Policy Type: Both internal and external

University of East Anglia Norwich

Artificial Intelligence Governance Policy

Date approved:	19 November 2026
Approving body:	Senior Information Risk Owner
Date of next review:	November 2027
Review frequency:	Annually
Policy/Regulation Owner:	James Bentley *Associate Director for Information Compliance (& Data Protection Officer)*

If you require more information about this policy/regulation, please contact the Governance Office by emailing Governance@uea.ac.uk – The Governance Office will direct your query to the relevant team or individual.

1. Overview and Purpose

AI can create new opportunities, improve productivity and support different ways of working, researching, and delivering education, including learning, teaching and assessment.

The University promotes the ethical, legal, and responsible use of AI to enable innovation, create new opportunities and overcome challenges in the pursuit of public benefit, and the advancement of education and research.[1]

The University also recognises that failing to make appropriate use of AI may carry strategic, educational, research and operational risks. This Policy therefore provides an enabling baseline for responsible AI use. It sets clear requirements where control or assurance is necessary and should be applied alongside practical guidance, training, digital fluency and accessible support.

The purpose of this Policy is to establish the University's authoritative baseline requirements for the lawful, responsible, ethical, secure and effective use of AI. It expands upon the Information Risk Management Policy and Acceptable Use Policy and should be read alongside other applicable University policies, regulations, and domain guidance.

This is the University's principal AI policy. It is supplemented by user-focused domain policies and guidance for research and innovation, education, and administration. Additional domain policy should be created only where a clear case is made that formal policy status is necessary.

Mandatory requirements applicable to a particular domain may be set out in this Policy, another approved University policy or regulation, or formally approved guidance where the University's governance arrangements give that guidance mandatory effect.

The University will maintain an accessible central AI resource that brings together this Policy, related requirements and guidance, approved tools and permitted uses, training, support, and routes for advice or approval.

Any queries relating to this Policy should be raised via IT SupportHub. Any queries relating to the processing of personal data should be directed to the Information Compliance Team, by emailing DataProtection@uea.ac.uk

[1] UEA's vision as stated in UEA Strategy 2030 and objective of the University as stated in our Royal Charter.

2. Scope

This Policy applies to anyone using AI at, or on behalf of, the University of East Anglia. It should be read and understood by all users including:

- all staff – including contractors - employed or working for UEA
- all students – processing information on behalf of the University
- any individual who has, by virtue of their role or relationship with the University, any degree of access and/or use of Personal Data the University holds.

This Policy does not establish how AI capability development will be coordinated with the University's wider digital-fluency, information literacy and professional development provision.

Adherence to approved policy and regulations is fundamental to the effective operation of the University. This policy has been developed in alignment with sector best practice to promote consistency, accountability, and compliance with relevant standards. Observing the policy ensures that decisions and actions are informed, equitable, and aligned with institutional values.

3. Definitions

Term	Definition
Artificial Intelligence (AI)	A term for a range of technologies, tools, and approaches that often attempt to mimic human thought using algorithm-based technologies.
Agentic AI	An artificial intelligence system that aims to accomplish specified goals with limited human supervision. It is comprised of a machine learning model(s) that replicate human decision-making to solve problems.
Automated Decision Making	A determination made automatically without human involvement.

Term	Definition
Confidential Data	Information that is operationally valuable. It may contain personal data, but must not contains special category personal data.
Confidential-Sensitive Data	Information that is operationally more valuable than Confidential; contains sensitive or special category personal data; or consists of a large volume of Confidential information.
Generative AI	Any type of artificial intelligence system that identifies patterns and structures in data / information / material and generates content, including: audio, code, images, text, simulations, and videos in response to instructions ('prompts'). Generative AI tools can be trained on large language models (LLMs). These models are the algorithmic basis for generative AI tools such as ChatGPT, Claude, Copilot, Gemini, and Grok.
Information Asset	A collection of any type of data, irrespective of type or format that is processed for a specific function and has shared risks and ownership. The asset is the data / information held, not the format / system it is stored within.
Information Asset Owner (IAO)	The person who acts as the principal authority and has overall responsibility for an information asset and how it is managed.
Personal AI Account	Accounts, log-ins or licenses created by an individual in a personal capacity rather than through the University.
Personal Data	Any information relating to an identified or identifiable living person such as – but not limited to - a name, identification number, contact details, location data, image, or description.

