


CURRICULUM FOR THE MICRO-CREDENTIAL

Ethical and Responsible Use of AI for University Students

AI in Academia, Innovation and Business



EmpowerAI



Funded by
the European Union



P.PORTO

2024-1-BG01-KA220-HED-000252605



UNIVERSITY OF
PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΑΣ



“europeanprofiles”^{SA}



SUCCESSOS
criativos

TABLE OF CONTENTS

<i>I. INTRODUCTION.....</i>	<i>4</i>
<i>II. GENERAL CHARACTERISTICS OF THE MICRO-CREDENTIAL.....</i>	<i>5</i>
<i>III. RATIONALE AND PEDAGOGICAL APPROACH.....</i>	<i>6</i>
<i>IV. AIM AND SPECIFIC OBJECTIVES.....</i>	<i>7</i>
<i>IV.1. Aim of the programme.....</i>	<i>7</i>
<i>IV.2. Specific objectives</i>	<i>7</i>
<i>V. EXPECTED LEARNING OUTCOMES</i>	<i>8</i>
<i>V.1. Knowledge</i>	<i>8</i>
<i>V.2. Skills.....</i>	<i>9</i>
<i>V.3. Competences</i>	<i>9</i>
<i>VI. STRUCTURE AND LEARNING WORKLOAD</i>	<i>10</i>
<i>VI.1. Guided Instruction</i>	<i>12</i>
<i>VI.2. EmpowerAI Chatbot – Guided Practical Activity.....</i>	<i>12</i>
<i>VI.3. Independent Case-Based Assignment.....</i>	<i>13</i>
<i>VI.4. Final Assessment and Reflection</i>	<i>13</i>
<i>VI.5. Integration of Learning Components</i>	<i>14</i>
<i>VII. MODULE DESCRIPTIONS.....</i>	<i>14</i>
<i>VII.1. Module 1 – Foundations of AI and Ethics</i>	<i>14</i>
<i>VII.2. Module 2 – AI in Academia</i>	<i>17</i>
<i>VII.3. Module 3 – AI in Business across Industries</i>	<i>20</i>
<i>VIII. PRACTICAL LEARNING COMPONENT</i>	<i>23</i>
<i>VIII.1. EmpowerAI Chatbot – Guided Practical Activity.....</i>	<i>23</i>
<i>VIII.2. Independent Case-Based Assignment.....</i>	<i>24</i>
<i>VIII.3. Role of the Learner and the AI Tool</i>	<i>26</i>

IX.	<i>ASSESSMENT AND COMPLETION REQUIREMENTS</i>	26
IX.1.	Formative Assessment	26
IX.2.	Summative Assessment	27
	A. Independent Case-Based Assignment	27
	B. Knowledge Assessment	27
IX.3.	Assessment of the Case-Based Assignment	28
IX.4.	Final Reflection	29
IX.5.	Requirements for Successful Completion	29
IX.6.	Pilot Evaluation and Participant Feedback	29
X.	<i>REFERENCES AND KEY RESOURCES</i>	29
	<i>APPENDIX 1 – CASE STUDIES</i>	31
	<i>APPENDIX 2 CASE STUDY RESPONSE TEMPLATE</i>	66

I. INTRODUCTION

Artificial Intelligence (AI), and particularly generative AI, is rapidly transforming the ways in which university students learn, conduct research, develop ideas, solve problems, and prepare for their future professional lives. While AI offers significant opportunities for learning, innovation, and productivity, its use also raises important ethical, academic, and societal challenges related to academic integrity, reliability of information, transparency, bias, privacy, human responsibility, and overreliance on automated systems.

This curriculum has been developed within the Erasmus+ project “EmpowerAI – Ethical AI in Academia” (Project No. 2024-1-BG01-KA220-HED-000252605), under Work Package 5 – Micro-Credentials for University Students, Task 5.1. It provides the framework for a transdisciplinary micro-credential designed for university students from different fields of study and educational levels.

The curriculum builds directly on the results produced in the preceding work packages of the EmpowerAI project. The Guidelines for Ethical Use of AI in Academia, developed under WP2, provide the main ethical principles, risks, use cases, and recommendations underpinning the learning content. The curriculum also draws on the pedagogical experience gained through the WP3 Seminar Series for Academic Staff, particularly the use of case-based learning, ethical reflection, and the application of the Guidelines to authentic academic situations. The practical learning component is further supported by the EmpowerAI chatbot developed under WP4, which enables students to explore ethical questions, consult guidance based on the project resources, and apply the acquired knowledge when solving realistic cases.

The micro-credential therefore represents a further step in the progressive implementation of the EmpowerAI results: from the development of ethical guidance, through capacity building for academic staff and the development of practical AI-supported tools, to strengthening students’ competences for the responsible use of AI.

The programme has a total learner workload of 10 academic hours and combines guided instruction, practical interaction with the EmpowerAI chatbot, independent case-based work, assessment, and reflection. Its purpose is not to provide advanced technical training in AI systems, but to develop AI literacy, critical thinking, ethical judgement, academic integrity, and responsible decision-making in academic, innovation, and professional contexts.

