Guidance on Generative AI use for the Nuffield Department of Primary Care Health Sciences

Version: 0.80
Effective Date: [To be confirmed]
Review Date: [12 months from approval]
Approved by: [Departmental Strategy Committee – TBC]

What this guide does

This guide aims to help you use generative artificial intelligence (AI) tools safely and effectively in your work at the Nuffield Department of Primary Care Health Sciences (NDPCHS). Whether you're writing grants, analysing data, preparing meeting minutes, or teaching students, you will find practical advice for your specific situation.

This document sets out the official guidance for all departmental members on using generative AI. It establishes mandatory requirements alongside practical guidance to ensure responsible AI use across all departmental activities.

We recognise AI as a powerful tool that can enhance your work. This guidance ensures you can harness those benefits while protecting patient data, student data, and maintaining research integrity.

Who needs to read this

Everyone in the department. Even if you have been using AI tools for a long time, you need to be familiar with this guidance:

- All academic, research, and administrative staff, including those with teaching responsibilities
- Students at all levels (DPhil, MSc, undergraduate)
- Honorary contract holders and visiting researchers
- External collaborators using departmental resources

Contents

1 Five principles to guide every decision 4
1. You own the output 4
2. Be transparent 4
3. Protect data ruthlessly – match your tool to your data 4
4. Maintain academic integrity 4
5. Enable innovation responsibly 5
2 Quick tool selector 5
Step 1: What type of data do you have? 5

Step 2: If personal data, STOP 6
Current situation: 6
⚠️ **Don't confuse these:** 6
✅ Confidential but NOT personal: 6
❌ Confidential AND personal: 6
✅ Confidential but NOT personal: 6
❌ Confidential AND personal: 6
3 Critical: Personal data and AI – Current restrictions 6
4 Using AI in research 7
Before you start 7
What works well 8
What you must never do 8
Required disclosure 8
Funder requirements 8
Publisher and peer review requirements 8
Intellectual property (IP) and commercialisation 9
Special considerations for health research 9
5 Using AI in teaching and learning 9
Teaching governance and coordination 9
The decision hierarchy: 9
If you're teaching 10
If you're a student 10
The AI Assessment Scale (AIAS) 11
Fostering formative use and AI literacy 11
Resources: 11
6 Using AI in professional services 11
⚠️ **Important:** 11
Quick reference examples by function 12
Communications and engagement 12
Grants and project management 12
HR and administration 12
Finance and compliance 12
Executive support 12
7. Visual content and AI 12
Tier 1: Low risk 13

Tier 2: Requires greater care to review AI outputs 13
Tier 3: Unchecked outputs could mislead or risk causing harm (exceptional cases only) 14
Quick decision tree: 14
Never use AI to 14
8. Data security essentials 14
The golden rules 14
Special security considerations 15
Local models 15
Meeting transcription 15
If you suspect a data breach 15
9. Intellectual property and AI 15
Key principles 15
Practical guidance 16
Commercialisation and patents 16
Teaching and communications materials 16
10 Getting help 16
- AI use in summative assessment 17
11 Governance and compliance 17
Appendices 18
A. Disclosure templates 18
B. Further common scenarios and examples 18
D. Substantive use definition 19
Substantive AI use requiring disclosure includes: 19
Non-substantive use (no disclosure required): 19

1 Five principles to guide every decision

These five principles guide all AI use. When you encounter new tools or scenarios not explicitly covered here, apply these principles to make responsible decisions.

1. You own the output

AI can make mistakes, perpetuate biases, or misunderstand context. You are the expert in your field and are accountable for the accuracy, integrity, and final content – not the AI.

Verify all outputs for accuracy and ensure they meet the professional standards expected and that you are willing to put your name to them.

2. Be transparent

When in doubt, be transparent. People deserve to know when they're reading AI-assisted content, and funders and publishers require disclosure of "substantive use" (i.e., use that shapes the intellectual work, such as analysis or idea generation – see Appendix D for the full definition).