Term	Definition
Processing	An operation or set of operations which is performed on data or information: • collection, recording, organisation, structuring or storage; • adaptation or alteration; • retrieval, consultation or use; • disclosure by transmission, dissemination or otherwise making available; • alignment or combination; or • restriction, erasure or destruction.
Secret Data	Information marked as 'secret' by HM Government or similar bodies.
Special Category Personal Data	Personal Data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person's sex life or sexual orientation. Low-volume Special Category Personal Data refers to individual or limited instances of Special Category Personal Data, rather than bulk datasets.
Unrestricted Data	Information that does not contain any personal data, that is not operationally sensitive, or is already legitimately in the public domain. Unrestricted Data cannot contain Personal Data, Special Category Data, business sensitive information or Intellectual Property.
Users	Any student, employee, worker, contractor, partner, franchisee or other individual who has, by virtue of their role or relationship with the University, any degree of access to the University's devices, information, data, files, or systems.

4. Roles and Responsibilities

The proper management of University data is the responsibility of everyone; we all have a role to play in keeping it safe and secure. This section sets out the responsibilities of all users, including those at UEA with additional responsibilities.

4.1 All Users

4.1.1 Any individual processing information on behalf of the University must to read and understand this Policy, the applicable relevant Operational Guidance set out in Section 7, and any other University policy, regulation or professional requirement relevant to the proposed use of AI.

4.2 Staff with additional responsibilities

4.2.1 Council

Council oversees the University's strategic risk exposure and receives assurance that material AI risks are being governed effectively. Council is responsible for:

i. Approving the University's overall risk appetite
ii. Receiving assurance on material AI risks and opportunities
iii. Considering AI within institutional strategy and long-term sustainability.

4.2.2 Executive Team (ET)

The Executive Team is accountable for institutional adoption of AI. ET is responsible for:

i. setting the strategic priorities for AI
ii. resolving material conflicts involving benefit, cost, ethics or risk
iii. Ensuring that adequate resources are available.

4.2.3 Senior Information Risk Owner (SIRO)

The SIRO provides leadership for information and AI risk, and has responsibilities for:

i. approving information risk management policies including this Policy
ii. reviewing and determining the outcome of any exemption requests to this Policy or its associated procedures
iii. reviewing and determining the outcome of any proposals for high-risk processing activities in which the IAO seeks to accept a risk against the advice of the DPO or CISO.

4.2.4 Chief Information Security Officer (CISO)

The CISO has responsibilities for:

i. advising the University on information security practices generally and in relation to specific applications – including AI tools

ii. co-authoring and reviewing this Policy

iii. supporting Information Asset Owners in the completion of third-party vendor reviews relating to AI use

4.2.5 Data Protection Officer (DPO)

The DPO has responsibilities for:

i. advising the University on data protection and records management practices generally and in relation to specific applications – including AI tools

ii. co-authoring and reviewing this Policy

iii. supporting Information Asset Owners in the completion of DPIAs relating to AI use

4.2.6 Senior Information Asset Owners (SIAOs) & Information Asset Owners (IAOs)

SIAOs and IAOs are responsible for:

i. identifying, owning and managing the information assets they oversee and understanding when these processes use AI

ii. overseeing the completion of required data protection, information security and information risk assessments where AI is introduced or used within the information assets they oversee, e.g, Data Protection Impact Assessments (DPIAs) when business sensitive data or high-volume Special Category Data is proposed to be input into an AI

iii. ensuring that AI-related information risks are managed and escalated through the University's existing information risk arrangements

iv. escalating any major or high-risk information processing matters to the CISO, DPO and / or SIRO.

4.2.7 Information Technology & Computing Service (ITCS)

ITCS has responsibility for:

i. the provision of Microsoft Copilot to end users
ii. monitoring the use of Copilot and other AI tools
iii. logging tickets with Microsoft and liaising with Microsoft to resolve a fault
iv. implementing and maintaining governance policies within the M365 Cloud at the direction of the Senior Information Risk Owner, IT Leadership Team, or Data Protection Officer.

