

Oxford Policy on Responsible Use of AI

1. Purpose

The purpose of this policy is to set out how employees, workers and contractors may use artificial intelligence tools safely, responsibly and lawfully in the course of their work.

AI tools can support productivity, creativity, analysis and decision-making. However, they also create risks relating to confidentiality, data protection, accuracy, bias, intellectual property, security, over-reliance and hallucination.

This policy is designed to help us use AI in a way that supports the business while protecting our clients, employees, customers, suppliers and confidential information. AI tools must be used only where they have been approved by Oxford for the relevant business purpose, account type, workspace or configuration, and information sensitivity.

Approved AI tool means an AI tool, account, workspace, integration or use case that Oxford has authorised for use for a specified business purpose, subject to any applicable conditions.

Closed AI system means an AI environment with controlled organisational access and data-handling arrangements intended to limit access to Oxford-authorized users and systems. A closed system is not automatically approved for all purposes or all categories of information.

2. Scope

This policy applies to all employees, workers, contractors, consultants and anyone else working on behalf of Oxford.

It applies to the use of all AI tools, including but not limited to:

- Generative AI tools
- AI tools embedded in business software
- Automated decision-making tools
- AI-enabled recruitment, HR, performance, monitoring or analytics tools
- Image, audio, video or document-generation tools

3. General Principles

When using AI, we follow these principles:

- Use AI responsibly and professionally. AI should support our work, not replace our judgement.
- Protect confidential information. Do not enter confidential, commercially sensitive or personal information into an AI tool unless Oxford has approved that

tool, account and configuration for the relevant type of information and intended purpose.

- Use the correct Oxford account, workspace and permissions. Approval for a tool does not automatically permit every use of that tool or every type of information.
- Check AI outputs carefully. AI can hallucinate and produce inaccurate, misleading, biased or outdated information. Check its output as if it were the work of a junior colleague.
- Keep people involved. We do not make important decisions solely using AI.
- Be transparent where appropriate. Do not present AI-generated work as entirely your own if that would be misleading.
- Comply with Oxford policies. This policy should be read alongside our data protection, information security, confidentiality and privacy policies.

4. Permitted Use of AI

Subject to this policy, we may use AI to support a wide range of tasks. AI may support the production of work, but work shared externally must not be presented as wholly generated by Oxford where that would be misleading. Any AI-generated output must be subject to appropriate review, judgement and approval by people.

When we automate work, we consider the appropriate guardrails and risks. We allow time for review, judgement and rework where necessary.

We review AI-generated outputs before using them. The person using the AI tool remains responsible for the final work product.

5. Prohibited Use of AI

We do not use AI to:

- Enter personal data, employee data, client data, confidential information or commercially sensitive information into an AI tool, account or workspace that has not been approved for the relevant information type and intended purpose.
- Input passwords or login details
- Make final employment decisions, including recruitment, promotion, disciplinary, redundancy, performance or dismissal decisions
- Create discriminatory, offensive, misleading, unlawful or inappropriate content
- Impersonate another person without their express consent. For example, consent for digital twin avatars must be given through the approved platform
- Produce deepfakes or manipulated audio, images or video without clear authorisation
- Bypass security controls or Oxford approval processes
- Generate legal, medical, financial or regulatory advice without appropriate professional review

- Infringe copyright, trademarks or other intellectual property rights. See our Copyright Guidelines policy for more detail
- Conceal the use of AI where disclosure is required by Oxford, a client, a regulator or law

6. Confidentiality and Data Protection

We do not enter personal data, confidential information or commercially sensitive information into an AI tool unless the specific tool, Oxford account, workspace and configuration have been approved for the relevant type of information and intended purpose.

This includes special category data, commercially sensitive information, financial information, client documents, contracts, internal strategy documents and unpublished or proprietary materials.

We do not enter passwords, credentials or security information into an AI tool.

Where an AI system is described as closed, secure, private or enterprise-grade, this does not remove the need for Oxford approval, appropriate access controls, data protection assessment and appropriate review of whether a proposed use is appropriate.

Where we use AI to process personal data, we must ensure that its use complies with UK data protection law and US federal and state data protection laws, including requirements around lawful basis, transparency, fairness, security, accuracy and accountability.

7. Accuracy, Review and Accountability

We treat AI-generated content as a draft, not a final answer.

Before using AI-generated content, we check:

- Factual accuracy
- Tone and appropriateness
- Legal, regulatory or technical accuracy, where relevant
- Whether the output is biased or discriminatory
- Whether the output includes invented sources, statistics or references
- Whether the AI tool has used, retrieved or relied on the correct source material
- Whether any cited or summarised information can be traced back to an appropriate source
- Whether access controls, permissions and client-data separation have been maintained
- Whether information from one client, project or restricted workspace has been inappropriately included in another output

When using AI to search, summarise or analyse Oxford or client information, we use an approved Oxford account, workspace, data source and permissions setting. We must confirm that we are authorised to access, use and share the information for that purpose; we do not rely on the AI tool to make that determination.