Always disclose AI use when:

- Publishing research or presenting findings
- Submitting work for assessment or review
- Creating formal or official content
- The audience would benefit from knowing

Disclosure generally not needed for:

- Routine internal emails and administrative tasks
- Personal productivity uses
- Accessibility support (e.g. speech-to-text for disabilities)
- Grammar and spell-checking of your own work

3. Protect data ruthlessly – match your tool to your data

Different AI tools have different data protections. Student data, patient data, and any personal information covered by regulations such as UK GDPR require the highest security and must only be used with University-approved tools.

Public information, as long as it does not include any personal data (e.g. names, locations of people, email addresses and so on), can be used with any tool. Never enter patient, student or any kind of confidential data or information into public AI tools.

4. Maintain academic integrity

AI cannot be an author – it cannot make original intellectual contributions or take responsibility for research findings.

Human expertise and judgment must drive research conclusions, teaching assessments, and critical analysis. This does not preclude the use of automated tools for objective marking (e.g. multiple-choice questions), but the design, validation, and oversight of such assessments remain a human responsibility.

Use AI to enhance, not replace, critical thinking.

5. Enable innovation responsibly

AI offers significant potential to advance health sciences research and education, but this power comes with broad responsibilities. As we innovate, we must consider the broader implications of our choices. These include:

- **Societal and professional impact:** Consider how our use of AI affects professional roles and skills, both within academia (e.g. research, teaching) and in wider society (e.g. creative professions). We should strive to use AI to augment human expertise, not simply to replace it without considering the cost.
- **Health equity:** Be mindful of how AI tools could perpetuate or worsen health inequalities through biased data or unequal access.
- **Environmental impact:** Recognise the significant energy consumption of large AI models.
- **Research and academic integrity:** Reflect on how these tools are changing our research practices and the skills we value in our students.

Experiment with AI within these guidelines and share successful, responsible approaches with your colleagues.

2 Quick tool selector

Step 1: What type of data do you have?

Data Category	Examples	Permitted Tools***
Public	• Published papers and reports, public websites	Any AI tool
Internal (no-personal data)	• Meeting agendas, draft policies, general admin	MS Copilot, ChatGPT Edu
Confidential (non-personal)	• Unpublished research, novel ideas, IP	MS Copilot, ChatGPT Edu, Local models*
Includes personal data**	Anything about identifiable people, e.g.: • Names, • email addresses • Meeting attendees • Student information • Patient records	Not permitted until dept DPIA complete, see section 3 below

Includes patient data**	• Health records, clinical information	Exceptional approval only and only with a DPIA

**Local models run on your own device, offering the greatest security for confidential data. You are responsible for their setup and secure use. See section 8 – Data Security for more details.*

***Requires a full Data Protection Impact Assessment (DPIA) and formal approval.*

****Please note that more* university approved *tools may be added over time. If this information conflicts with newer information, please contact IT to confirm.*

Step 2: If personal data, STOP

- Contact datasecurity@phc.ox.ac.uk
- Complete *a full Data Protection Impact Assessment (DPIA*) before proceeding – a high-bar process reserved for specific research projects and is not for general use.
- Consider if AI is necessary at all

Current situation:

The department is developing DPIAs for common AI uses. Until these are in place:

⚠ **Don't confuse these:**	
☑ Confidential but NOT personal:	Your unpublished research on lung function → Can use ChatGPT Edu / MS Copilot
✗ Confidential AND personal:	Interview transcripts from your qualitative study → Needs DPIA first
☑ Confidential but NOT personal:	Your novel statistical methodology → Can use ChatGPT Edu or MS Copilot
✗ Confidential AND personal:	Student dissertation drafts with feedback → Needs DPIA first

3 Critical: Personal data and AI – Current restrictions

The department has not yet completed Data Protection Impact Assessments (DPIAs) for AI processing of personal data.

What counts as personal data?

Any information relating to an identifiable person, including:

- Names (even just first names in context)
- Email addresses, phone numbers, or user IDs
- Job titles with names (e.g. "Professor Smith said...")
- Meeting attendees or participant lists

- Any combination of details that could identify someone (e.g. "the new research fellow in cardiac studies")
- Photos, recordings, or transcripts featuring people
- Student work, grades, or feedback
- Patient information (even if "pseudonymised" with codes)

Simple test: If someone could work out who you're referring to from the content, it's personal data.

