Title	**Policy for using Generative AI in Research**: guidelines for researchers and professional staff
Version	1.0
Date	Aug 2025
Review timetable and process	Policy will be reviewed on a termly basis by the Policy Advisory Group
Audience	All researchers, including research students, and professional staff who support research
Service responsible for maintenance and review	Research Strategy and Policy Unit (RSPU), Research Services
Contact	Dr Lotte Boon, Research Planning Lead, RSPU

1. Introduction

1.1 Statement of intent

Generative Artificial Intelligence (GenAI) is a fast-evolving technology that offers tremendous opportunities and, as with all new technologies, comes with risks. The University of Oxford wishes to enable and support the safe and productive use of GenAI by the Oxford research community.

The aims of this policy are to ensure the responsible use of GenAI in research and to provide clear guidelines for its integration into the research process. **The policy is supported by, and should be read alongside, a list of FAQs for GenAI in research that provides good practice advice covering different parts of the research process.**

The University promotes adherence to relevant funder and publisher policies and supports the Concordat to Support Research Integrity. In preparing this policy, we have drawn on existing funder policies, and guidelines from funders, publishers, and scholars (a list of sources can be found at the end of this document), and where funder policies exist, we have referenced these.

The policy and FAQs were shaped, through roundtable meetings, by researchers who use GenAI in research and domain experts, working with the academic divisions and the Artificial Intelligence and Machine Learning (AI and ML) Competency Centre; the Research Practice Sub Committee (RPSC) reviewed the policy and FAQs.

Users are advised to interpret the use of GenAI in the context of research practice standards on transparency, rigour and respect; data protection, intellectual property and export control legislation; and information security and compliance guidance. Use of GenAI is also dependent on the fitness for purpose of the relevant tool in terms of functionality and performance. Users are advised to choose the tool that is most relevant for their discipline and application, and to be aware that both technology and practice will evolve over time.

Use of GenAI in research is within the scope of research conduct, and actions that are found to be in potential breach will be reviewed in line with the University's Code of practice and procedure on academic integrity in research.

The University has guidelines in place on using GenAI in communication, and a policy on the use of AI in summative assessment, setting out the expectations for those setting summative assessment and for students undertaking summative assessment.

1.2 Scope

The policy encompasses the **substantive use of GenAI in the research lifecycle**, particularly, in applying for funding, conducting research, disseminating and sharing research outputs, and review and assessment, but also its use in research management to enhance service provision and support

strategic insight. Researching GenAI, and training (configuring) and fine-tuning GenAI models, are out of scope of this policy.

Substantive use is defined as using GenAI in research to:

- interpret and analyse data (including audio and image files) and texts,
- undertake a literature review or translation,
- identify research gaps,
- formulate research aims,
- develop hypotheses,
- assist with generating ideas or to develop one's own thinking,
- generate code and synthetic data,
- undertake transcription of interviews, meetings etc. to save time and effort rather than overcoming barriers (see exclusions below)
- produce documents using other people's work or from transcripts or recordings.

Substantive use excludes using GenAI for tasks such as assisting non-native speakers with translation, transforming transcribed spoken language to written language by users who face barriers to writing, formatting documents, or improving one's own standard of English or a foreign language.

See section 5 for a glossary of other terms used in this policy.

This policy applies to all researchers, i.e. members of the University, including staff and students, who are conducting research, and those who are not members of the University but who are conducting research on University premises (including in overseas research units), or using University facilities or funding for their research, such as visiting staff, fellows, students and contractors. It also applies to Professional staff, University staff who in some capacity support research, e.g. research facilitators, technicians, IT support staff, library staff, and administrative support staff.

The policy also applies to all internal research funding schemes managed by Research Services, Divisional Offices, and departments.

2. Expectations from users of generative AI in Research

All researchers and professional staff supporting researchers are expected to meet the expectations below, and to comply with relevant funder and publisher policies, and professional guidelines. In addition, researchers are expected to follow the guidance in the GenAI FAQs document that sets out good practice in the research process.

Professional staff supporting researchers might be interested to use GenAI in order to simplify and improve the operational efficiency of processes, and the speed and quality of analysis; to accelerate the production of documents for different audiences, including summaries; to provide more accurate and targeted guidance to funding applicants.

Professional staff should consult the GenAI FAQs document and when supporting researchers, ensure they are aware of any restrictions to the use of GenAI in research.

All users should:

2.1 Remain ultimately responsible for GenAI content used in research

Users of GenAI tools are accountable for the integrity of the content generated by, or with the support, of these tools, maintaining a critical approach to using the output produced by GenAI and an awareness of the tools' limitations, such as hallucinations, or social biases that may be embedded in training data, which could perpetuate misrepresentation of social categories, protected groups, or historical inaccuracies.

