AI Researchers

Embedding Responsible AI in Research



AI READINESS TOOLKIT

Practical resources for responsible AI adoption in research

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author or authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them.

CC BY SA



CONTENTS

1 About the toolkit

2 Literature review and synthesis

3 Structured drafting and reporting

4 Data preparation and content analysis

5 Academic writing

6 Paper presentations

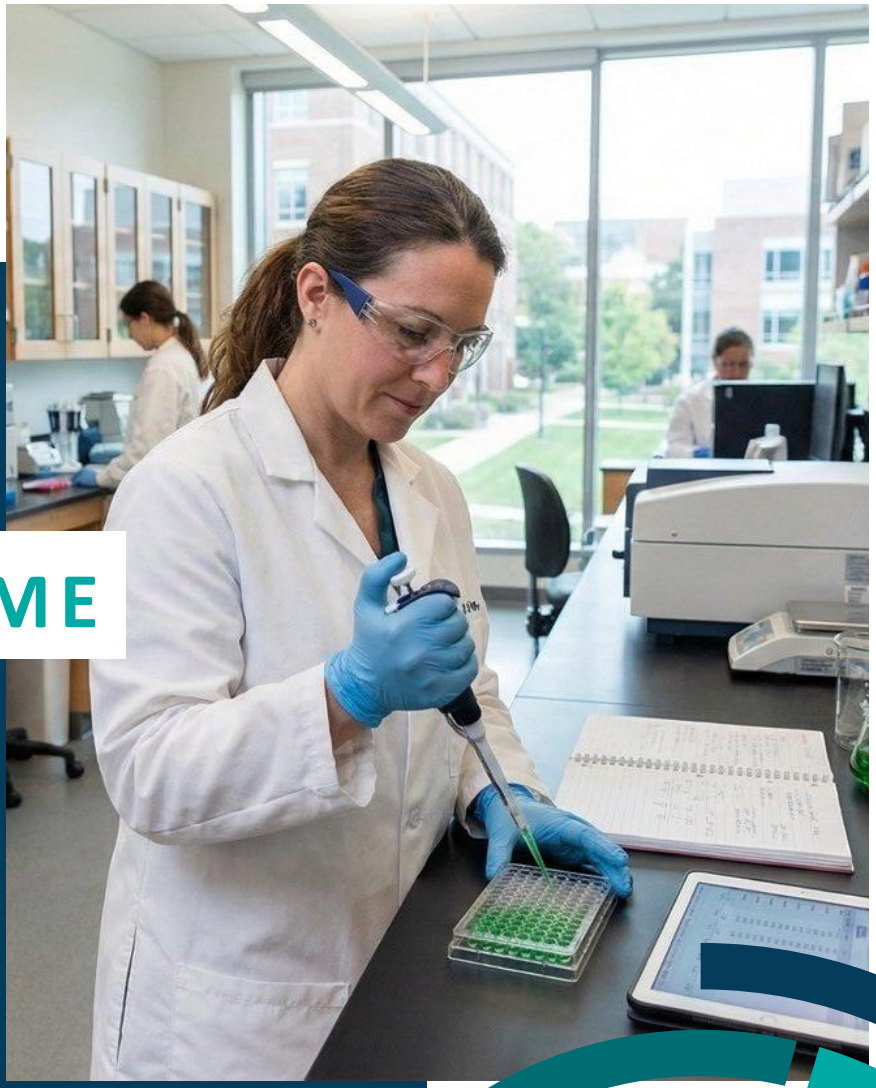
7 Peer review and refereeing

8 Quality, standards and the AI Researchers project

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author or authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them.

CC BY SA



WELCOME

A practical toolkit for responsible AI in research

The AI Readiness Toolkit brings together ready-to-use case studies, workshops, templates, peer-led discussions and a video that help academic researchers and supervisors adopt AI responsibly in everyday research practice.

It is written for professors, lecturers and research mentors who guide early-career researchers. Many resources are also useful to early-career researchers themselves.

Every resource presents AI as an enhancement to research, not a replacement.



Adoption is now a question of method and confidence, not principle.

The approach behind this toolkit

From Understanding To Certification

AI Researchers is an Erasmus+ Cooperation Partnership of seven European organisations working to embed AI literacy in research supervision and training. The toolkit is the second step in a simple chain.

Understand

Why researchers and supervisors hesitate, explored in the AI Engagement Framework.