This means:

- You currently CANNOT use any AI tool (including MS Copilot or ChatGPT Edu) for:
- Meeting minutes that include names or attributions
- Student work, feedback, or grades
- Staff performance discussions or HR matters
- Interview transcripts or research data containing identifiable information
- Any document that mentions specific individuals

This restriction applies even to university-licensed tools.

While these tools are *security-approved* for confidential data, they are not yet compliance-approved for personal data processing in our department.

What you CAN do now:

- Use AI for completely anonymised content
- Use AI for public information and your own work (with no personal data)
- Use the non-AI Microsoft Teams transcription feature for meetings
- Continue using traditional methods for tasks involving personal data

The department is developing processes / DPIAs for common administrative uses. Once approved, we will update this guidance with specific permitted uses and communicate them to the department.

Until then, if you need to process personal data for a specific research project, contact datasecurity@phc.ox.ac.uk to discuss initiating a project-specific DPIA.

4 Using AI in research

Before you start

Define what you are trying to achieve. In line with University research policy, you should first consider whether a locally run model is appropriate and feasible for your task, especially when handling sensitive data.

Document your process for transparency and rigour. To support research integrity, keep a log of substantive AI use in your research notebook / log, including the tool/version, key prompts, and the outputs received. This creates a valuable audit trail, even if the tool itself is not perfectly reproducible.

What works well

- Debugging code for statistical analysis
- Generating initial literature search strategies
- Improving manuscript clarity and readability
- Creating accessible summaries of published papers
- Initial thematic coding of fully anonymised transcripts (document how this shapes interpretation)

What you must never do

- Enter patient data into any AI tools
- Let AI write your discussion or conclusions, and so on
- Process confidential peer review materials
- Let AI interpretations override your analytical judgment

Required disclosure

In all research outputs, state:

- Which AI tools and versions you used
- The specific purpose (e.g. "language editing", "code debugging")
- Your validation and oversight process

Example: "We used [AI tool, version] to [specific purpose]. All scientific content, analysis, and conclusions remain the authors' original work and were subject to full human review and validation."

Funder requirements

Major funders have agreed a joint statement on AI use. Key points:

- AI tools cannot be listed as authors
- You must declare AI use transparently
- You are accountable for all content

Always check individual funder policies, as requirements vary. Research Services can provide guidance.

Publisher and peer review requirements

Individual journals / publishers may have their own specific guidance on the use of AI on publications or the peer review process.

You should familiarise yourself with these when selecting journals to publish in, or when being asked to peer review.

Intellectual property (IP) and commercialisation

The use of AI has significant implications for intellectual property, especially regarding authorship, inventorship, and commercialisation. Key points for researchers include:

- **Authorship:** AI cannot be listed as an author on a publication. You are responsible for all intellectual content.
- **Patents:** AI use can affect the patentability of an invention. It is critical to document the human intellectual contribution.
- **Protecting IP:** Never input unpublished manuscripts, confidential data, or potentially commercialisable ideas into public AI tools.

All researchers must read and adhere to the full guidance in Section 8: Intellectual property and AI. If your work may have commercial value, contact Oxford University Innovation (OUI) immediately.

Special considerations for health research

Clinical research: Ensure AI outputs align with clinical guidelines. Consider whether AI might introduce or worsen health inequalities.

Qualitative research: Document AI's role in your analytical process. Be aware that AI may miss cultural nuances. Ensure participant voice remains authentic. Even when anonymised, qualitative transcripts can contain sensitive or re-identifiable information. These should be treated as confidential data and only processed using secure tools, with local models being the preferred option.

Data sharing: Include AI use in your Data Management Plan. Archive prompts, modles, and parameters for reproducibility where possible.

5 Using AI in teaching and learning

AI offers significant opportunities in education – from helping educators to design curricula and create materials, to providing students with on-demand explanations and feedback. The University has adopted the Russell Group principles for AI use, emphasising integrity, appropriate assessment design, and clear guidance.

Teaching governance and coordination

Our teaching operates within a complex landscape. We deliver education through our own programmes and in collaboration with other departments, each with specific governance structures.

