in detail.

The responsible office must explain the actual basis of a consequential decision in terms the affected person can use, including relevant information, the role of AI and the route to correct an error. Corrections must address affected institutional records and downstream decisions where necessary. A vendor's inability to explain a recommendation does not remove this responsibility.

7 Teaching learning and workforce preparation

Faculty and programs determine how AI advances learning objectives and when independent work is required to demonstrate competence. Instructions must state permitted, required and prohibited assistance, disclosure expectations, assessment criteria and available support. Assignment-specific rules must be communicated clearly and in time for students to follow them.

The institutional default permits AI for general learning support, such as practice and explanation, subject to data safeguards and clearly communicated course restrictions. AI-generated answers or substantive assessed work require permission in the course or assignment instructions. When instructions are unclear, faculty must clarify the expectation and address ambiguity fairly before imposing a misconduct finding.

Programs must provide appropriate opportunities to demonstrate independent understanding and, where relevant, responsible performance with AI. A polished AI-assisted product is not sufficient evidence that a student has acquired the knowledge or skills being assessed. Faculty may use demonstrations, explanations, process evidence or other suitable assessment methods.

Credit and noncredit workforce programs must connect AI permissions to occupational competencies, safety requirements and employer or placement obligations. Clinical care, laboratory activity and technical work may require tighter rules than general classroom use. AI may not substitute for a required competency demonstration or authorize a procedure beyond a learner's training and supervision.

Program leaders must document additional requirements for clinical placements, apprenticeships, employer projects and dual-enrollment learners. These requirements must address relevant partner agreements, information restrictions, age or account conditions, supervision and assessment. A partner's preferred tool does not override institutional safeguards.

When AI use is required, the institution must provide a reasonably accessible route that does not depend on a student buying an unapproved personal subscription. Approved accommodations remain available. An equivalent alternative should be provided when it meets the learning objective. Essential AI requirements and limits on alternatives must be addressed through established academic and accessibility processes.

Faculty remain responsible for course content, assessment and feedback. AI may support these activities within authorized data and assessment conditions. Material use affecting evaluation must be explained to students. AI-generated feedback does not relieve faculty of their responsibility to provide the teaching and interaction required by the course.

Academic and workforce leaders will periodically review how AI changes relevant disciplinary and occupational tasks. Through established academic processes, they will seek appropriate employer, worker, learner and community input and identify where curricula, continuing education, apprenticeships or other learning pathways need revision. The review must preserve broad educational aims, human judgment and transferable capability, including the ability to work independently when needed.

Preparation plans must consider adult and noncredit learners, underserved communities, smaller employers and others who may face barriers to participation. Claims about graduate readiness, placement or economic benefit must be supported by relevant evidence. A vendor credential or tool demonstration alone does not establish occupational competence.

8 Integrity disclosure and intellectual property

Users must represent their work and sources honestly. Fabricated evidence or citations, deceptive impersonation, concealed prohibited assistance and the presentation of synthetic data as observed data are prohibited. Legitimate simulations and synthetic data must be identified and used with appropriate safeguards.

Substantive AI assistance must be disclosed when it materially shapes a work's reasoning, findings, interpretation or original expression and the audience could reasonably attribute that contribution to a person. The disclosure should identify the tool and its role in terms useful to the audience. Users remain responsible for the result.

Routine spelling, formatting and comparable assistance do not require separate disclosure unless a course, professional standard, publisher, sponsor or other applicable requirement says otherwise. Assistance that alters assessed reasoning or substantive meaning is not routine merely because the tool describes it as editing. Disclosures must not require unnecessary revelation of a person's disability or accommodation.

Academic and research work must follow applicable disciplinary, course, sponsor, publisher and institutional requirements. Researchers must verify sources, protect confidential material, preserve appropriate methodological records and obtain required approvals. AI cannot assume an author's responsibility. Confidential review material must not be shared without the necessary authority and safeguards.

Allegations of misuse must follow existing fair procedures. The institution must establish the applicable rule, examine reliable evidence, allow the person to respond and provide the appropriate review or appeal route. An AI detector score or AI opinion about authorship must not, by itself, establish misconduct. Educational correction should be considered when unclear instructions or insufficient preparation materially contributed to the problem.

9 Automated actions and AI agents

An AI agent may act only within documented authority. The owner must specify its task, accessible information and systems, permitted actions, spending limits, duration and escalation conditions. Access must be limited to what the task requires. Content encountered in documents, messages or websites cannot authorize expansion of those permissions.

Before deployment, the owner must test likely failures and misuse, establish proportionate activity records, provide a way to interrupt operation and plan recovery or reversal where feasible. Material changes to connected systems, models or permissions require reassessment.

Low-consequence, reversible actions may be automated within an approved scope and monitoring plan. Actions with material financial, academic, employment, legal or personal consequences require the authorization specified for that action. Agents may not make final determinations reserved to people under section 6. Permission to draft does not itself authorize sending, publishing, purchasing or changing a record.