2.2 Understand the implications of uploading material subject to Intellectual Property Rights

Users might be considering whether to upload published or unpublished work, i.e. work protected by copyright and related rights, to a GenAI tool for analysis, or a submitted manuscript or grant application, when acting as a reviewer. Uploading published or unpublished work may constitute a breach of copyright, if it cannot be covered by a copyright exception, e.g. non-commercial research and private study, or criticism and review. Users are advised to ensure the tool being used will not use the uploaded content for training purposes or finetuning the model, and are recommended to use a licence of an enterprise GenAI tool governed by an agreement between the University and the supplier; experienced users may prefer a locally deployed model. Further guidance is available in the University's guide to copyright or through emailing copyright@bodleian.ox.ac.uk.

2.3 Follow information security advice on using generative AI safely

While many popular generative AI tools are cloud-hosted services, locally deployed models may also be available for use. Cloud-hosted services present specific confidentiality and disclosure risks since data is accessible to the service provider and potentially other third parties. Information that is commercially sensitive and where the Intellectual Property is unprotected, or is subject to a confidentiality agreement or export control, or personal data, should only be uploaded to a generative AI tool when confidentiality can be guaranteed. Users should consult Information Security and Information Compliance advice, use a locally deployed model or a licence of an enterprise GenAI tool governed by an agreement between the University and the supplier, or contact the AI and ML Competency Centre for guidance. Uploading personal data is considered processing, needs to comply with data protection legislation and, when done as part of undertaking research, requires ethics approval.

2.4 Declare use of GenAI in research and communicate openly

When employing a discrete GenAI tool, rather than where GenAI is a functionality in existing software, users should declare the substantive use of GenAI tools in their work. Declarations could include the name, version, date, and how it was used and affected the research process. If relevant, users must make the input (prompts) and output available, in line with open research principles.

2.5 Aim for reproducibility/transparency and robustness in results and conclusions

Users should consider the probabilistic nature of GenAI tools, which includes the tendency to produce a different output from the same input. Cloud-hosted third-party models inherently present a reproducibility and transparency issue, as they are continuously updated with new versions, cannot be archived and do not make their training data available. Locally run models and fully open models offer a better alternative, where feasible and appropriate, to improve reproducibility and transparency. Users should disclose or discuss in methods or acknowledgements, or an additional section determined by the guidelines or policies of the funder or publisher, the limitations of, and interactions with, the GenAI tools used and mitigation measures, e.g. validating outputs, iterative prompting by systematically refining and adjusting the prompts to improve the quality of outputs.

2.6 Learn how to use GenAI tools responsibly to maximise their benefits, including by undertaking training

GenAI tools are evolving quickly, and new ways to use them are regularly discovered. Users should make use of the training, guidance and advice offered by the University (see section 3) and other reputable sources to stay up to date on good practice. Users might be interested in the University's Generative AI Special Interest Group (SIG) on Teams run by the AI and Machine Learning (ML) Competency Centre as a source for finding out about emerging practice and new releases, and to ask specific questions.

2.7 Be mindful of the environmental impact of using GenAI tools

Building GenAI models is extremely energy-intensive, while the energy consumed through the use of the resultant tools depends on the task's complexity and whether it involves image generation or

reasoning. Where feasible and appropriate, users are encouraged to favour smaller, locally run GenAI models for routine or low-complexity tasks; this can reduce the environmental and data privacy impacts of remote, cloud-based AI use. Users should consider whether a local AI model, a non-GenAI tool or intelligent automation, is fit for purpose before opting for large, cloud-hosted models.

3. Support provided by the University

The University will support users in meeting the expectations set out in section 2, by providing:

- Through the AI and ML Competency Centre, up-to-date and informed support, advice and information to users on all aspects of using GenAI in research, with more in-depth support for tools selected for their suitability in research, education and administration
- Through Research Data Oxford, advice and guidance on use of research data with GenAI
- Through the Advanced Research Computing service and the Biomedical Research Computing facility, advice on running locally deployed models
- Through the Compliance team in the Assurance Directorate, guidance and support to ensure compliance with data protection legislation
- Through the Information Security (InfoSec) team in the Assurance Directorate, guidance and support to ensure the use of GenAI does not pose a cyber-security threat
- Two major generative AI tools for purchase by departments and colleges: ChatGPT Edu, and Microsoft 365 Copilot, offering the ability to use GenAI for research and to collaborate with other users across the University, while protecting data used and generated
- Training to enable all staff and students to use the capability of GenAI tools in research and to promote best practice; training provided by the AI and ML Competency Centre and the Bodleian Libraries

Sources of related guidance and policy are listed at the end of this policy.

