Brunel University of London

Information Governance Framework:

AI Governance & Security Standard

Contents

1. Document Control

Title	AI Governance & Security Standard

Version	1.0
Publish date	September 2026
Next review date	September 2027
Policy owner	Deputy COO and Director of Business Operations
Stakeholders consulted in development	Digital Services Data Protection Governance & Secretariat AI Strategy Committee Associate Directors and Heads of Department Information Assurance Committee
For information & action	All staff, contract staff, student and third-parties handling university data.
Supersedes	N/A
Supporting framework, policies & standards	Information Governance Framework IT Acceptable Usage Policy Data Classification & Handling Policy Cyber-Supply Chain Risk Management Policy Software/AI Tool & Use Case Approval Standard

2. Purpose

The AI Governance and Security Standard establishes proportionate governance and security arrangements for the procurement, development, deployment and responsible use of Artificial Intelligence (AI) systems and tools at Brunel University of London (BUoL). It is intended to

enable safe, confident and innovative use of AI across teaching, research and professional services, while ensuring that legal, ethical, data protection, information security and assurance requirements are met.

This Standard is not intended to restrict innovation or routine low-risk use of approved AI services. Low-risk use may proceed under defined guardrails and local oversight once recorded in the AI use case register. Additional assessment, review and approval apply where the proposed use case falls outside those guardrails, including where the AI service is unapproved, the use is outside existing approval conditions, or the risk, data, integration, automation or potential impact requires it.

AI can enhance teaching, research and professional services, but must be managed through existing University governance, cyber security, data protection, procurement, research ethics and risk management arrangements.

This Standard ensures that:

- AI is adopted safely, lawfully and responsibly without unnecessary barriers to beneficial use.
- Controls are proportionate to the risk, data, context, automation and impact of the use case.
- Sensitive, personal, confidential, research-sensitive and regulated data is protected.
- Human accountability, oversight, assurance and incident response remain clear.

3. Scope

This standard applies to:

- All staff, students, contractors, consultants, visiting academics, research partners and third parties using AI on behalf of or in connection with the University;
- All AI systems and capabilities, including generative AI, machine learning, predictive analytics, embedded AI functionality and agentic AI services;
- AI used across teaching, learning, research, student services, administration, decision support, automation, cyber defence and operational activities;
- AI capabilities that are procured, developed internally, integrated with University systems, or embedded within third-party platforms;
- AI use regardless of delivery model, including free, paid, open source, commercial, internally developed, cloud-hosted or locally hosted systems, browser-based tools, APIs or embedded software features;
- Security and risk management controls required to protect the University from external AI-enabled threats.

Personal use of AI services that does not involve University business, University systems, University data, University obligations or University-managed devices falls outside the scope of

this standard. However, staff and students remain responsible for complying with all relevant University policies when acting in a University capacity.

4. Key definitions

Term	Definition
Artificial Intelligence (AI)	A broad term covering systems or functionality that perform tasks normally requiring human intelligence, including systems that generate text, code, images, predictions, classifications, recommendations or decisions.
AI system	An engineered system that uses AI techniques to process inputs, identify patterns, make inferences, support decisions or actions, or produce outputs such as content, predictions, classifications or recommendations for defined objectives.
AI-enabled system	A system, platform or service where AI functionality is embedded within broader software or operational capability.
AI tool	An AI service, application, model, platform, plug-in, API, embedded feature, script or capability, whether cloud-based, locally hosted, open source, commercial or internally developed, used to generate, classify, analyse, predict, recommend, automate or support decisions or outputs.
Approved AI Service	An AI tool, platform, system or embedded AI capability that has been approved for defined use under University governance arrangements.
Generative AI	AI that generates text, images, code, audio, video, data or other content in response to prompts or inputs.
Internally developed AI	AI functionality, models, scripts, workflows or applications developed by University staff, students, researchers or contractors for teaching, research, operational or administrative purposes.
Prototype / exploratory AI use	Experimental, sandboxed or developmental AI activity that is not used for operational decisions, does not process sensitive data, and is not deployed beyond a controlled teaching, research or test environment.
Deployment	Making an AI system or capability available for operational, institutional, external-facing or repeatable use beyond individual experimentation or controlled teaching, research or test activity.
High-risk AI use case	A use case where failure, misuse, inaccuracy, bias, loss of control, data exposure or security compromise could materially affect individuals, University operations, legal compliance, research integrity, finances or reputation.
Human Oversight Model	A control model where an appropriately accountable person reviews, approves, can intervene in, or can stop an AI output, decision or action before material reliance is placed on it.

