University of Exeter - AI Policy

Version 1.2

Reviewer Group	Review Date
Audit & Risk Committee	27/03/2025
IT Operating Board	30/05/2025
IT Strategy Group	06/05/2025
Information Steering Group	30/04/2025
Compliance Committee	12/06/2025
PSDLT	10/03/2025
UEB	11/06/2025

1. Purpose of the policy and scope

1.1 Purpose

The University of Exeter's Artificial Intelligence (AI) Policy aims to harness the transformative potential of AI to enhance academic research, learning and administrative excellence. AI technology is an enabler for our Strategy 2030 work, as detailed in our AI Enablement Strategy. As per our terms of use policy we expect all our community to engage with AI in a responsible, legal, and safe manner, the policy seeks to:

- **Empower Innovation:** Enable staff, students, and researchers to leverage AI tools and applications to advance education, streamline administrative processes, and drive groundbreaking research. The policy links to an agreed list of Generative AI tools and data levels to help our communities make the right choices when using AI.

- **Ensure Ethical Use:** Ensure that AI adoption aligns with the University's core values, upholding fairness, transparency, and accountability in all AI-related activities. The University will ensure AI adoption complies with the University's legislative environment, such as education laws and regulations, GDPR, global AI regulations, etc. AI should not replace or override the decisions or insights of our people. It is a supplementary tool, not a primary decision maker.
- **Manage Risks:** Mitigate risks associated with AI, including bias, misinformation, hallucinations, data security, and privacy concerns, while encouraging its use in a manner that safeguards the University's reputation and compliance with applicable laws and regulations.
- **Promote Collaboration:** Facilitate interdisciplinary collaboration across Professional Services, Faculties, student groups and external partners to integrate AI solutions into the University's broader digital transformation strategy. The policy links to training materials and support functions. In the future as our AI community grows staff leads on specific techniques and toolsets will be highlighted in the AI Catalogue.
- **Prepare for the Future:** Equip the University community with the knowledge, tools, and governance structures needed to responsibly navigate the rapidly evolving landscape of AI and Generative AI technologies. Our approach will equip our community for their AI journeys within the University of Exeter whilst also supporting their personal AI development. The policy will contribute to the University's efforts to digitally transforming the student lifecycle to enhance both academic and administrative excellence.

1.2 Scope

This policy applies to all AI systems, tools, and applications used, developed, or procured by the University, including those used by staff, students, administrative staff, as well as external collaborators when working on university-led projects. This policy should be used in conjunction with existing university policies and guidance, particularly those related to Copyright, Intellectual Property (IP), Ethics and Academic Conduct and Practice policy.

The principles for responsible use of AI detailed in this policy should be considered and followed by all staff utilising AI in any area of their work. This policy however is not intended to restrict use of AI for research, or research efforts into new AI systems.

*When utilising Generative AI specifically, researchers are asked to make use of the tools provided in the AI Catalogue (and associated Information Classification Schedule) wherever possible. T*he university accepts however, that researchers may need access to a broader suite of AI tools and approaches for research purposes. If researchers require new Generative AI tools to be added to the AI Catalogue, please contact your IT Partner for further advice.

Generative AI systems are split into three groups:

- Direct-use AI systems:

- Technologies providing a Generative AI capability (e.g., Copilot, ChatGPT)
- AI embedded in other systems:
  - Systems with Generative AI components (e.g., chatbots in a room booking solution, data mining in storage solutions)

- AI used in education and research:
  - Creating in-house Large Language Models (LLM) using open-source tools, and local data sets
  - Utilising tools like Azure AI Studio and AWS SageMaker for Gen AI outputs

It is important to note that the data processed across those three groups of Generative AI systems can be very different. This policy links to the University's guidance on the Information Classification Schedule level allowed for specific use cases and tools.

The policy will address both the opportunities and risks across the following areas:

Education and Learning:

- Enhancing teaching methodologies, such as personalised learning tools, automated grading systems, and AI-driven course content recommendations.
- Supporting students with AI-powered tutoring, accessibility tools, and academic guidance systems.
- Enhancing student experience, by making use of innovative tools to improve students' learning, enable them to be creative and personalise their educational journey.

Research and Innovation:

- Driving innovation through Generative AI applications across the research lifecycle; from hypothesis generation, grant writing and literature review, to interdisciplinary collaboration, data analysis, simulation, output creation and dissemination. Safeguarding research integrity by promoting ethical data usage, avoiding biases, and adhering to regulatory standards.