4. Policy implementation and reviewing mechanisms

This policy and supporting FAQs are published on the Research Practice pages, with links to the policy and FAQs on the Research Support webpages and the AI at Oxford website.

A policy advisory group will, on a termly basis, undertake a review of the policy, as the use of GenAI will change over time with certain uses becoming universal and new uses emerging; it will also review the FAQs, including commissioning the development of new FAQs and the removal of redundant FAQs.

Recommendations by the policy advisory group that do not materially alter the policy, e.g. additional clarification to sections 1–3, new levels of support and minor changes to existing levels of support (section 4), updates to glossary (section 6), list of existing University policies and guidance, and sources may be approved by the Co-Chairs for RPSC and submitted for noting to RPSC. Substantial changes will be reviewed by RPSC, which may refer the matter to Research and Innovation Committee for approval, or approve the change itself, in accordance with RPSC's Terms of Reference.

Policy advisory group membership: **Dominik Lukes**, Lead Business Technologist for AI and ML Competency Centre, Oxford e-Research Centre; **Chris Morrison**, Head of Copyright and Licensing, Bodleian Libraries; **Rowan Wilson**, Head of Research Computing and Support Services**,** IT Services**; Dr Sara Ratner**, PI of the AI in Education at Oxford University (AIEOU) Interdisciplinary Hub and post-doctoral researcher on the Learning in Families through Technology (LiFT) project, Education; **Professor Eleanor Stride**, Professor of Biomaterials, Engineering Science and the Nuffield Department of Orthopaedics, Rheumatology and Musculoskeletal Sciences; **Dr Hazem Zohny**, Senior Research Fellow in Practical Ethics, Neuroscience, and Society, the Oxford Uehiro Centre for Practical Ethics.

5. Glossary of terms

Fully open models	Class 1 models within the three-tier Model Openness Framework (MOF) are the gold standard in reproducible research practice; MOF has 17 components to fulfil completeness of model artefacts, covering code, data, and documentation that are part of the model development lifecycle as well as sharing that enables reproducibility, auditing, and downstream use. https://arxiv.org/pdf/2403.13784
Generative AI (GenAI)	Novel and rapidly developing AI models and tools that identify patterns and structures in data and generate novel content such as text, images and other media in response to instructions ('prompts').
GenAI model	A computational system trained on existing data that can generate new content or novel variations of existing content when provided with additional input. A model is created through training on existing data and documents. The learned patterns and relationships in a model are stored as model 'weights'. If these weights are available publicly then it is an open-weight model, such as Meta's LLaMa or OpenAI's CLIP models. If the weights are not available publicly then it is a closed-weight model, such as OpenAI's GPT-4 or Google's Gemini-2.5.
GenAI tool	The interface through which a user can interact with a GenAI model, sending instructions to the model and receiving generated outputs. Common examples include ChatGPT (powered by GPT models), which began as a text-to-text interface but now supports image-to-text and text-to-image capabilities. Similar tools include Claude and Meta AI's chatbots. • Tools can be hosted by a third party or by the maintainers of a model; ChatGPT is an example of a model hosted by the maintainers • HuggingFace is a popular website that hosts research, open-weight, and industry models; many models can then either be run locally or run in interfaces on HuggingFace or other third-party websites • It is possible in many cases to download and run models on one's computer or on University servers, i.e. a **locally run model**, and thereby creating a GenAI tool, but this depends on the size of the model, the specifications of the computer/server, and context-specific resources like technical know-how and internet hosting capacity.
Hallucinations	GenAI outputs that might initially appear to be believable but are in fact highly inaccurate or untrue, e.g. references that sound plausible but do not exist; a common characteristic of AI language models when producing responses based on limited or incomplete knowledge. Sometimes also termed confabulation.
Information Compliance	Means compliance with data protection legislation, i.e. any personal data needs to be processed in an open, accurate way and in accordance with the UK's data protection legislation.