Read the AI Engagement Framework ›

Show

What responsible use looks like in practice. This toolkit.

Teach

How to train others. An upcoming trainer manual.

Certify

Structured, assessed learning. A forthcoming online course.

Five barriers, five responses

The AI Engagement Framework names five barriers that hold researchers and supervisors back, pairs each with a practical response, and defines where human judgement must stay in control. Every resource is tagged with the barriers it addresses.

Trust in AI

Epistemic uncertainty

RESPONSE

Trust scaffolding

Protecting expertise

Identity and role threat

RESPONSE

Cognitive reframing

Know-how

Knowledge gaps

RESPONSE

Capability building

The rules

Institutional ambiguity

RESPONSE

Structural support and boundary setting

This task

Task-specific resistance

RESPONSE

Workflow embedding

How to use this toolkit

The toolkit is organised around research tasks rather than tools. Tools change quickly, the underlying tasks do not.

Start from your task. Find the use case that matches what you are doing, for example reviewing literature or preparing a grant report.

Pick a resource. Case studies, workshops, templates, discussion packs and a video. Every resource is free, openly licensed and ready to download.

Check the boundary. Each resource names the framework barriers it addresses and states the human-only boundary: the parts of the task that must stay with the researcher.

Every resource has a card on the toolkit page at **airesearchers.eu**. *The links in this guide open those cards.*



USE CASE 1 OF 6

AI can screen, cluster and summarise large volumes of literature and surface themes for the researcher to interpret, reducing the time spent on screening.

02

LITERATURE REVIEW AND SYNTHESIS



HUMAN-ONLY BOUNDARY

Every source and interpretation is checked against the originals. The final analysis stays with the researcher.

FRAMEWORK LINK

Know-how (knowledge gaps), met by capability building. A conditional domain.

IN THIS SECTION

UC1-01 From literature mapping to hypothesis development with generative AI

UC1-02 From summary to synthesis

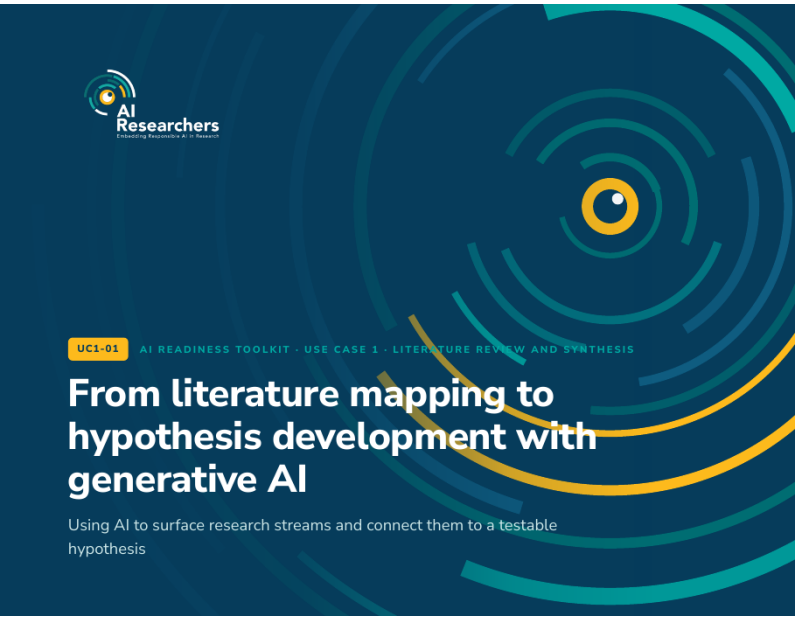
UC1-03 AI literature reviews: source enrichment and validation

From literature mapping to hypothesis development with generative AI

UC1-01 Case study

Use case 1: Literature review and synthesis

A researcher uses generative AI to map the literature around a topic, identify the most relevant research streams and link their findings to testable hypotheses. The case shows how to prompt for research streams, check every suggested source against the original and decide which evidence supports each hypothesis. AI speeds up the mapping. The researcher selects the sources, judges the evidence and owns the interpretation.



CONTRIBUTING PARTNER

UPF Barcelona School of Management

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Know-how, This task

HUMAN-ONLY BOUNDARY

Every source and interpretation is checked against the originals. The final analysis stays with the researcher.