Security Operations Centre (SOC)	The function responsible for security monitoring, detection, triage and response.
University Data	Any information created, received, processed, stored or shared by staff, students, contractors, researchers or third parties in connection with University business, teaching, learning, research, administration or operations. This includes, but is not limited to, teaching materials, assessment content, research data and outputs, student and staff information, operational records, intellectual property, emails, documents, slide decks and data held in University systems. University Data may be public, internal, confidential, restricted or sensitive, and controls must be proportionate to its classification, sensitivity and context of use.

5. Governing Principles

The use of AI within BUoL shall be governed by the following principles.

- **Enable innovation responsibly.** AI governance shall support safe and effective adoption while keeping risks within the University's risk appetite.
- **Proportionality.** Controls shall reflect the risk, sensitivity, impact and context of use. Low-risk use may proceed under defined guardrails once recorded in the AI use case register; higher-risk, integrated, automated or sensitive use requires proportionate review and approval.
- **Human accountability.** Humans remain accountable for AI deployment, access, safeguards, outputs and decisions made in response to AI-generated or AI-enabled content.
- **Security and data protection by design.** AI systems must be designed, configured and operated securely, with appropriate controls for personal, confidential, commercially sensitive, research-sensitive or restricted data.
- **Least privilege and bounded access.** AI systems shall have only the access, functionality and privileges needed for the approved purpose, particularly where they connect to systems, data sources, APIs or automated workflows.
- **Transparency, oversight and auditability.** AI usage, risks, decisions and ownership must be documented and subject to oversight sufficient to support assurance, investigation and incident response.

6. Accountability

The Deputy COO and Director of Business Operations have overall accountability for this Standard. Day-to-day responsibilities are defined through the Responsible, Accountable, Consulted and Informed (RACI) model below and supporting procedures.

6.1 Key roles & responsibilities

The following RACI model defines the main accountability and responsibility arrangements for AI governance and security. It should be read alongside the supporting AI use case process, delegated approval matrix and relevant University policies.

Activity	Deputy COO	IAC	Head of Information Security	Head of Privacy	Digital Services / Digital Education	Academic / Research Lead	AI System Owner	Staff / Students
Maintain AI Governance & Security Standard	A	C	R	C	C	C	I	I
Maintain AI use case process and register	A	C	R	C	C	C	C	I
Record low-risk AI use cases	I	I	C	C	C	R	A/R	R
Assess and approve medium-risk AI use cases	I	C	C/R	C/R	C/R	C/R	A/R	I
Review and approve high-risk AI use cases	C	A/R	C/R	C/R	C	C/R	A/R	I
Define AI security requirements	I	C	A/R	C	C/R	C	C	I
Complete privacy assessment or DPIA	I	I	C	A/R	C	C/R	R	I

where required								
Review supplier and contractual arrangements	I	I	C	C	C/R	C/R	A/R	I
Approve teaching and learning AI use	I	C	C	C	C/R	A/R	R	I
Approve research AI use	I	C	C	C	C	A/R	R	I
Monitor AI-related security risks and incidents	I	C	A/R	C	R	I	R	I
Provide training, user guidance and awareness	A	C	C/R	C/R	C/R	C/R	C	I
Review exceptions	A	C	C/R	C/R	C/R	C/R	C/R	I
Report assurance and compliance status	A	C/R	R	C	C	C	C	I

7. Context of Use

7.1 AI Tool Usage and Approval

AI services used for University business, teaching or research where University data, systems or obligations are involved must be approved, recorded or otherwise covered by University governance arrangements appropriate to the risk and context of use.