Agents must use institutionally managed identities and credentials appropriate to the task. Owners must define which systems and tools may be used, where permissions are enforced, when a person must confirm an action and how long authority remains valid. An agent may not approve its own expansion of

access or authority. Any delegation to another agent must remain within the approved scope and preserve an accountable institutional owner.

Testing must address misleading instructions in retrieved content, unauthorized information disclosure, incorrect tool use, partial completion and failure across connected systems. Owners must verify that critical limits are enforced by system controls and authorized processes, rather than relying solely on instructions given to the model.

Activity records must support reconstruction of material actions without retaining unnecessary personal information. Recovery plans must address downstream records, recipients and transactions affected by an error. Where an action cannot be reversed, approval must consider that consequence before execution. New tools, delegated agents, persistence, transaction limits or action capabilities require review proportionate to the changed risk.

10 Access learning and institutional support

Institutional AI services must identify themselves clearly, explain relevant limitations and provide an appropriate route to human assistance. Their design must not mislead people about their identity, capability or authority.

Student-facing services must publish their service boundaries, human contact route and response expectations. They must not represent themselves as a licensed professional or promise crisis response they cannot provide. Services involving health, wellbeing, younger learners or other heightened vulnerabilities require relevant professional review and an appropriate referral process. Design must support user choice and avoid manipulative pressure to continue an interaction.

The institution will provide role-appropriate learning, practical guidance and support for the tasks people are expected to perform. This includes choosing suitable uses, protecting information, checking outputs, recognizing limitations, exercising judgment and remaining accountable. Adjunct faculty, part-time staff, student workers and people in noncredit programs must receive relevant guidance when their roles fall within this policy.

Required deployment must be matched with sufficient tools, staff time and support. Owners must consider disability access, affordability, language, connectivity and other barriers relevant to the community. Procurement and deployment choices should account for total cost, workload, service continuity and resource use at a level proportionate to the proposed benefit.

Before requiring a new AI-enabled workflow, the responsible leader must assess effects on roles, workload, supervision, professional development and performance evaluation, and involve the applicable employee representatives and authorities. Access to an approved tool or volume of AI use is not, by itself, a defensible measure of employee performance. Required instructional use must provide accessible participation or an appropriate equivalent route that does not penalize a learner for an unmet access need.

Where communities hold distinct rights or interests in information or knowledge, appropriate representatives and institutional experts must be involved in decisions affecting those interests.

Leaders must support practical experimentation within authorized boundaries and publish a route for proposing improvements. Where a useful proposal cannot proceed, the responsible office should explain the condition to resolve or the reason for declining it. Capability building and fair workforce transition must be included when AI materially changes institutional work.

11 Monitoring incidents and records

The institution will maintain a register of institutional AI services, integrations and consequential deployments. It must identify purpose, owner, provider, information categories, affected groups, permissions, approvals, testing and review dates. Routine individual uses within published standard permissions do not each require registration.

Owners must define the intended benefit and unacceptable performance before deployment. Monitoring must examine relevant outcomes, errors, complaints, disparities, accessibility, cost and workload. Learning applications require evidence of learning. Service applications require evidence of service quality. Administrative uses require evidence relevant to accuracy, timeliness, control effectiveness and workload. Workforce and regional initiatives require proportionate evidence of capability, access and mission contribution; employment or economic claims must be independently supportable. Changes to purpose, information, users, permissions or material system behavior trigger reassessment.

Suspected harmful outcomes, unauthorized actions, information exposure and material failures must be reported promptly through the institution's published incident route. Owners must be able to pause use, preserve necessary evidence, protect affected people and activate a service fallback. Responsible offices determine notification, remediation and restart conditions. Good-faith reporting must not result in retaliation.

Records must be sufficient to explain approvals and material decisions without collecting unnecessary personal information. Access, retention and deletion follow existing schedules and obligations. Routine learning or experimentation does not justify unrestricted surveillance of students or employees.

12 Implementation exceptions and review

The policy owner will publish implementation responsibilities, supporting standards, review forms, approved uses and contact routes. Existing deployments will be assessed under a transition plan that prioritizes consequential uses and protected information. Unresolved high-consequence uses must be restricted until the necessary safeguards and approvals are in place.

Exceptions must state their purpose, scope, compensating safeguards, accountable owner, approving authority and expiration. An exception cannot authorize conduct contrary to applicable requirements or override another body's authority. Exceptions must be reviewed when circumstances change and recorded with the relevant deployment or decision.

The policy will be reviewed at least annually and after significant incidents or material changes. The policy owner's report to the approval authority must address outcomes, unresolved risks, access, staff and student capability, review timeliness and proposed improvements. Revisions must involve the relevant academic, employee and student governance processes. Violations are addressed through existing proportionate conduct procedures.

For consequential deployments, the approval authority must assign a reviewer sufficiently independent of day-to-day operation to challenge the evidence and conditions of continued use. A small institution may use qualified cross-unit, district or external expertise. Reviews must record conflicts of interest, corrective actions and decisions to continue, restrict, suspend or retire a service. Expired pilots and exceptions do not renew automatically.