Operational Efficiency:

- Streamlining operations by automating routine administrative tasks, allowing staff to focus more on strategic activities, and enhancing overall operational efficiency across the university.
- Managing data and processes securely to enhance operational resilience and decision-making.
- Improving business development practices through support with marketing.

Ethical and Legal Compliance:

- Provide assurances on usage of AI and its accuracy in line with the relevant laws and ICO expectations.
- Provide, additional security controls to protect both personal and commercially sensitive data.

- Ensure any newly configured AI algorithms are checked against introducing bias or harm to individuals.
- Completion of relevant documentation to meet legislative requirements, such as DPIA.
- Ensuring compliance with relevant laws, such as the General Data Protection Regulation (GDPR), and emerging AI legislation (e.g. EU AI Act).
- Embedding principles of Responsible AI, focusing on fairness, accountability, transparency, and sustainability.

Exclusions:

- This policy does not cover detailed technical implementation guidelines, which are addressed in supporting documentation and practical guidance.
- Non-AI-related technologies or systems governed by other University policies.
- The policy does not cover personal use of AI for non-university data sets or non-university owned equipment

2. Sustainable AI

2.1 Responsible AI within our University Culture

To ensure a sustainable use of AI within the University of Exeter, our policy suggests the following principles are to be applied to properly consider impact on affected individuals:

Human-centric and socially beneficial: The University should prioritise AI systems that support human decision-making, ensuring AI remains a tool for empowering users, rather than a substitute for human expertise. AI applications must be designed and procured with accessibility in mind, ensuring that systems benefit users equally, and do not reinforce social inequalities or exclude groups of individuals. Diverse user testing will be embedded into the development and procurement process for AI tools, ensuring inclusive representation and usability across demographic groups and accessibility needs.

Fair and unbiased: The University will conduct regular bias audits and equality impact assessments against protected characteristics on AI models deployed or procured, and for all significant AI applications or implementations, particularly where AI interacts with students or staff in decision-making contexts, to detect and mitigate algorithmic bias in decision-making processes.

This includes auditing content provided to LLMs (e.g., support articles, knowledge bases, feedback datasets) to ensure it represents diverse perspectives, avoids perpetuating stereotypes, and promotes inclusivity across academic and cultural contexts.

Transparent and documented: All decision-making processes based on the output of AI must be transparent and clearly document functionality, data sources and any known limitations.

Secure and safe: AI tools and systems will comply with robust cyber security protocols to prevent unauthorised access, data breaches and AI-driven cyber threats.

Accountable: The University should be accountable for ensuring that AI systems operate in compliance with the guiding principles.

Auditability and recourse: The University will ensure that AI processes are transparent and reviewable, and that individuals can challenge and appeal AI-supported decisions. By enforcing regular reviews of AI accuracy and compliance, and by providing clear mechanisms for human oversight and appeals, University will maintain principles of accountability, fairness, and transparency in all AI-related activities.

2.2 Managing the Environmental Impacts of AI

The environmental impact of AI is increasingly being recognised as a key sustainability concern, particularly due to the energy-consuming nature of training large AI models, data storage and water consumption levels. The University with its commitment to Net Zero and sustainability initiatives will consider these risks when integrating AI solutions into research, education and operations.

The policy askes that all AI enterprise agreements are procured within the context of sustainable operation and where possible selections are made on the prioritising low-energy AI models and opting for smaller fine-tuned AI models instead of training new models from scratch. This policy aligns with our procurement regulations and processes.

The University will help adopt a sustainable AI supply chain through looking into adopting the following four recommendations.

1. AI workloads will be run in data centres powered by renewable energy
2. Data Centres running University AI models will minimise energy consumption
3. AI models and tools chosen will minimise the data movement between data centres
4. Our selected AI models will do run heavy processing outside of core power hours to minimise peak demands

All users are urged to apply AI thoughtfully, using it where it adds genuine value, and to avoid default or excessive use. To support informed decision-making, the AI Catalogue will link to sustainability insights for suppliers, where possible. Researchers and Developers should actively design AI systems to optimise code and model design to reduce compute loads and energy demand.

3. Definition of AI and use of AI across the University

This diagram describes the embedded layered structure of the AI ecosystem and clearly defines Generative AI as the key term in this documentation.