4.2.8 AI Steering Group

The University's Executive Team (ET) has established an AI Steering Group which reports to ET. It does not provide specialist assurance or approval, accept information risk, or replace the responsibilities of established University decision-making and governance bodies. Its purpose is to:

i. maintain a strategic, institution-wide view of AI opportunities, risks, activity and dependencies
ii. bring together education, research and innovation, organisational development, information governance and technology perspectives
iii. identify recurring or systemic barriers, gaps between existing processes, significant initiatives that are being lost from institutional sight, and capacity constraints that are preventing responsible adoption or delivery
iv. advise ET on strategic priorities, material barriers, institutional risk-appetite questions and resource or investment choices, including where additional enabling capacity may be required
v. support clearer guidance and navigation into existing specialist assurance, digital governance and prioritisation routes
vi. promote connections between relevant functions while avoiding duplication of the responsibilities of SPMO, ITCS, the DPO, IAOs, the SIRO, procurement, legal, finance, accessibility and academic governance bodies.

5. Policy Statement

5.1 Human Accountability

5.1.1 A named university employee must remain accountable for every decision that is supported by AI.

5.1.2 The University prohibits the use of automated decision-making that would have an impact on a human.

5.1.3 AI may support human decision-making, provided that a named person retains responsibility and exercises meaningful judgement. Routine drafting, summarising or administrative support that does not materially determine an outcome is not, by itself, a consequential decision-making use.

5.1.4 AI supported decision making may be authorised on a case-by-case basis. Such use must only be conducted following the completion of a DPIA and approval from the relevant Senior Information Asset Owner.

5.1.5 AI-generated recommendations must not be accepted automatically where they could materially affect an individual's education, employment, access to services, rights, safety or reputation. The accountable person must review the relevant evidence, consider the limitations of the AI output and determine the decision. Appropriate reviewing activities are set out in Section 5.8 below.

5.1.6 AI may assist with marking or the drafting and structuring of student feedback only where the educator retains and exercises academic judgement, reviews the output, uses an approved tool and follows the applicable education and assessment guidance. AI must not determine marks or independently generate evaluative feedback.

5.2 Effective Governance

5.2.1 AI is a digital tool; its implementation and use follows existing digital governance processes.

5.2.2 Governance needs to be proportionate to potential harm. Low-risk productivity tools will follow an efficient approval route. High-impact systems will require enhanced assurance including but not limited to: vendor assurance checks and a DPIA.

5.2.3 AI risks will be assessed using the University's current risk management framework and scoring methodology. A separate AI risk-scoring system will not be used. A summary of the risk scoring matrix is included in Annex 1.

5.2.4 Formal risk assessments will be completed through the relevant University process.

5.2.5 The end-use of Copilot is not supported by ITCS. ITCS do not provide assistance on implementing Copilot, nor in resolving issues arising from its use. Functionality issues arising from the use of Copilot will not be resolved by ITCS. The use of such Teams Applications is done at the users own risk; users should consult with the relevant Team Owner and/or Information Asset Owner before adding such an App to a Team.

5.3. Data Privacy

When processing information using AI, like with any other software, the sensitivity of the information is paramount. Information must be handled in line with its Information Security Classification.

5.3.1 Unrestricted Data (Security Classification)

5.3.1.1 The University permits the processing of Unrestricted Data using any AI tool that does not integrate into the wider University IT infrastructure.

5.3.2 Personal Data or Confidential Information

Any processing of Personal Data must conform to the University's Data Protection Policy. Additionally, in relation to AI:

5.3.2.1 Where personal data is being processed, only the minimum necessary data required to achieve a task should be processed.

5.3.2.2 The University permits the processing of personal data by Microsoft Copilot when using a University-issued Microsoft365 licence.

5.3.2.3 Confidential information or Personal Data may be processed using Microsoft Copilot when accessed through a University-issued Microsoft 365 licence.

5.3.2.4 Personal data processed by the University must not be used to train AI models.

5.3.2.5 Where processing is related to a research project, such activities must be approved by the appropriate S-REC as part of the ethics approval process.

5.3.3 Special Category Personal Data

Any processing of Special Category Personal Data must conform to the University's Data Protection Policy. Additionally, in relation to AI:

5.3.3.1 Where personal data is being processed, only the minimum necessary data required to achieve a task should be processed.

5.3.3.2 The University permits the processing of individual and low-volume special category personal data by Microsoft Copilot for transcription and summary functions.

5.3.3.3 This permission includes the use of Copilot to summarise emails containing special category data, and for MS Teams to transcribe meetings at which special category characteristics are discussed. This permission includes the use of Copilot in administering staff and student disciplinary meetings.

5.3.3.4 This permission does not extend to bulk datasets or other large-volume processing of Special Category Personal Data. This includes but is not limited to: health data sets (i.e. staff absence data, Clinical-trials data, NHS patient-level data), EDI datasets, and Trades Union membership data.