Where the output relates to clients, employees, finance, legal matters, HR, health and safety, regulatory compliance or reputationally sensitive issues, it must be reviewed by an appropriately competent person before it is used. The reviewer must be able to assess the underlying source material and the basis for the final output.

8. AI in Recruitment, HR and People Decisions

We do not use AI as the sole basis for decisions affecting individuals. AI-supported people decisions must include meaningful review and decision-making by appropriately authorised people. We ensure that the basis for each decision can be explained and do not simply rely on AI output.

This includes decisions about:

- Recruitment and shortlisting
- Interview selection
- Promotion
- Pay or bonus decisions
- Performance management
- Disciplinary action
- Grievance outcomes
- Redundancy selection
- Absence management
- Dismissal

We take care to avoid unfairness, discrimination and bias. AI tools may reflect bias in the data they were trained on or in the way they are used.

9. Transparency and Disclosure

We are open about our use of AI where transparency is required or where non-disclosure could be misleading.

For example, disclosure may be appropriate where:

- AI has materially contributed to a report, proposal or client deliverable
- AI has been used to summarise or analyse information for decision-making
- AI-generated images, video, audio or synthetic content are used
- A client, regulator or business partner requires disclosure

Suggested wording:

“One or more AI tools approved by Oxford may have been used to support the preparation of this document. The final content has been reviewed and approved by the Oxford team.”

10. Intellectual Property and Copyright

We take care when using AI-generated content. AI outputs may unintentionally reproduce or closely resemble third-party material.

Before using AI-generated content externally, we consider whether:

- The content may infringe copyright or other intellectual property rights
- The prompt included confidential or copyrighted material
- The output is original enough for the intended purpose
- Oxford or the client needs ownership or usage rights

We do not use AI-generated content in client work, marketing materials, training materials or external publications without appropriate review by a member of the relevant client team.

11. Approved AI Tools

Oxford approves AI tools, accounts, integrations and use cases through its internal governance processes. Approval may be limited by role, business purpose, information type, account or workspace, technical configuration, location of processing, contractual terms, licensing, training requirements or other conditions.

The current record of approved AI tools and any applicable conditions is maintained separately by the relevant Oxford function. It is not reproduced in this policy because it may change over time.

A tool, account, integration or use case must be treated as **not approved** unless Oxford has confirmed that it is approved for the intended business purpose. Employees do not assume that approval of one AI tool, account, feature or use case extends to another.

Approval does not remove the need to comply with this policy, including requirements relating to confidentiality, data protection, security, intellectual property, accuracy, appropriate review and transparency.

Oxford may amend, suspend or withdraw approval at any time where circumstances, risk assessments, contractual arrangements, legal requirements, security considerations or business needs change.

12. Use of Personal AI Accounts

We do not use personal AI accounts, free trials, consumer accounts or personally managed subscriptions for Oxford work where this involves any Oxford, client, employee, supplier or other non-public business information.

For business purposes, we use the relevant Oxford-approved account, workspace and configuration. We do not use a personal account as an alternative to an Oxford-approved account, even where the same provider or model is used.

13. Record Keeping and Source Transparency

We do not need to retain a record of every individual AI prompt, chat or output. However, where AI has materially contributed to a client deliverable, significant business decision, HR-related matter or other high-risk work, we ensure that the final work can be explained and appropriately evidenced.

For client-facing summaries, analysis, reports or recommendations, we maintain a clear source record showing the key documents, data sources, meeting notes, slide numbers, pages or other materials relied on. The source record should enable a reviewer to verify the basis of the final output without needing access to every AI interaction.

Where AI has materially influenced a decision affecting an individual, or a significant internal business decision, the relevant decision record, meeting minutes or approval documentation must record:

- The role AI played in the work or decision
- The key source material or evidence relied on
- The review and approval undertaken by people
- The final decision or action taken

This requirement is proportionate to the risk and significance of the work. Where the appropriate record is unclear, we seek guidance from the Operations Team and/or IT.

14. Security

We follow Oxford's security requirements when using AI.

We do not use AI tools to:

- Analyse company systems without approval
- Generate or test code that could compromise security
- Process security credentials
- Bypass access controls
- Investigate vulnerabilities without authorisation

We stay alert to malicious AI-generated impersonations of colleagues or clients. If in doubt, we verify communications through an alternative, trusted channel.

We report any suspected data breach, security concern or accidental disclosure through an AI tool immediately to the Head of Business Operations.

15. Training and Support

We provide appropriate guidance and training on responsible AI use.

If you are unsure whether a tool, account, workspace, integration or proposed use is approved, or whether particular information may be used with it, seek guidance from the Operations Team and/or IT before proceeding.

16. Breach of This Policy

Failure to follow this policy may result in disciplinary action.

Serious breaches, including unauthorised disclosure of confidential information, misuse of personal data, discriminatory use of AI or deliberate concealment of AI use, may be treated as gross misconduct.

17. Policy Review

This policy will be reviewed regularly to reflect changes in technology, business use, legal requirements and regulatory guidance.