Approval or guidance applies to the defined tool, context and use case. Approval of a tool does not automatically permit all possible uses of that tool, particularly where sensitive data, automation, integration or material decisions are involved. Low-risk use does not require separate approval before each use where it uses an approved AI service and remains within defined guardrails. Where an AI service is unapproved, or the proposed use falls outside approved conditions or low-risk guardrails, proportionate security, privacy and supplier triage is required before the use can proceed.

Users must not expose personal, special category, confidential, sensitive research, credential, security, commercially sensitive or legally privileged information to AI services unless the use is authorised, documented, controlled and consistent with the relevant approval conditions.

Where an alternative AI tool is required, the requester should provide a proportionate justification explaining the intended use, expected benefits, data involved, risk considerations and why existing approved tools are insufficient.

7.2 Use Cases Requiring Review and Approval

Use cases involving University data exposure, system integration, automation, operational decision support, material outcomes for individuals, or financial, legal, regulatory, security or compliance impact require proportionate review and approval before deployment.

7.3 Human Review and Accountability

AI outputs must be subject to appropriate human review and oversight before being relied upon for material decisions or actions. Full requirements are defined in Section 12.

7.4 Change and Reuse Controls

New or materially changed use cases should be reassessed where the change affects data access, integration, automation, decision impact, security risk, supplier arrangements or the approved context of use. Approval is therefore granted for a specific use case and context, not blanket use of a tool across materially different purposes.

7.5 Academic and teaching

The University supports innovation, including specialised, experimental or discipline-specific AI use. Controls should be proportionate to the data, context, student impact and level of automation involved.

Teaching use must still comply with all relevant University requirements, including data protection, information security, accessibility, academic integrity, intellectual property and acceptable use policies.

Routine teaching activities, including labs, coursework, demonstrations, student exercises, final-year projects, dissertations and supervised student prototypes, may be treated as low-risk academic use where they use approved, local, sandboxed or otherwise controlled AI tools, involve only synthetic, anonymised, public or otherwise non-sensitive data, and are not deployed beyond the learning environment or used to make decisions affecting individuals.

Where AI tools are recommended or required for students, the School or Department is responsible for ensuring that appropriate consideration is given to privacy, accessibility, licensing, equity of access, and assessment integrity.

7.6 Research use

Research use may involve specialised tools, experimental models, high-performance computing, external collaborators or domain-specific platforms. Controls must be proportionate to the nature, sensitivity and impact of the research while managing data protection, security, ethical, contractual and legal risk.

Low-risk research experimentation, including local model testing, offline experimentation, open-source tool evaluation, sandboxed development, PhD prototypes and early-stage research using approved, local, sandboxed or otherwise controlled tools and synthetic, dummy, anonymised, public or otherwise non-sensitive data, may proceed under defined guardrails and local oversight once recorded in the AI use case register. Separate approval is not required for each use case where the activity uses an approved AI service or otherwise controlled tool and remains within the approved conditions and low-risk guardrails. Where the AI service is unapproved, or the proposed use falls outside approved conditions or low-risk guardrails, proportionate security, privacy and supplier triage is required before the use can proceed. Higher-risk research use should be assessed individually, taking account of the data, purpose, collaborators, funder or contractual obligations, ethical considerations, external transmission of University Data, system integration, automated decisions and the potential impact of outputs or decisions.

Research involving personal or special category data, human participants, sensitive research data, funder or contractual obligations, dual-use concerns, export controls, NDA-bound data or external collaborators may require DPIA, research ethics, legal, contractual, supplier and information security review.