The decision hierarchy:

1. **Programme requirements come first** – Follow guidance from your assessment officer, course director, and Board of Examiners.
2. **Departmental standards apply where programme guidance doesn't exist** – Use this policy as your baseline.

3. **University policies set the boundaries** – All teaching must comply with the University's AI in summative assessment policy.

If you're teaching

For every summative assessment, you must:

- **Declare AI permissions clearly** – Specify whether and how students can use AI for each assignment. In line with University guidance, which presents several frameworks, this department has adopted the AI Assessment Scale (AIAS) as its standard where no programme-specific guidance exists.
- **Design assessments appropriately** – Review your assessments using the Centre for Teaching and Learning (CTL) triage tool to ensure they are robust and align with permitted AI use.
- **Ensure equitable access** – If AI use is permitted or required, all students must have access to appropriate tools (e.g. University-licensed ChatGPT Edu or MS Copilot).
- **Specify declaration requirements** – Tell students exactly how to acknowledge any permitted AI use.

Your responsibilities with AI-generated teaching materials:

- Always review, verify, and enhance AI-generated content before sharing it with students.
- Stay within the University's secure AI ecosystem for all student data and materials.
- Write personalised feedback yourself – AI cannot replace your expert judgment on student progress.
- Take full responsibility for all materials students see, regardless of how they were created.

A note on AI detection: The University has not endorsed any AI detection tools due to their unreliability. *You must not use them*. Design assessments that do not rely on detection. If you suspect misconduct, follow existing Academic Integrity regulations.

If you're a student

- **Check requirements for each assignment** – AI policies vary by module and assessment.
- **Ask if you are unclear** – Your instructor must specify which AIAS level (1–5) applies.
- **Declare all permitted AI use** – Include the tool name, version, prompts used, and how you modified the outputs.
- **Understand the consequences** – Unauthorised AI use constitutes academic misconduct.
- **Use AI to enhance your learning** – Treat it as a starting point, not an endpoint.
- **Consider the limitations** – Critically reflect on potential biases and inaccuracies in AI outputs.

See the University's student guidance on AI use for detailed support.

The AI Assessment Scale (AIAS)

This is our department's recommended framework for defining AI use in assessments:

1. **NO AI** – Complete the assessment without any AI assistance.
2. **AI PLANNING** – Use AI for brainstorming and initial research only.
3. **AI COLLABORATION** – Use AI for drafting with your critical evaluation and modification.
4. **FULL AI** – Use AI extensively, focusing on directing and critiquing its output.
5. **AI EXPLORATION** – Use AI creatively to solve novel problems.

See https://aiassessmentscale.com/ for more information on this scale.

Fostering formative use and AI literacy

Beyond assessment, we encourage using AI as a formative learning tool. This can include helping students to summarise complex papers, generate practice questions, or debug code, as outlined in the University's student skills guidance. The goal is to develop critical engagement, not just to prevent misuse. Support your students by:

- Teaching them how to evaluate AI outputs for accuracy and bias.
- Discussing the ethical implications specific to primary care health sciences.
- Preparing them for AI's role in their future careers.
- Sharing examples of good and poor AI use in your field.

Resources:

- CTL AI resources
- AI Assessment Scale

6 Using AI in professional services

> **⚠ Important:**
> Until department DPIAs are complete, no personal data can be processed with AI tools. The examples below assume no personal data is included.

Professional services staff can use AI to enhance efficiency, improve communication quality, and accelerate document production. The key is matching the right tool to your data classification. For these roles, the centrally supported MS Copilot and ChatGPT Edu tools are the primary recommended options.

Quick reference examples by function

Communications and engagement

- **Public information (any AI tool):** Social media post variations of public content; newsletter content from published sources.
- **Internal/Confidential (MS Copilot or ChatGPT Edu only):** Grant methodology sections; anonymised budget narratives; general project templates.

Grants and project management

- **Internal/Confidential (MS Copilot or ChatGPT Edu only):** Funding application drafts and reviews; budget narratives; project risk assessments; progress reports for funders. Compilation of department Standard Operating Procedures (SOPs) or policies.
- **Never:** Share details that could compromise competitive advantage or reveal confidential reviewer comments.