5.3.3.5 Where processing is related to a research project, such activities must be approved by the appropriate S-REC as part of the ethics approval process.

5.3.3.6 Large-volume processing of Special Category Personal Data in any AI tool requires a Data Protection Impact Assessment and prior approval through the University's information risk processes.

5.3.3.7 Special Category Personal data processed by the University must not be used to train AI models.

5.3.4 Confidential-Sensitive Data

5.3.4.1 Confidential-Sensitive Data must not be processed using Personal AI Accounts or AI tools that have not been formally approved for that purpose. The use of such tools or licences for processing personal data must first be formally approved by the University. The use of such tools will be considered on a case-by-case basis.

5.3.4.2 The processing of Confidential-Sensitive information will normally require approval from the relevant Information Asset Owner.

5.3.4.3 Where processing is related to a research project, such activities must be approved by the appropriate S-REC as part of the ethics approval process.

5.3.4.4 Confidential-Sensitive data processed by the University must not be used to train AI models.

5.3.5 Secret Data

5.3.5.1 Secret data must not be processed using Personal AI Accounts or AI tools that have not been formally approved for that purpose. The use of such tools or licences for processing personal data must first be formally approved by the University. The use of such tools will be considered on a case-by-case basis.

5.3.5.2 The processing of Secret information must be escalated to the relevant Senior Information Asset Owner or the SIRO, according to the assessed risk.

5.3.4.3 Where processing is related to a research project, such activities must be approved by the appropriate S-REC as part of the ethics approval process.

5.3.5.4 Secret data processed by the University must not be used to train AI models.

5.4 Assessment and Approval

5.4.1 The following must follow the University's digital governance process before use:

i. a new AI tool that will process Personal Data, Special Category Personal Data, Confidential, Confidential Sensitive, or Secret information

ii. an AI-enabled system that will connect to, or be embedded within, the University's digital estate

iii. an AI capable of taking actions with limited human supervision

iv. a new integration between an AI tool and a University system or dataset

v. large-volume processing of Special Category Personal Data

vi. use of an AI tool other than Copilot to process Personal Data, Confidential, Confidential-Sensitive, or Secret information.

5.4.2 The proposed use will be screened to determine which existing assessments or approvals are required, which may include data protection, information security, IT architecture, supplier or procurement assurance, accessibility, research ethics or academic governance.

5.4.3 The use of publicly available or Open Data does not normally require a DPIA. However, research ethics review or other approval may still be required depending on the research question, methodology, reasonable expectations of individuals, potential for identification or harm, intellectual property considerations, contractual terms, the AI tool used and the intended outputs.

5.4.4 Where AI use is intended as part of a research project, this must be highlighted in the relevant ethics application and data management plan. Research ethics review will be determined using the University's normal ethics criteria; use of generative AI does not, by itself, create an automatic requirement for ethics approval. Where both ethics review and DPIA screening are required, the applicable guidance must explain their sequence and dependencies.

5.4.5 To commence such an assessment please raise an Information Compliance Support Request.

5.5 Security by Design

5.5.1 Security requirements must be identified before procurement, development, or deployment.

5.5.2 Security must be maintained throughout design, development, deployment, operation, change and retirement.

5.5.3 Research ethics review will be determined using the University's normal ethics criteria; use of generative AI does not, by itself, create an automatic requirement for ethics approval. Where both ethics review and DPIA screening are required, the applicable guidance must explain their sequence and dependencies.

5.6 Equality and Fairness

5.6.1 Where AI systems are being used to inform decision making – such as by reviewing data to flag concerns – they must be assessed for discriminatory, exclusionary, or unequal outcomes.

5.6.2 Testing must consider relevant groups, accessibility requirements, language, culture and differing levels of digital capability.

5.7 Transparency

5.7.1 UEA will be clear when users are interacting with AI or where AI materially influences a service or decision. Individuals have a responsibility to declare the use of AI when used in supporting decision-making, and in education, and research and innovation, endeavours.

5.7.2 The degree of explanation required will increase with the significance of the outcome. Where personal data is processed, or data is processed for a research project, this will include referencing the AI tool in the relevant privacy notice / participant information sheet.

5.8 Accuracy and Reliability

5.8.1 AI outputs must be treated as potentially inaccurate unless verified.

5.8.2 Verification methods must be proportionate to the consequences of potential errors; it may include the deployment of controls such as:

i. Authoritative source grounding
ii. Human review
iii. Confidence thresholds
iv. Structured testing
v. Citation validation
vi. Sampling and quality assurance
vii. Independent verification.