Where external submission platforms, publishers, funders or conference systems use AI screening or assessment functionality outside the researcher's control, researchers are not responsible for configuring or approving that AI functionality. However, researchers must take reasonable steps to understand any disclosed AI use, platform terms and data handling implications before submission, and must ensure that submission remains consistent with research ethics, confidentiality, intellectual property, funder, contractual, data protection and University requirements. Concerns about sensitive, restricted, unpublished, personal, NDA-

bound or otherwise high-risk material should be escalated through the appropriate University channels before submission.

7.7 Student use

Student use of AI is governed by the IT Acceptable Use Policy, academic integrity rules, assessment guidance and any approved academic-level policies. These frameworks collectively define the acceptable, ethical and lawful use of AI in learning and assessment.

Students must not be instructed to upload personal data, third-party confidential information or University restricted data into AI tools.

Where AI is used in learning or assessment, expectations must be clearly defined and communicated to students. This includes explaining what AI use by students is permitted, restricted or prohibited, and where AI may be used by staff, the University or digital platforms to support teaching, feedback, assessment, academic integrity checks, screening or other processing of student work. Transparency must apply in both directions so that students understand both how they may use AI and how AI may be used in relation to their work.

7.8 Acceptable Use & Prohibited Use

For the purposes of this standard, University Data includes teaching materials, documents, slide decks and other content created, handled or used for University purposes. Not all University Data carries the same level of risk. The use of AI to support low-risk activities, such as improving the layout, formatting or readability of non-sensitive teaching materials, may be permitted where an approved AI service is used and no personal, confidential, restricted, unpublished, assessment, student, commercially sensitive, research-sensitive or otherwise sensitive information is exposed.

The following activities are not permitted:

- Use of free-tier, consumer or unapproved AI tools for University business, University data or University-controlled activity where the tool and use case have not been registered and approved.
- Inputting personal data, special category data, confidential information, sensitive research data, credentials, security configuration details, commercially sensitive information or legally privileged material into unapproved AI services.
- Granting AI services unrestricted or excessive access to University systems or data, including mailboxes, file storage, SharePoint sites, student records, HR records, finance systems, research repositories or security systems.
- Deploying AI systems where outputs or actions cannot be adequately monitored, constrained, explained, challenged or stopped.
- Using AI outputs as the sole basis for material decisions affecting students, staff, applicants, research participants, finances, compliance obligations, safety or legal rights.

- Using AI to bypass University controls, access restrictions, monitoring arrangements, licensing conditions, academic rules or security safeguards.
- Using AI in a manner that breaches University policies, contractual obligations, intellectual property rights, data protection legislation, research ethics requirements or academic integrity rules.

8. Risk Assessment and Classification

8.1 AI use case assessment

AI use cases must be recorded in the AI use case register through the University's AI use case process. Routine low-risk use may proceed through simplified registration and fast-track approval where it uses an approved AI service, involves only synthetic, anonymised, public or otherwise non-sensitive data, remains within defined guardrails, and does not involve system integration, automation, material decisions or significant impact on individuals or University operations. Additional documented assessment is required where the risk-based model requires it.

Where assessment is required, it shall be proportionate and consider the purpose, data, AI capability, system integration, decision impact, human oversight, security risks, supplier arrangements, and whether DPIA, ethics, legal or procurement review is required.

8.2 Risk classification

AI use cases shall be classified as Low, Medium or High through the AI use case assessment process. Classification shall consider the nature of the data, intended purpose, level of automation, degree of integration, human oversight, supplier arrangements and potential impact on individuals, University operations, legal obligations, research integrity, finances, reputation or regulatory compliance.

The classification should be increased where the use case involves sensitive or personal data, material decisions, student or staff outcomes, external AI services, integration with University systems, autonomous actions, limited explainability, or increased exposure to AI-enabled security threats.

Risk classification should also take account of the deployment model, including whether the AI capability is locally hosted, offline, open source, sandboxed, cloud-hosted, externally managed, consumer-grade, integrated with University systems or connected to third-party services. Higher risk is likely where University Data leaves controlled environments, where vendor reuse or model training may occur, or where access controls, provenance, logging or contractual protections are unclear.

The resulting classification determines the level of assessment, approval, assurance and review required.