OPEN THE RESOURCE ›

From summary to synthesis

UC1-02 Workshop

Use case 1: Literature review and synthesis

A 60-minute workshop that helps researchers move from AI-generated summaries to a traceable, balanced literature synthesis. Participants build a manual evidence map, audit each AI claim against its source, improve a prompt and rewrite a synthesis paragraph. The pack includes a facilitator guide with timings and a script, a 14-slide deck and four participant worksheets. AI proposes structures. Researchers authorise every interpretation.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

UPF Barcelona School of Management

FORMAT

Facilitator guide (PDF), slides (PPTX and PDF)

FRAMEWORK BARRIERS

Trust in AI, Know-how, Protecting expertise, This task

HUMAN-ONLY BOUNDARY

Source selection, assessment of evidence quality, interpretation, the final narrative and responsibility for claims remain with the researcher.

AI literature reviews: source enrichment and validation

UC1-03 Case study

Use case 1: Literature review and synthesis

A practical guide to finding, validating and adding sources that an author has not yet discovered. It shows how AI tools can match a draft paper against the published literature, suggest relevant work and help the author check that each new source is real, relevant and correctly cited. Use it when preparing or revising a literature review. Every source and interpretation is checked against the original.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

UPF Barcelona School of Management

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Know-how, This task

HUMAN-ONLY BOUNDARY

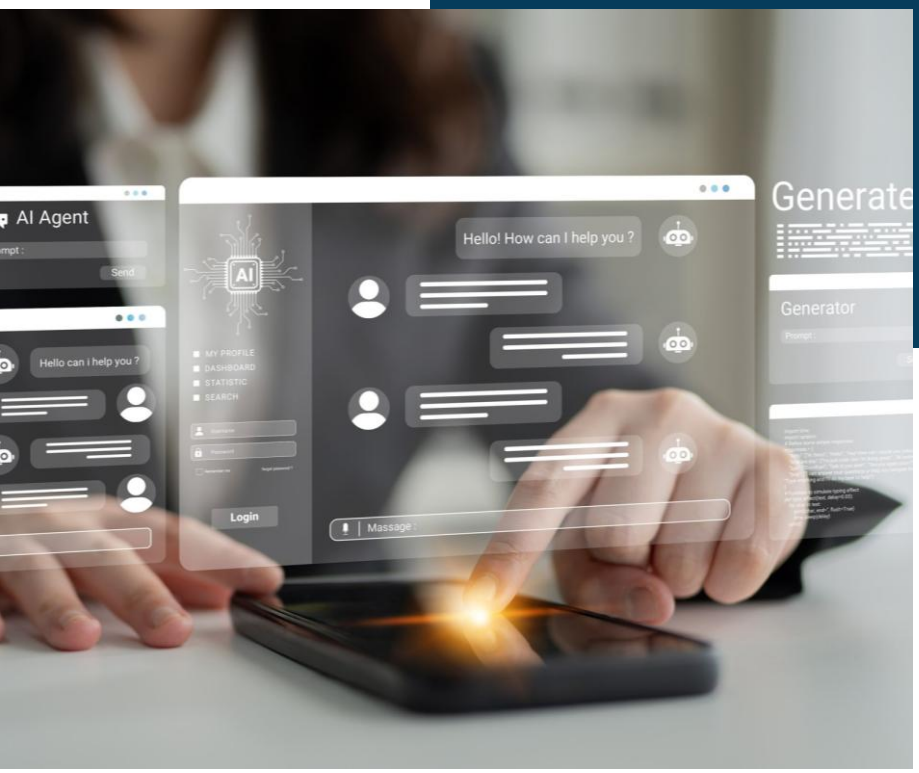
Every source and interpretation is checked against the originals. The final analysis stays with the researcher.

USE CASE 2 OF 6

For grant applications and project reports, AI can generate alternative structures, outlines and first drafts of administrative text.

03

STRUCTURED DRAFTING AND REPORTING



HUMAN-ONLY BOUNDARY

Authorship, claims and intellectual contribution stay human. AI supports structure, not substance.

FRAMEWORK LINK

Protecting expertise (identity and role threat), met by cognitive reframing. A high-acceptance domain.

IN THIS SECTION

UC2-01 From notes to a structured project report

UC2-02 When an AI-generated project report sounds convincing but is wrong

UC2-03 Responsible AI prompt template for grant and project reporting

From notes to a structured project report

UC2-01 Workshop

Use case 2: Structured drafting and reporting

A 60-minute workshop in which researchers turn verified project notes and reporting requirements into a structured report outline. Participants ask AI for alternative structures, draft selected administrative sections and then test every statement against the evidence. The workshop separates what AI can organise from what only the project team can claim. Authorship, claims and intellectual contribution stay with the researchers.