5.9 Academic Integrity

5.9.1 AI governance will help protect academic freedom while maintaining academic standards, research integrity and appropriate attribution.

5.9.2 Requirements for permitted, restricted and prohibited AI use in education, - including teaching, learning and assessment practices - must be in line with the Generative AI Policy For Teaching and Learning and clearly communicated to students and staff.

5.10 Sustainability

5.10.1 Material AI proposals should consider energy use, computational demand, supplier sustainability and whether a less resource-intensive solution could achieve the intended outcome.

6. Compliance and Monitoring

6.1 UEA undertakes regular and routine monitoring of compliance with this Policy.

6.2 Routine use that is authorised by this Policy does not require individual registration.

6.3 The University will maintain a record of:

- pre-approved AI tools and the parameters of such use
- approved AI tools and uses that are not pre-authorised by this Policy
- exemptions and authorised deviations approved by the SIRO; and
- identified unauthorised or non-compliant AI use.

6.4 Activities that deviate from this Policy – without a prior approved Exemption from the SIRO – must be reported to the CISO and, where Personal Data is affected, the DPO. Failure to report such activities also constitutes failure to adhere to this Policy.

6.5 Where actual or suspected deviation from this Policy is identified the SIRO, CISO, DPO, or their deputies, are authorised to instruct Users to cease further activities and / or take remedial action(s) to ensure compliance with this Policy. Failure to adhere to such directions constitutes failure to adhere to this Policy.

6.6 Failure to comply with this policy may result in disciplinary action, up to and including dismissal or termination of contract, in accordance with university HR and disciplinary procedures.

6.7 In addition to the standard review cycle, this policy will be reviewed sooner where:

- there is a significant incident or trend that changes UEA's risk exposure
- there are material changes to law, regulation, or sector expectations
- there are significant technology changes; or
- audit/assurance findings require policy amendment.

6.8 Significant or substantive changes to this Policy will only be approved by the Approving Body. However, the Policy Owner is authorised to make minor amendments to the Policy for the purpose of formatting, updating links, or correcting errors.

7. Related Documents

This is the authoritative University baseline for AI use.

The documents below should be read alongside this Policy and provide supporting requirements and/or audience-focused guidance. This guidance is considered to be extension of this Policy and carries the same mandatory compliance requirements.

Where guidance conflicts with this Policy, this Policy takes precedence, unless a higher legal, regulatory, contractual, funder or professional requirement applies.

Associated UEA Policies

- Generative AI Policy for Research and Innovation
- Generative AI Policy for Teaching and Learning
- Academic Integrity Policy
- Acceptable Use Policy
- Data Protection Policy
- Information Risk Management Policy
- Records Management Policy
- Information Security Classification Scheme
- Research Ethics Policy
- Risk Management Policy
- Freedom of Speech

Associated UEA Guidance for Staff

- UNIVERSITY OF EAST ANGLIA – RESEARCH ETHICS GUIDANCE NOTE
- Operational Guidance – Using AI to generate copy for Marketing and Recruitment

Associated UEA Guidance for Students

- Thinking skills for success at University and beyond | Rise 360

8. Version History

Version number	Approval date	Approval mechanism	Details of change
1.0	19/11/2026	SIRO	Initial version, no previous AI Governance Policy in place.

Annex 1: UEA Risk Scoring Matrix

The below is a summary of the University's risk scoring definitions that form part of the Risk Management Framework. Queries about the risk management framework should be directed to the Risk and Resilience Team.

<75%	**Very Likely (5)**	**5**	**10**	**15**	**20**	**25**
51-75%	**Likely (4)**	**4**	**8**	**12**	**16**	**20**
11-50%	**Possible (3)**	**3**	**6**	**9**	**12**	**15**
1-10%	**Unlikely (2)**	**2**	**4**	**6**	**8**	**10**
<1%	**Rare (1)**	**1**	**2**	**3**	**4**	**5**
		Insignificant (1)	**Minor (2)**	**Moderate (3)**	**Major (4)**	**Catastrophic (5)**
		<£10k	£10k-£100k	£100k-£1M	£1M-£10M	£10M+
		Isolated/ minor	Routine low-level	Short-term	Significant/ sustained	Total Sustained
		Minor internal deviation	Internal deviation	Possible legislative breach	Legislative breach	Significant legislative breach