8.3 Risk-based Usage Model

AI governance shall operate on a proportionate, risk-based model.

Low-risk use may include routine productivity, teaching or research activity using synthetic, dummy, anonymised, public or otherwise non-sensitive data, where outputs are reviewed by a human and do not materially affect individuals, legal obligations, finances, regulatory compliance or University operations. Low-risk deployment models may include approved University AI services used within published guardrails, local or offline model testing, sandboxed teaching or research environments, controlled open-source tool evaluation, or isolated prototypes that do not connect to University systems, process sensitive data or transmit University Data outside controlled environments.

Medium-risk use may include use cases involving University information, internal processes, limited system integration, structured outputs, repeatable workflows or decision support requiring increased oversight. Medium-risk deployment models may include approved cloud-based AI services used with internal or controlled data, AI functionality embedded within University-supported platforms, limited integrations with University systems, managed supplier services with contractual controls, or sandboxed environments that involve controlled data flows, user groups or repeatable business, teaching or research processes.

High-risk use includes use cases involving sensitive data, personal data at scale, special category data, student or staff outcomes, material decisions, operational automation, regulatory reporting, financial impact, legal obligations, significant research sensitivity, external integrations or autonomous actions. High-risk deployment models may include externally hosted or consumer AI services processing University Data, AI systems integrated with core University systems or authoritative data sources, agentic or automated workflows able to take actions on behalf of users, supplier-managed AI services with unclear data reuse or model training arrangements, externally facing AI services, or deployments where errors, bias, unauthorised access, data exposure or loss of control could materially affect individuals, University operations, compliance, finances, research integrity or reputation.

The level of assessment, approval, assurance and review shall increase with risk. All University AI use cases, including low-risk use, must be recorded in the AI use case register through the AI use case process.

8.4 Approval and Oversight

Approval, registration and oversight requirements shall be proportionate to the assessed risk. All AI use cases must be recorded in the AI use case register through the University's AI use case process.

Low-risk use may proceed under defined guardrails and local oversight once recorded in the AI use case register. It does not normally require wider approval unless additional risk factors are identified.

Medium-risk use must be registered, supported by a documented assessment, and approved by the appropriate accountable owner(s).

High-risk use must be supported by a documented assessment and undergo proportionate formal review and approval through the appropriate governance route, which may include Information Security, Data Protection, Legal, Procurement, Research Ethics, the Information Assurance Committee (IAC), or other relevant bodies.

Detailed approval routes, delegated authority levels, assessment requirements, registration processes and review workflows shall be defined in supporting procedures and guidance.

9. Procurement & Suppliers

Supplier AI risk must be considered where a third party provides an AI product or AI-enabled service, or where any supplier may use AI when processing University Data or delivering services to the University.

Supplier assessment is mandatory for high-risk AI use cases and may also be required where a supplier processes University data, integrates with University systems, accesses University information assets, uses AI to support service delivery, introduces embedded AI functionality, or presents material legal, regulatory, security, operational, academic or research risk.

Supplier due diligence must therefore consider both AI functionality provided to the University and the supplier's own use of AI in handling, processing, analysing, storing, accessing or otherwise supporting University services.

The depth of supplier assessment must be proportionate to the risk, data involved, level of integration and potential impact on University operations and people. Where relevant, due diligence should consider security posture, data processing, model training or vendor reuse of University Data, data residency, access control, logging, incident response, vulnerability management, subcontractors, contractual rights, transparency of supplier AI use and exit arrangements.

Supplier use of AI introduced through new services, product updates, new modules, feature releases, operational process changes or subcontractor arrangements must be reassessed where it materially changes data access, processing, model behaviour, data reuse, user impact, risk, assurance obligations or the University's ability to monitor, configure, restrict or control how University data is used.

10. Security Control Baseline

Security controls must be proportionate to the AI system's risk level. The following requirements provide the minimum baseline for material AI systems and AI-enabled services.