CONTRIBUTING PARTNER

University of Lodz

FORMAT

PDF

FRAMEWORK BARRIERS

Know-how, Protecting expertise, Trust in AI, This task

HUMAN-ONLY BOUNDARY

Authorship, claims and intellectual contribution stay human. AI supports structure, not substance.

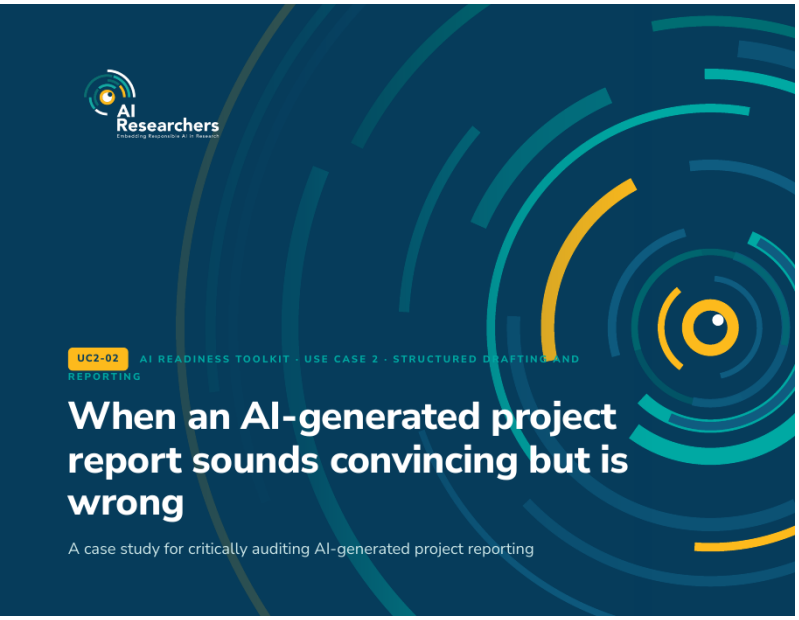
OPEN THE RESOURCE ›

When an AI-generated project report sounds convincing but is wrong

UC2-02 Case study

Use case 2: Structured drafting and reporting

A fluent, professional-looking AI-generated progress report is not necessarily accurate. In this case study, participants audit an AI-drafted report section for a fictional EU-funded project against the project notes it was based on. They test each statement against the evidence, then write a defensible version. The case shows how AI can turn provisional information into confident claims that the evidence does not support.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

University of Lodz

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Know-how, The rules

HUMAN-ONLY BOUNDARY

AI may help organise and phrase verified project information. Humans must verify facts, decide what can legitimately be claimed, interpret project progress, assess evidence of results or impact, and approve the final report.

Responsible AI prompt template for grant and project reporting

UC2-03 Template

Use case 2: Structured drafting and reporting

A reusable template for using AI on grant applications and project reports without handing over claims, evidence or intellectual contribution. It takes the user from preparing verified information, through a structured prompt, to reviewing the output and recording how AI was used. Checklists and worked prompts make each step quick to apply. It suits researchers, project managers and research support staff.



CONTRIBUTING PARTNER

University of Lodz

FORMAT

PDF

FRAMEWORK BARRIERS

Know-how, The rules, Trust in AI

HUMAN-ONLY BOUNDARY

Authorship, claims and intellectual contribution stay human. AI supports structure, not substance.

OPEN THE RESOURCE ›

USE CASE 3 OF 6

AI can format and clean data and support preliminary qualitative coding and transcript analysis, reducing the mechanical work of getting data ready for analysis.

04

DATA PREPARATION AND ANALYSIS



HUMAN-ONLY BOUNDARY

No unpublished or sensitive data in unapproved tools. Interpretation and conclusions stay human.

FRAMEWORK LINK

The rules (institutional ambiguity), met by structural support and boundary setting. High to conditional.

IN THIS SECTION

UC3-01 Data quality and audit

UC3-02 AI-assisted mixed-methods evidence integration

UC3-03 AI for hypothesis testing

Data quality and audit

UC3-01 Workshop, self-guided

Use case 3: Data preparation and content analysis

A self-guided 60-minute workshop on cleaning and preparing quantitative data with AI support. Using 17 fictional training records, participants build an issue log, ask AI to audit the data without editing it, agree a cleaning plan and run an independent check before analysis begins. A person traces and approves every change. The pack includes the workshop slides and the training dataset.