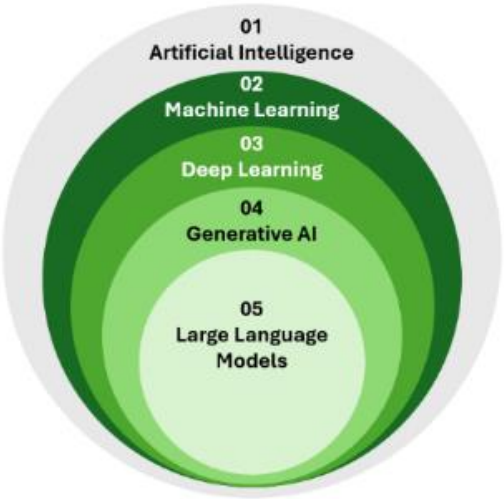

01 Artificial Intelligence (AI): A technology in which a computing system is coded to 'think for itself', adapting and operating autonomously. AI is increasingly used in more complex tasks, such as medical diagnosis, drug discovery, and predictive maintenance.
Example tool in the higher education sector: Perlego digital library - www.perlego.com

02 Machine Learning (ML): The set of techniques and tools that allow computers to 'learn and adapt' by creating mathematical algorithms based on accumulated data. Example technologies include spam filters in emails using words in subject line, predictive mathematical models.
Example tools in the higher education sector: Coursera ML-Based Recommendation System, Jisc Learning Analytics, Turnitin AI detection, etc.

03 Deep Learning (DL): A subset of machine learning where systems 'learn' to detect features that are not explicitly labelled in the data. Example technologies include facial recognition technologies, autonomous vehicle navigation, etc.
Example tools in the higher education sector: Otter.ai, Labster.

04 Generative AI (Generative AI): AI models that can create new content e.g. text, computer code, audio, music, images, and videos. Typically, these models are trained on extensive datasets, which allows them to exhibit a broad range of general-purpose capabilities. *Example tools are LLMs (see below) and Google Gemini, DALL-E (image generation).*

05 Large Language Models (LLMs): Type of Generative AI systems designed to learn grammar, syntax and semantics of one or more languages to generate coherent and context-relevant language. *Example tools are ChatGPT, Google Gemini, or Microsoft Copilot.*

→ Other definitions in GLOSSARY

4. Policy principles for AI usage

Actions to take when using AI

1. Don't fear AI usage or this policy seek help from AI Enablement, your IT Partner or Research IT team if you need support.
2. Read and comply with this policy and associated guidance before using AI tools
3. Complete the mandatory training on the universities online learning platform for information governance.
4. Use a toolset from the university AI Catalogue, making sure the information you plan to provide the platform meets the guidance given for the associated Generative AI tool
5. Read the Information Classification Scheme
6. Report any concerns and inappropriate AI content to Information Governance or your line manager

7. Clearly indicate when AI is in use either in meetings or outputs (documents, images, etc)
8. Test your AI results and look to confirm benefits, be ambitious but think critically, AI platforms can make errors
9. Maintain a record of AI usage and be able to share those records with your service director or department manager or other authorised company personnel upon request
10. Don't sign-up for new AI platforms, use the ones within the AI Catalogue or reach out to your Research IT team or IT Partner
11. Don't rely solely on AI for decision making, always apply human judgement
12. Don't share any information classified above the level associated with the tool you intend to use, please see the AI Catalogue for further information

5. Roles and responsibilities

It is the responsibility of everyone in our university that all information is handled and protected appropriately this includes any personal data or commercially sensitive information and includes any form of processing.

5.1 University of Exeter Community (staff, students, researchers, teaching staff)

All individuals making use of the University of Exeter's information assets are responsible for ensuring that they are aware of the requirements of the University's policies in relation to AI use, information governance, data protection and data security and adhere to them as they relate to the University's commitment to responsible AI use.

All users are responsible for highlighting areas of perceived risk, or where AI practices could be improved to their line managers and/or Information Governance teams.

All users are responsible for reporting any incidents that may involve a data breach, unauthorised processing of personal data, or violations of university policies or legislation. Reports should be made to the Data Protection Officer (DPO), Manager, or Senior Staff by contacting information-security@exeter.ac.uk or informationgovernance@exeter.ac.uk.

5.2 IT Services

- The policy will be managed by the Director of IT Services with the support of the IT Operating Board, with changes reviewed and approved by the Compliance Committee. The policy will be included within the policy audit schedule.
- IT Services are responsible for all AI security assurances, and the provision of the enterprise supported AI platforms (within the University's AI Catalogue)

5.3 Human Resources and Legal

- The polices connections and dependencies on HR and Legal terms and conditions will be managed according to their existing arrangements.
- Mandatory staff training will be considered in association with AI

- The General Counsel and Deputy Director of Compliance, Governance and Risk is responsible for the legalities of copyright ownership of AI and misconduct.