Threat modelling and risk assessment: The AI system owner is responsible for ensuring appropriate risk assessment before deployment. Low-risk use may be addressed through the AI use case process; medium- and high-risk use, sensitive data, system integration, elevated

access, automation or material impact require specialist input where relevant. AI-specific threats such as prompt injection, data leakage, model manipulation, unsafe automation, insider misuse and AI-enabled attacks should be considered proportionately.

Secure configuration: Tenant, model, prompt, API, connector and service settings must be configured securely. Unsafe features should be disabled by default where appropriate. Configuration decisions and material changes must be documented.

Identity and access control: AI systems must use appropriate identity and access controls, including strong authentication, role-based access, privileged access management where relevant, separation of duties and periodic access review.

Data protection controls: Data minimisation, classification, approved data sources, retention controls, masking, anonymisation and pseudonymisation should be applied where feasible. Controls must restrict inappropriate vendor training, downstream reuse or unauthorised exposure of University data.

Connector and tool control: Plugins, connectors, APIs, agents and workflow automation must be reviewed for data access, action rights, permission inheritance, auditability, containment and revocation. Particular care must be taken where embedded AI functionality can access mailboxes, calendars, document repositories, collaboration platforms, student records, HR records, finance systems or security systems.

Testing and validation: AI systems must be tested proportionately before go-live. Testing should consider security, privacy, accuracy, bias, misuse, abuse cases, access control, data leakage, hallucination risk, resilience, performance and rollback.

Logging and monitoring: Logs and monitoring must be sufficient to support accountability, investigation, anomaly detection and assurance, including unusual data access, misuse, unexpected outputs, data leakage, actions outside approved scope and indicators of prompt injection or manipulation.

Vulnerability and patch management: Supporting infrastructure, libraries, models, APIs, integrations and supplier services must be covered by appropriate vulnerability management and patching processes.

Change and version control: Material model versions, prompts, configuration, data sources, connectors, integrations, approval conditions and review dates must be recorded where relevant. Major updates must be reassessed before release where they change the risk profile.

Incident response and recovery: AI-related incidents shall be managed through the University's existing incident management arrangements, with additional playbooks where necessary for prompt injection, unauthorised data exposure, unsafe automation, harmful AI-generated content, impersonation, AI-assisted fraud, authorised-user misuse or supplier AI incidents.

11. Data protection and IP

11.1 Data handling

Users must not enter personal, special category, confidential, sensitive research, credential, security, commercially sensitive or legally privileged information into unapproved AI services.

Approved AI use must define what data can be used, where it is processed, whether it is retained, whether it is used for model training, who can access it, and how it will be deleted or exported at exit.

Data minimisation must be applied. Where feasible, use synthetic, dummy, anonymised, aggregated or pseudonymised data rather than identifiable or sensitive data.

11.2 DPIA requirements

A Data Protection Impact Assessment (DPIA) is required where AI processing is likely to result in high risk to individuals, particularly where it involves personal or special category data at scale, profiling, monitoring, decisions affecting individuals, vulnerable groups, sensitive student or staff data, research participant data or other high-risk processing. The use of AI, experimental methods or novel research techniques does not automatically require a DPIA where no personal data or high-risk processing is involved. Where appropriate, a department, programme or platform-level DPIA may be used to cover routine activity, with project-specific addenda required only where an individual use case introduces materially higher risk.

11.3 Intellectual property and research outputs

AI use must not disclose unpublished research, proprietary methods, confidential business information, source code, contractual information or third-party IP to unapproved services.

Where AI is used to generate content, code, research material or teaching material, users must consider ownership, licence terms, attribution, originality, plagiarism, copyright and contractual restrictions.

12. Human Oversight

AI outputs must be reviewed by an appropriately accountable human before being relied upon for material decisions, formal communications, advice, assessments, recommendations or actions that could affect individuals or University operations.

The accountable owner must define who reviews outputs, what checks are required, when escalation is needed, and when AI outputs must not be used.