CONTRIBUTING PARTNER

FEP, University of Porto

FORMAT

Slides (PPTX and PDF), training dataset (PDF)

FRAMEWORK BARRIERS

Trust in AI, Know-how

HUMAN-ONLY BOUNDARY

No unpublished or sensitive data in unapproved tools. Interpretation and conclusions stay human.

OPEN THE RESOURCE ›

AI-assisted mixed-methods evidence integration

UC3-02 Case study

Use case 3: Data preparation and content analysis

A case study on combining survey results with qualitative evidence from interviews or focus groups. Using a fictional example, it shows how to prepare the material, ask AI for defensible analysis options, build an integration matrix and avoid conclusions the data cannot support. It suits researchers new to mixed methods. No unpublished or sensitive data enters an unapproved tool, and interpretation stays with the researcher.



CONTRIBUTING PARTNER

FEP, University of Porto

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Know-how, The rules, This task

HUMAN-ONLY BOUNDARY

No unpublished or sensitive data in unapproved tools. Interpretation and conclusions stay human.

OPEN THE RESOURCE ›

AI for hypothesis testing

UC3-03 Case study

Use case 3: Data preparation and content analysis

A supervisor uses a bounded AI methods referee to check whether a proposed statistical test fits the hypothesis and the study design. The case gives the prompt, shows how to question the AI's reasoning and how to record a transparent decision that a person approves. It helps supervisors and researchers discuss test choice with confidence. Interpretation and conclusions stay with the researchers.



CONTRIBUTING PARTNER

FEP, University of Porto

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Know-how

HUMAN-ONLY BOUNDARY

No unpublished or sensitive data in unapproved tools. Interpretation and conclusions stay human.

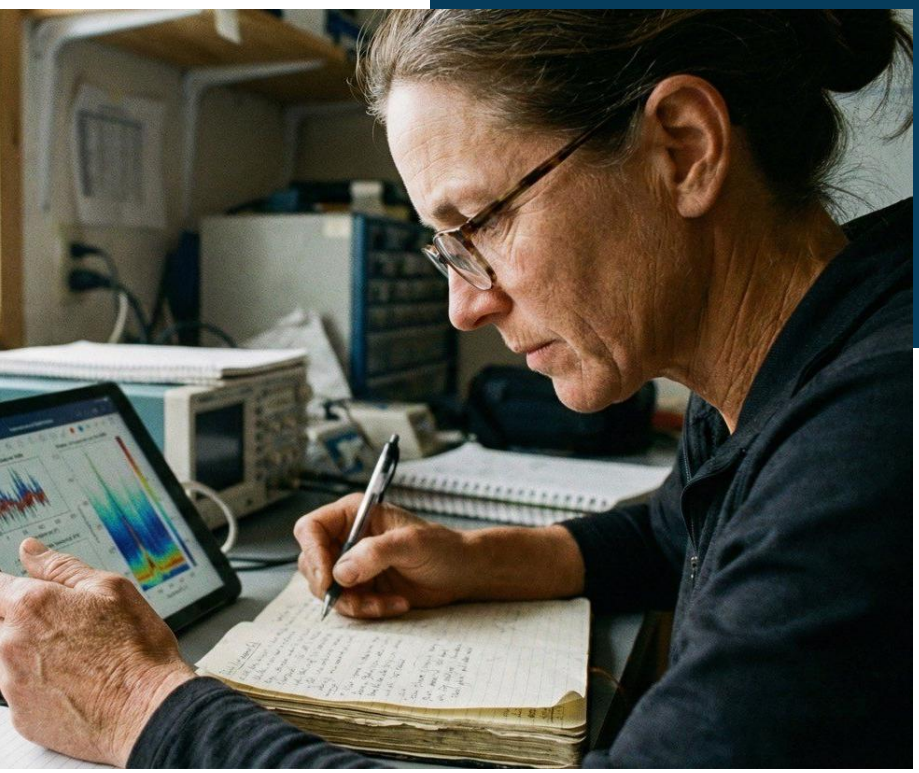
OPEN THE RESOURCE ›

USE CASE 4 OF 6

AI can refine the structure, clarity and language of academic text and suggest alternative phrasings for the author to accept or reject. This is especially useful when publishing in a second language.