5.4 DPO/Information Governance

- The DPO is responsible for the protection of data for security and mitigation of AI bias.
- The DPO is responsible for reporting to the (Senior Information Risk Owner) SIRO all incidents, breaches, risks or circumstances that can create a risk to data or potential damage to the University's reputation
- Changes to the Information Classification Scheme need to be agreed by our DPO

5.5 Digital Transformation

- The policy will be cross-referenced in support of our AI Strategy.
- The Director of Digital Transformation is responsible for the upskilling of AI tools across the University.

5.6 Research

- Research leadership teams will use this policy in grant submissions and research groups ways of working.

5.7 Senior Information Risk Officer (SIRO)

- They are the chair of the Information Steering Group Chair and hold monthly meetings with DPO.

5.8 Education

- Monitoring Teaching Quality Assurance policies and practices related to student engagement and assessment.
- Creating and maintaining student and educator training on the appropriate use of Gen AI in teaching and learning.

6. Monitoring and reporting

The University retains the authority to oversee the utilisation of AI to guarantee adherence to this policy. Monitoring measures include:

- Total AI tool usage (by category: academic, research, admin, security, IT)
- AI-driven academic misconduct cases
- Number of AI research disclosures
- AI security incidents and data breaches
- Completion rates of AI training before usage of tools / mandatory AI / IG training
- Student and staff feedback on AI usability & risks
- Number of third-party AI vendor security audits completed
- Investment on AI systems and products within the AI Catalogue
- AI Gateway (future implementation)

- Carbon footprint of AI use in university infrastructure (to consider in the future when AI metrics are in place)

Our AI maturity will be measured every year using the NIST AI Framework

7. Training and awareness

7.1 Mandatory training

Prior to obtaining a licence and using an AI tool for the first time people should ensure that they have completed the mandatory training associated with Information Governance to confirm they will be adhering to university rules. Tools maybe removed if this training expires.

7.2 Continuous learning and digital upskilling

The policy will ensure that it provides the University community with a platform for continuous learning and digital upskilling, to further embed a culture of innovation across research, learning, teaching and operational activities.

As part of our commitment to digital inclusion, the University will develop differentiated support and learning pathways to reflect the diverse needs of our community. This includes tailored AI guidance for users with neurodivergence, individuals undergoing cultural transitions (e.g., international students or staff), and those requiring additional scaffolding to use AI confidently. AI training and support will incorporate inclusive design principles, and user feedback will be used to refine and extend available resources.

8. Policy updates

The University will adopt a dynamic review approach to ensure the policy remains up to date with the fast-evolving AI landscape. The policy will be reviewed as needed, with a minimum of annually, to remain aligned with rapid changes to relevant laws, contractual obligations and best practice.

9. Liability and indemnification

Users must take responsibility for their actions while using AI. Policy breaches may result in disciplinary action taken against the employee as outlined in our disciplinary policy.

Glossary

Academic integrity: The ethical and moral standards that govern academic work. It is the practice of honest and responsible conduct in all aspects of academic work, including research, writing, teaching, and learning. This means always giving full credit for the contributions of others to our own achievements (i.e. by full and correct referencing of borrowed words or ideas) and never falsifying the results of any research. (University of Exeter)

Cyber security: The protection of devices, services and networks – and the information on them – from unauthorised access, theft or damage. (NCSC)

Data privacy: The protection of personal data in a way that respects individual liberties and the free flow of data. (OECD)

Data Protection Impact Assessment (DPIA): A mandatory process under UK GDPR (Article 35) to identify and mitigate privacy risks before deploying AI systems that process personal data. The DPIA is required for any new system, new processing, collection of data, or changes of use, use of AI or data sharing. It is important that you contact IG/DPO early in the project bid or planning.

General Purpose AI (GPAI) [EU AI Act]: AI systems that have a wide range of possible uses, both intended and unintended by the developers. They can be applied to many different tasks in various fields, often without substantial modification and fine-tuning. Example of technologies include ChatGPT (due to its ability to perform many tasks such as text generation, translation, answering questions, etc.), Claude, DALL-E.

Infrastructure: Collection of the University's hardware, software and networks that it uses to run its technology.

Intellectual property (IP): something that you create using your mind - for example, a story, an invention, an artistic work or a symbol (UK Gov)

Our Network: Multiple computers and/or devices connected to shared resources and information. (NCSC)

Our environment: The collection of hardware, software, networks, systems, and digital services that the University of Exeter uses to support its academic, research, administrative, and operational functions. This includes on-premises and cloud-based infrastructure, cyber security frameworks, data storage solutions, digital learning platforms, and AI-driven tools.

Our systems: a database which is used to store data and present it to a user as information.

Use case: An AI application or the problem an AI system intends to solve.