AI systems must not be treated as authoritative where outputs are probabilistic, incomplete, biased, unverified, unsupported by evidence or outside the approved use case.

Where AI-generated or AI-assisted content is used for teaching, research, assessment, formal communications, advice, reports or decision support, users should be transparent about its

use where appropriate, verify outputs before reliance, and make clear where outputs remain unverified or should not be treated as authoritative.

For high-risk use cases, human oversight must be meaningful rather than ceremonial. Reviewers must have the time, competence, context and authority to challenge or reject AI outputs.

13. Assurance, Review and Ongoing Monitoring

Approval of an AI use case is not the end of governance. AI system owners must ensure approved use cases remain appropriate as models, suppliers, integrations, data sources, outputs or user behaviour change.

High-risk AI use cases shall be reviewed periodically to confirm that approval remains valid, use remains within scope, data handling and access remain appropriate, controls and human oversight remain effective, supplier arrangements remain acceptable, and incidents or material changes have been addressed.

Medium-risk use cases may also be reviewed periodically where required by the approval conditions, assurance plan or risk profile.

AI-related assurance may make use of existing evidence, including supplier assurance reports, penetration testing, privacy assessments, audit findings, access reviews, monitoring outputs, incident records, DPIAs and research ethics documentation.

14. Training and awareness

AI risk awareness and acceptable use training must be embedded into existing security awareness, data protection and information governance programmes.

Training and awareness should cover appropriate use of approved AI tools, data that must not be entered into unapproved services, human review and accountability, output risks, AI-enabled phishing, impersonation and reporting routes.

Short-form user guidance, practical examples, FAQs and “dos and don’ts” should be maintained separately from this standard to support consistent and confident adoption across the University.

15. External and Internal AI Threats

AI-enabled cyber threats are managed through the University’s existing cyber security, monitoring and incident response arrangements. This section summarises the main AI-specific considerations and should be read alongside the wider Information Security Management System.

Controls should be applied proportionately and may include threat intelligence, Security Operations Centre (SOC) monitoring, strong authentication, least privilege access, data loss prevention, user awareness, supplier assurance, guardrails for AI systems and incident response playbooks for AI-assisted phishing, impersonation, prompt injection, unauthorised data exposure, unsafe automation and misuse by authorised users.

16. Exceptions

Where it is not possible to apply or enforce any part of this standard, an exception request must be raised through the appropriate University route.

The request must explain the requirement affected, the business, teaching or research justification, the risks created, compensating controls, duration and accountable owner.

The Head of Information Security will review the security risk. The Head of Privacy, relevant business or academic owner, and other stakeholders shall review the exception where their areas are affected.

Exceptions shall be time-bound, documented and subject to review.

17. Non-Compliance

Non-compliance with this standard may expose the University to information security, data protection, legal, academic integrity, contractual, financial, operational or reputational risk.

Suspected or confirmed breaches may be handled through relevant University procedures, including disciplinary procedures, depending on their nature and severity.

18. Supporting Procedures and Guidance

This standard defines the University's minimum governance and security requirements for AI. Detailed operational material, checklists and examples shall be maintained separately so that this Standard remains concise, accessible and focused on key requirements.

Supporting materials may include user guidance, dos and don'ts, AI use case assessment forms, approval workflows, the AI use case register operating procedure, delegated approval matrix, AI guidance, teaching and research guidance, DPIA guidance, supplier assessment guidance and incident response playbooks.

These supporting materials should be reviewed regularly to reflect changes in technology, risk, regulation and University-approved tooling.

19. Review and Maintenance

This standard and relevant supporting materials shall be reviewed at least annually, or sooner where there are material changes to technology, regulation, risk appetite, University operating model or AI-related threat landscape.

The review shall ensure that the standard:

- Remains operationally fit for purpose
- Supports responsible innovation
- Reflects current AI technologies and usage patterns
- Aligns with information security, data protection, procurement, research ethics and risk management requirements
- Supports continued regulatory, contractual and legal compliance
- Remains proportionate and practical for teaching, research and professional services