HR and administration

- **Public information (any AI tool):** Generic job description templates; generic job description templates; general recruitment guidance
- **Currently NOT permitted**: ANY processing of individual employee data, performance information, or specific recruitment decisions.

Finance and compliance

- **Internal (MS Copilot or ChatGPT Edu only):** Policy templates; general compliance frameworks.
- **Never:** Input actual financial figures, account details, or confidential budget information into any AI tool.

Executive support

- **Currently NOT permitted:** Meeting minutes, correspondence drafts mentioning individuals, briefings with names.
- **Special consideration:** When drafting on behalf of senior staff, always disclose substantial AI use in external communications.

7. Visual content and AI

The use of AI to generate visual content offers significant opportunities. To ensure this is done responsibly, we use a risk-based approach. The level of human oversight required depends on the context and potential impact of the visual.

Specifically for external communications representing the department (i.e. official materials such as brochures or annual reports, as well as website news or blog stories and social media posts), inline with the university guideline for communications using generative AI, we may publish images or videos fully generated by AI, but *only when they meet a specific need that*

cannot be met by using existing photographs or illustrations, taking or commissioning new ones or using human-created stock images, and fit into the tier structure below. We will always be transparent about our use of AI-generated images and must work to ensure copyright, IP and GDPR considerations are accounted for.

Tier 1: Low risk

For example: Image enhancement (upscaling, colour correction); background removal; simple abstract graphics.

Tier 2: Requires greater care to review AI outputs

For example: Data visualisation beyond standard charts; conceptual illustrations; infographics for presentations.

These should ideally include disclosure: "Generated with [tool] and reviewed for accuracy by the authors".

Example of a tier 2 usage:

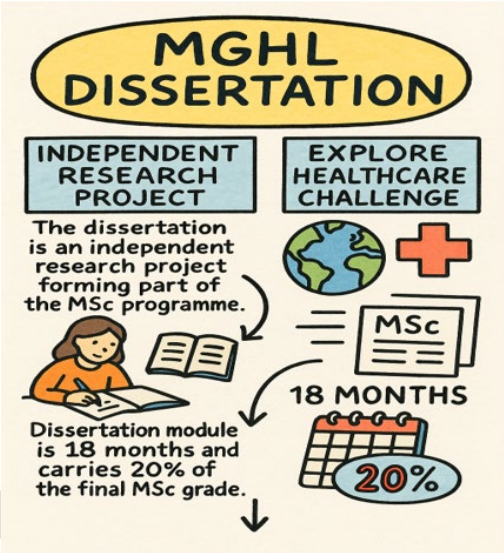

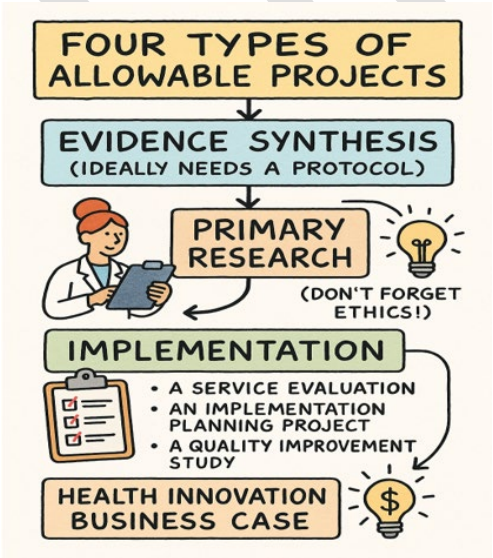

Tier 3: Unchecked outputs could mislead or risk causing harm (exceptional cases only)

For example:

- **Clinical accuracy critical:** Anatomical illustrations, medical device guides.
- **Data representation critical:** Disease hotspot maps, research participant personas.
- **Direct patient impact:** Patient education materials.

For both Tiers 2 and 3, you should consider including explicit disclosure in all cases.

Right: Example of the risks of an unchecked Tier 3 output.