05

ACADEMIC WRITING



HUMAN-ONLY BOUNDARY

Arguments, claims, original meaning and authorship remain the researcher's.

FRAMEWORK LINK

Protecting expertise (identity and role threat), met by cognitive reframing. A conditional domain.

IN THIS SECTION

UC4-01 AI prompt template library for academic writing

UC4-02 Accept, adapt or reject

UC4-03 Specialised AI tools for academic writing

AI prompt template library for academic writing

UC4-01 Template

Use case 4: Academic writing

Twenty ready-to-use AI prompts for seven stages of writing a research paper, from planning the structure to the abstract and final revision. Every prompt follows one principle: the researcher provides the knowledge and direction, AI assists with a defined task, and the researcher evaluates, verifies and decides. Placeholders make each prompt easy to adapt to any discipline. Arguments, claims and authorship remain the researcher's.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

UPF Barcelona School of Management

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Know-how, This task

HUMAN-ONLY BOUNDARY

Arguments, claims, original meaning and authorship remain those of the researcher.

Accept, adapt or reject

UC4-02 Template

Use case 4: Academic writing

A revision log for checking AI-suggested changes to academic text before accepting them. For each suggestion, the researcher records the original and proposed wording, checks meaning, evidence and voice, and decides to accept, adapt or reject it. A pre-use checklist covers approved tools, confidentiality and disclosure. The log gives a clear record of AI use and keeps arguments, claims and authorship with the researcher.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

UPF Barcelona School of Management

FORMAT

PDF

FRAMEWORK BARRIERS

Protecting expertise, Trust in AI, Know-how, The rules

HUMAN-ONLY BOUNDARY

AI may propose revisions. Arguments, claims, interpretation, evidential strength, contribution and final authorship remain with the researcher.

Specialised AI tools for academic writing

UC4-03 Case study

Use case 4: Academic writing

A stage-by-stage guide to purpose-built AI tools for academic writing: finding and reading literature, managing and checking references, improving language and choosing a journal. For each stage it explains what the tools bring, their risks and how to check what they produce. Tools covered include Semantic Scholar, Elicit, Zotero, scite and Writefull. Arguments, claims, meaning and authorship remain the researcher's.



CONTRIBUTING PARTNER

UPF Barcelona School of Management

FORMAT

PDF

FRAMEWORK BARRIERS

Know-how, Trust in AI, The rules

HUMAN-ONLY BOUNDARY

Arguments, claims, original meaning and authorship remain the researcher's.

OPEN THE RESOURCE ›

USE CASE 5 OF 6

AI can draft slide outlines, visual structures and summaries of research for the author to adapt.

06

PAPER PRESENTATIONS



HUMAN-ONLY BOUNDARY

The scientific content, claims and delivery remain the researcher's.

FRAMEWORK LINK

This task (task-specific resistance), met by workflow embedding. A high-acceptance domain.

IN THIS SECTION

UC5-01 Designing a research presentation with AI

UC5-02 Case 5.2, title to follow

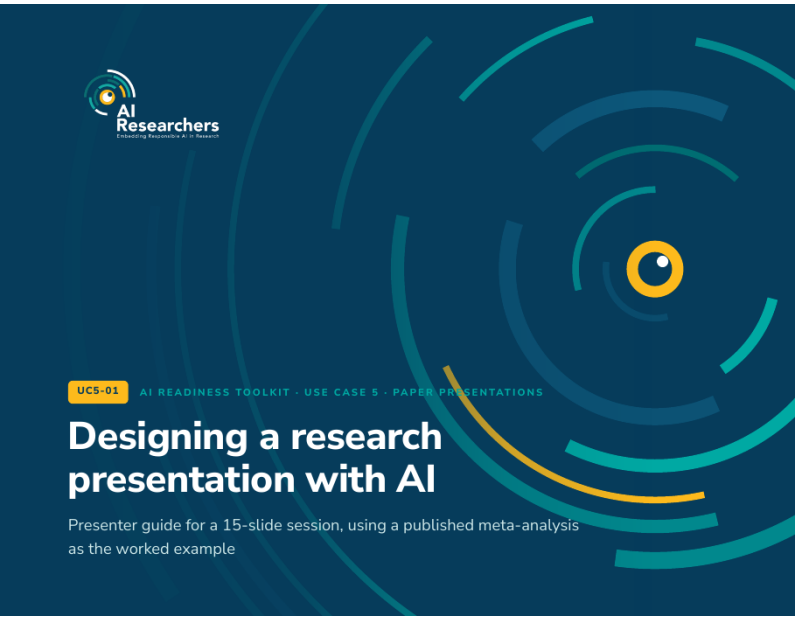
UC5-03 Would you use this AI-generated research presentation?