	When using GenAI involves processing of personal data, you will be required to consult the data protection by design framework; if this processing is part of undertaking research, you should obtain research ethics approval. The data protection by design framework often includes the completion of a Data Protection Impact Assessment (DPIA): the processing of personal data is of a type identified by the Information Commissioner's Office (the ICO) as likely to result in a high risk to individuals.
Information Security	Means safeguarding the integrity, availability, and confidentiality of information assets within the University. In the context of information security, information is classified as either: public, internal or confidential. Where a GenAI tool is an application connected to a third-party system or a website, information security considerations come into play when data that is internal or confidential is transmitted. • The supplier of the GenAI tool should be on the University's third-party security register, which assesses the regulatory and data compliance of external services. New suppliers may require the completion of a TPSA (Third Party Security Assessment) • The University periodically engages in licensing agreements with third parties that will facilitate the secure sharing of documents and has agreements with OpenAI for ChatGPT Edu and with Microsoft for CoPilot for Microsoft 365 • Users who have a personal account on ChatGPT should not transmit any internal or confidential information to ChatGPT (such as research data), but obtain a ChatGPT-Edu account via the University and read the InfoSec guidance on using GenAI services safely
Large Language Models (LLM)	A type of machine-learning model designed for natural language processing tasks such as language generation; LLMs are language models with many parameters, and are trained with self-supervised learning on a vast amount of text, from a variety of sources, including public sources; feeding data or text into an LLM is likely to become part of its training data and may show up for other users of the model; LLMs are also used in symbolic systems studies that might not seem to be languages, e.g. LLMs trained on genomic datasets for genome analysis
Substantive use	Using GenAI in research to: • interpret and analyse data (including audio and image files) and texts, • undertake a literature review or translation, • identify research gaps, • formulate research aims, • develop hypotheses, • assist with generating ideas or to develop one's own thinking, • generate code and synthetic data,

	• undertake transcription of interviews, meetings etc. to save time and effort rather than overcoming barriers etc. • produce documents using other people's work or from transcripts or recordings. **Substantive use excludes** using GenAI for tasks such as assisting non-native speakers with translation, transforming transcribed spoken language to written language by users who face barriers to writing, formatting documents, or improving one's own standard of English or a foreign language.
Users	Collective term for researchers and professional staff.

Relationship with existing University policies and guidance

- Academic Integrity in Research (https://hr.admin.ox.ac.uk/academic-integrity-in-research)
- Research Ethics Policy (https://researchsupport.admin.ox.ac.uk/governance/ethics/committees/policy)
- Open Access Publications Policy (https://openaccess.ox.ac.uk/oapp)
- Research Data Management Policy (https://researchdata.ox.ac.uk/research-data-management-policy-full-text)
- Intellectual Property Policy (https://governance.admin.ox.ac.uk/legislation/statute-xvi-property-contracts-and-trusts) and Regulations (https://governance.admin.ox.ac.uk/legislation/council-regulations-7-of-2002)
- Environmental Sustainability Strategy (https://sustainability.admin.ox.ac.uk/environmental-sustainability-strategy
- Data Protection Policy (https://compliance.admin.ox.ac.uk/data-protection-policy/)
- Information Security Policy (https://www.infosec.ox.ac.uk/guidance-policy)
- Publication and Authorship Guidance (https://researchsupport.admin.ox.ac.uk/governance/integrity/publication)
- Export Control Guidance (https://researchsupport.admin.ox.ac.uk/policy/export)
- Open Research Statement (https://www.ox.ac.uk/research/support-researchers/open-research)

List of sources

Funder statement and policies

- Funders joint statement: use of generative AI tools in funding applications and assessment (September 2023)
- UK Research and Innovation (UKRI) policy on use of generative artificial intelligence in application preparation and assessment (September 2024)
- Wellcome policy on use of generative artificial intelligence when applying for Wellcome grant funding (2024)
- Cancer Research UK (CRUK) policy on the use of generative artificial intelligence tools in Cancer Research UK funding applications (December 2024)

Guidelines and guidance

- International Association of Scientific, Technical & Medical Publishers (STM) Ethical and practical guidelines for the use of generative AI in scholarly communications (December 2023)
- Guidelines for ethical use and acknowledgement of large language models in writing https://doi.org/10.1038/s42256-024-00922-7
- CRUK guidance for researchers on the use of generative AI (December 2023)
- European Commission's Living guidelines on the responsible use of generative AI in

This policy was written by Lotte Boon, Tanita Casci and the AI in Research Policy Coordination Group, consisting of authors: Johannes Abeler, Kanza Basit, Helen Carstairs, Ferdousi Chowdhury, Andrew Cusworth, Megan Gooch, Marina Jirotka, Dominic Lukeš, Chris Morrison, Mary Muers, Michael Parker, Sara Ratner, Matthew Reynolds, Martin Robinson, Irene Scullion, Eleanor Stride, Kelly Webb-Davies, Rowan Wilson, Elizabeth Wonnacott, and Hazem Zohny