Quick decision tree:

Could an error in this visual:

- Cause or lead to physical harm? → **Tier 3**
- Mislead about data or concepts? → **Tier 2**
- Only affect aesthetics? → **Tier 1**

Never use AI to

- Create or alter images of real people without their explicit written consent.
- Generate fake patient photographs.
- Produce deepfake videos or voice clones of individuals.
- Create content intended to mislead or misrepresent.

8. Data security essentials

Always follow the University's Information Security guidance on AI.

The golden rules

1. Classify your data first (Public / Internal / Confidential).
2. Match the tool to the classification.
3. When in doubt, use University-approved tools (MS Copilot, ChatGPT Edu) – but remember these still cannot be used for personal data without department DPIA approval. Never put patient or identifiable student data into public AI tools.

Special security considerations

Local models

Local models run entirely on your own computer (e.g. via tools like Ollama or LMStudio) or University servers, meaning your data never leaves your device. As such, they are an excellent choice for confidential or sensitive research data, offering enhanced security, reproducibility, and a lower environmental impact. The University encourages their use where technically feasible.

While local *models* do not require the same Third-Party Security Assessment (TPSA) as cloud services, **you are responsible** for ensuring they are sourced from reputable providers (e.g. Hugging Face) and configured securely. However, the software used to run them, e.g. Ollama or LMStudio may require a TPSA to be completed.

Be aware that local models are often less powerful than large commercial tools, are prone to the same limitations in terms of things like hallucinations, and often require technical setup. For advice on running local models, contact Advanced Research Computing (ARC) or Biomedical Research Computing (BMRC).

Meeting transcription

Current status: AI-powered transcription or summarisation (including MS Copilot) *cannot* be used for meetings until department DPIA is complete, as these inherently contain personal data.

Alternative: Use Microsoft Teams' built-in transcription feature (non-AI) for recording purposes only.

If you suspect a data breach

Follow the PHC Incident Reporting Policy:

1. **STOP** using the tool immediately.
2. **REPORT** without delay to: data.breach@admin.ox.ac.uk.
3. **CC:** Your line manager, the Information Asset Owner (IAO), and datasecurity@phc.ox.ac.uk.
4. **DO NOT DELAY** reporting to gather more information.

9. Intellectual property and AI

Guidance on intellectual property is likely to develop over the coming months and years. The below is primarily general guidance for what you may need to consider.

Key principles

- **AI cannot be an author or inventor** – Only humans can hold copyright or be named on patents. You must maintain clear records of human intellectual contributions to any work.

- **Protect others' IP** – Never input unpublished work, manuscripts under review, confidential research, or proprietary information belonging to others into public AI tools. This is a serious breach of academic and professional integrity.
- **Your AI outputs have uncertain legal status** – Copyright in purely AI-generated content is legally ambiguous. The strongest claim to copyright exists where there is substantial human creative input, direction, and modification.

Practical guidance

When reviewing others' work

- **Never** use AI for peer review (journals, grants, theses) unless explicitly allowed by the publication.
- **Never** input others' unpublished manuscripts or data.
- **Never** share committee materials that include personally identifiable information.
- If asked to use AI for peer review, decline and explain that it violates University policy.

For your own work

- Keep detailed records of your original ideas vs. AI suggestions, the prompts you created, and how you modified AI outputs. This documentation may be crucial for establishing human authorship and inventorship.
- Archive these records with your research data or project files.

Commercialisation and patents

- **Contact Oxford University Innovation (OUI) immediately if:**
    - Your AI-assisted work might have commercial value.
    - You are developing something potentially patentable.
    - Industrial partners are involved.
- **Critical:** AI use can affect patent eligibility. *Early* and transparent disclosure to OUI is essential.

Teaching and communications materials

- You hold copyright in your creative arrangement, pedagogical design, and original explanations.
- Document which elements are your original contribution versus AI-assisted drafts.
- Be cautious with purely AI-generated images, as their copyright status and licensing are often unclear.