Designing a research presentation with AI

UC5-01 Slides with presenter guide

Use case 5: Paper presentations

A ready-to-present 15-slide session on turning a published paper into a research presentation with AI support. Using an open-access meta-analysis as the worked example, it shows how to request alternative talk structures, shorten slide messages and check every statistic against the source. A presenter guide gives the objective, cues and a full script for each slide. Content, claims and delivery remain the researcher's.



CONTRIBUTING PARTNER

University of Lodz

FORMAT

Slides (PPTX and PDF), presenter guide (PDF)

FRAMEWORK BARRIERS

This task, Trust in AI, Know-how

HUMAN-ONLY BOUNDARY

The scientific content, claims and delivery remain the researcher's.

OPEN THE RESOURCE ›

A research article is not yet a talk

UC5-02 Demonstration, video

Use case 5: Paper presentations

In this six-minute animated demonstration, a researcher prepares a seven-minute conference talk on an open-access meta-analysis. After writing a verified brief, they ask AI for three alternative six-slide structures, compare them, reject an unsupported headline and revise the chosen outline. The video shows why every statistic is traced back to the paper. The article is the evidence source, and the final interpretation belongs to the researcher.



WATCH THE VIDEO ›

CONTRIBUTING PARTNER

University of Lodz

FORMAT

Video (about 6 minutes), transcript (PDF)

FRAMEWORK BARRIERS

This task, Trust in AI, Know-how

HUMAN-ONLY BOUNDARY

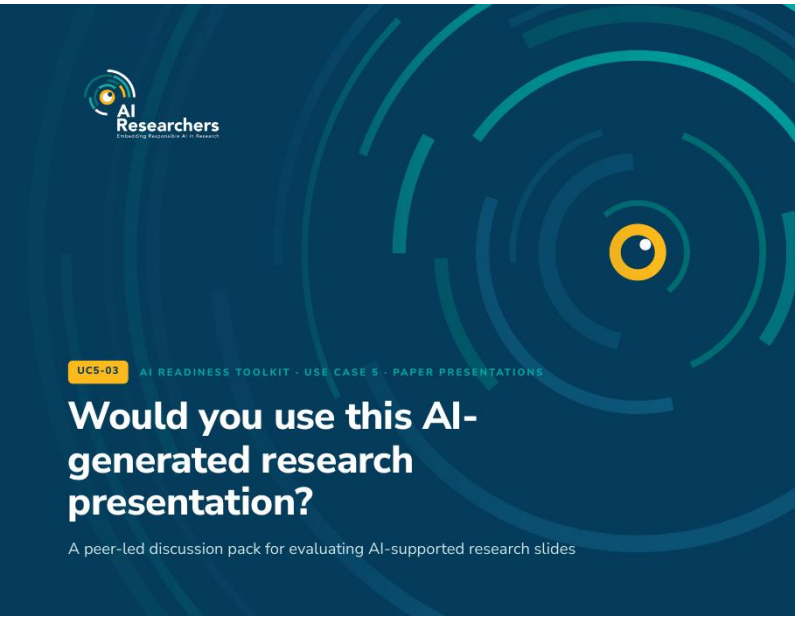
The article is the evidence source. The final interpretation belongs to the researcher.

Would you use this AI-generated research presentation?

UC5-03 Peer-led discussion

Use case 5: Paper presentations

A peer-led discussion built around an AI-generated slide deck that mixes accurate content with oversimplifications, unsupported claims and misleading visuals. In pairs or groups, participants review six slides against a short evidence summary of an open-access study and decide whether to retain, revise, reject or verify each one. The pack includes the sample slides, a review sheet and facilitator notes.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

University of Lodz

FORMAT

Discussion pack (PDF), sample slides (PPTX and PDF)

FRAMEWORK BARRIERS

Trust in AI, Know-how, This task

HUMAN-ONLY BOUNDARY

AI may help suggest slide outlines, summaries and visual structures. Scientific content, evidential claims, interpretation, limitations and delivery remain the researcher's responsibility.

USE CASE 6 OF 6

Within institution-approved environments only, AI can support the structuring of review comments and consistency checks. This use case also makes the limits of appropriate use clear.