10 Getting help

University policies on use of generative AI:

- Policy for using Generative AI in Research: guidelines for researchers and professional staff
- Use of generative AI tools to support learning
- InfoSec: Use Generative AI services safely
- AI use in summative assessment

For general AI guidance

- Departmental AI Ambassadors
- IT Support: it-support@phc.ox.ac.uk
- Information Governance: datasecurity@phc.ox.ac.uk
- Personal data and DPIA queries: datasecurity@phc.ox.ac.uk (check current DPIA status before processing any personal data)
- University AI Competency Centre

For specific scenarios

- Research ethics: CUREC
- Teaching: Department Director of Studies or CTL
- Data protection: Information Compliance Team
- Security assessments: InfoSec GRC Team

11 Governance and compliance

This guidance ensures compliance with key University policies, including Information Security, Data Protection, Research Integrity, Academic Integrity, and specific policies on AI in assessment, communications, and research. It also aligns with external requirements from UK GDPR, NHS Information Governance, funders, and publishers.

Consequences of non-compliance:

Non-compliance may result in additional training requirements, restricted access to AI tools, formal disciplinary procedures, and potential legal liability for data breaches.

Policy review:

This policy will be reviewed annually by the departmental AI Advisory Group.

Appendices

A. Disclosure templates

- **Publications**: "We used [AI tool, version] for assistance with [specific purpose, e.g. language editing and improving clarity]. All scientific content, analysis, and conclusions were generated by the authors and remain their sole responsibility."
- **Grant applications**: "AI assistance: [Tool] was used for [purpose]. All intellectual content and strategic decisions are the work of the named investigators."
- **Teaching materials**: "This material was developed with assistance from [AI tool] for [purpose]. All pedagogical content has been reviewed and validated by the instructor for accuracy and appropriateness."
- **External communications**: "This content was generated with the help of [AI tool] and carefully reviewed by our team for accuracy and tone."

B. Further common scenarios and examples

Research scenarios

- ✅ **Acceptable with disclosure:** Using ChatGPT to debug R code; asking AI to suggest alternative phrasings for manuscript clarity; creating plain-language summaries of your published papers.
- ❌ **Not acceptable:** Uploading patient interview recordings to any AI service; having AI write your discussion section; using AI to review others' confidential manuscripts or grant applications.

Teaching scenarios

- ✅ **Generally acceptable:** Creating practice problems based on learning objectives; generating diverse case scenarios for class discussion; drafting rubrics for assessment (with human modification and validation).
- ❌ **Never acceptable:** Using unapproved AI detection tools on student work; sharing student work with public AI tools; allowing AI to determine final grades.

Administrative scenarios

- ✅ **Acceptable:** Creating templates without names; drafting general policies; summarising confidential / public information
- ⚠️ **Awaiting DPIA approval:** Meeting minutes with names; internal emails mentioning individuals; any document with personal identifiers
- ❌ Never acceptable: Processing HR data; uploading student records; sharing confidential personal information.

Data handling scenarios

- ✅ **Public data (any tool):** Published research papers; open datasets; public websites.

- ⚠️ **Internal data (University tools only):** Draft manuscripts; meeting minutes; project planning documents.
- 🚫 **Confidential data (extreme caution):** Unpublished research data (local models preferred; University tools acceptable); personal staff/student information (only with explicit approval); patient data (never without full DPIA and IG approval).

D. Substantive use definition

Substantive AI use requiring disclosure includes:

- **In research:** Analysing or interpreting data, texts, or images; conducting literature reviews or identifying research gaps; developing hypotheses; generating ideas that substantively shape your thinking; creating code or synthetic data.
- **In other work:** Creating assessment materials; drafting official communications or policy documents; analysing institutional data for reports.

Non-substantive use (no disclosure required):

- Grammar and spell-checking.
- Improving clarity and flow of your own writing.
- Formatting documents.
- Translation to overcome language barriers.
- Accessibility support.
- Routine administrative tasks (e.g. "Draft a polite decline to this meeting invitation").

AI disclosure:
ChatGPT-5 was used to assist with document structuring, initial policy summaries, and drafting of example scenarios. All substantive content, policy interpretations, and departmental-specific guidance were developed by departmental staff. The AI outputs were extensively reviewed against original university policies, modified for accuracy and context, and validated by subject matter experts across research, teaching, and professional services.