07

PEER REVIEW AND REFEREEING



HUMAN-ONLY BOUNDARY

The assessment and recommendation are entirely human. Unpublished manuscripts never go into external tools.

FRAMEWORK LINK

Trust in AI and the rules, met by boundary setting and trust scaffolding. Conditional to prohibited.

IN THIS SECTION

UC6-01 AI referee agents for pre-submission review

UC6-02 AI-assisted major revision

UC6-03 A reflection on the ethical and practical use of AI in peer review

AI referee agents for pre-submission review

UC6-01 Case study

Use case 6: Peer review and refereeing

A doctoral researcher runs three independent AI referee agents over her own manuscript before sending it to her supervisor. One checks the conceptual contribution, one the methods and evidence, and one reads it as a journal editor would. The case supplies all three prompts, a decision log and a verification rule. The student and supervisor decide what to change and when the manuscript is ready.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

FEP, University of Porto

FORMAT

PDF

FRAMEWORK BARRIERS

Protecting expertise, The rules, This task

HUMAN-ONLY BOUNDARY

The AI agents may surface weaknesses, questions and revision options. The PhD student and supervisor verify the evidence, decide what to change, retain authorship, and decide when the manuscript is ready to circulate or submit.

AI-assisted major revision

UC6-02 Case study

Use case 6: Peer review and refereeing

A project-based example of how an approved AI assistant can help authors respond to a major revision decision. It turns reviewer and editor comments into a traceable register, tracks verified changes to the manuscript and supports a point-by-point response letter. Authors keep control of scientific judgement, additional analyses, which comments to accept or rebut, and the final decision to resubmit.



OPEN THE RESOURCE ›

CONTRIBUTING PARTNER

FEP, University of Porto

FORMAT

PDF

FRAMEWORK BARRIERS

Trust in AI, Protecting expertise, The rules, This task

HUMAN-ONLY BOUNDARY

Scientific judgement, additional analyses, interpretation of results, acceptance or rebuttal of comments, authorship and the final resubmission decision remain entirely with the authors.

A reflection on the ethical and practical use of AI in peer review

UC6-03 Testimonial, video

Use case 6: Peer review and refereeing

In this five-minute video, Artur Rodrigues, Full Professor at the School of Economics and Management of the University of Porto (FEP), reflects on how AI can assist peer reviewers. He covers practical uses, the risk of losing independent judgement, confidentiality and how research papers may need to change. The views are his own. In this toolkit, unpublished manuscripts never go into unapproved tools.



CONTRIBUTING PARTNER

FEP, University of Porto

FRAMEWORK BARRIERS

Protecting expertise, The rules, This task

FORMAT

Video (about 5 minutes), transcript (PDF)

HUMAN-ONLY BOUNDARY

The assessment and recommendation in peer review are entirely human. Unpublished manuscripts must never be entered into external tools, and many publishers prohibit it outright.

WATCH THE VIDEO ›

Quality and standards

Every resource meets the same standards, so it can be used and shared with confidence.

Open and reusable

Resources developed by the project are published under a Creative Commons Attribution-ShareAlike 4.0 licence (CC BY-SA 4.0). Third-party material is openly licensed or used with documented permission.

Reviewed by the consortium

Each resource was reviewed by project partners for quality, use case fit and alignment with the AI Engagement Framework before publication.

External review under way

Senior academics are reviewing the toolkit. The resources and this guide will be updated in 2027 in response.

Human judgement first

Every resource states a human-only boundary: the part of the task that stays with the researcher, whatever the tool.



ABOUT US

AI Researchers: Embedding Responsible AI in Research

An Erasmus+ Cooperation Partnership in Higher Education (KA220-HED), running from October 2025 to March 2028. It helps researchers and supervisors move from hesitation to confident, responsible use of AI.

University of Lodz
Poland, coordinator

FEP, University of Porto
Portugal

UPF Barcelona School of Management
Spain

ACEEU
Germany

NIPI (OPI PIB)
Poland

European E-learning Institute
Denmark

Momentum
Ireland

airesearchers.eu ›

AI Researchers

Embedding Responsible AI in Research

Use, share and adapt

This guide and the resources it links to are published under CC BY-SA 4.0. You may share and adapt them for any purpose, provided you credit the source and share your adaptations under the same licence.

airesearchers.eu

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author or authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them.

CC BY SA