Particular attention is given to the principle that AI should support rather than replace human thinking, learning, and responsibility. Students are therefore encouraged not simply to learn how to use AI tools, but to understand when, why, and under what conditions their use is appropriate, ethical, transparent, and reliable.

II. GENERAL CHARACTERISTICS OF THE MICRO-CREDENTIAL

Characteristic	Description
Title	Ethical and Responsible Use of AI for University Students – AI in Academia, Innovation and Business
Type	Transdisciplinary micro-credential developed within the EmpowerAI project.
Target group	Bachelor's, Master's and doctoral students from all fields of study, including both STEM and non-STEM disciplines.
Entry requirements	No prior specialised knowledge of AI, programming or data science is required.
Mode of delivery	Face-to-face, online or hybrid, depending on the arrangements of the participating higher education institution.
Total learner workload	10 academic hours, combining guided instruction, practical interaction with the EmpowerAI chatbot, independent case-based work, assessment and reflection.
Main thematic areas	Foundations of AI and Ethics; AI in Academia; AI in Business across Industries.
Learning approach	The programme follows a practice-oriented, interdisciplinary and case-based approach. Theoretical input is combined with ethical reflection, critical evaluation of AI-generated outputs, interaction with the EmpowerAI chatbot and independent work on realistic academic and professional scenarios.
Completion requirements	To successfully complete the micro-credential, learners are required to participate in the guided learning activities, complete

Characteristic	Description
	the assigned case-based task, submit their solution, and complete the final knowledge assessment and reflection activities.
Award	Upon successful completion of the learning activities and assessment, learners receive a Certificate of Completion and an EmpowerAI digital badge certifying completion of the 10-academic-hour transdisciplinary micro-credential.

III. RATIONALE AND PEDAGOGICAL APPROACH

Artificial Intelligence is becoming an increasingly common part of students' academic and professional activities. University students use AI tools for searching and analysing information, generating ideas, preparing written assignments, coding, language support, data analysis, and problem-solving. These opportunities can improve learning and productivity, but they also create risks related to academic integrity, misinformation, bias, privacy, overreliance on AI, and the loss of independent judgement.

For this reason, higher education should not approach AI only as a technological tool, but also as an ethical, pedagogical, and professional challenge. Students need to understand not only what AI can do, but also when its use is appropriate, how its outputs should be critically evaluated, what responsibilities remain with the human user, and how AI should be used without undermining authentic learning and academic integrity.

The pedagogical approach of this micro-credential is therefore based on the principle of learning through application and reflection. The programme combines concise theoretical input with practical activities, guided interaction with the EmpowerAI chatbot, independent case-based work, and reflection on realistic academic and professional situations.

The curriculum adopts a case-based and learner-centred approach, enabling students from different disciplines to analyse ethical dilemmas, identify risks, compare possible courses of action, and justify responsible solutions. The case method allows students to apply general ethical principles to concrete situations rather than treating ethics as an abstract set of rules.

The design also builds on the pedagogical experience developed through the EmpowerAI WP3 Seminar Series for Academic Staff, where case-based learning and reflective practice were used to translate the Ethical Guidelines into authentic academic situations. In

WP5, this approach is adapted to the student perspective and complemented by the practical use of the EmpowerAI chatbot developed under WP4.

The programme is guided by the following pedagogical principles:

- practice-oriented learning – students apply concepts to realistic academic and professional cases;
- critical reflection – learners are encouraged to question AI-generated outputs and assess their reliability and ethical implications;
- human responsibility and oversight – AI is treated as a support tool, while responsibility for decisions and outputs remains with the human user;
- interdisciplinarity – the programme is designed for students from both STEM and non-STEM fields;
- active and independent learning – students combine guided instruction with chatbot-supported and self-directed work;
- ethical and responsible AI use – the learning process is grounded in the principles of transparency, fairness, privacy, accountability, academic integrity, and human agency.

The final goal of the pedagogical model is to help students move from passive use of AI tools toward informed, critical, and responsible engagement with AI in academic, innovation, and professional contexts.

IV. AIM AND SPECIFIC OBJECTIVES

IV.1. Aim of the programme

The aim of the micro-credential is to develop students' knowledge, skills, and competences for the ethical, critical, transparent, and responsible use of Artificial Intelligence in academic, innovation, business, and professional contexts.

The programme seeks to enable students not only to use AI tools effectively, but also to critically evaluate their outputs, recognise ethical and academic risks, make informed decisions about appropriate AI use, and retain human responsibility for the final outcome.

IV.2. Specific objectives

Upon successful completion of the micro-credential, learners should be able to:

- understand the basic principles, capabilities, and limitations of AI, generative AI, and Large Language Models;

- recognise key opportunities and risks associated with the use of AI in higher education and professional contexts;
- apply the core principles of ethical AI use, including fairness, transparency, accountability, privacy, inclusivity, and human oversight;
- distinguish between acceptable, questionable, and inappropriate uses of AI in academic work;
- apply principles of academic integrity and transparency when using AI-supported tools;
- critically evaluate AI-generated information and identify risks such as hallucinations, bias, misleading content, and fabricated sources;
- recognise privacy, data protection, copyright, and confidentiality risks related to AI use;
- understand the basic relevance of the EU AI Act and related ethical requirements for the use of AI in education and other professional settings;
- use AI as a support for learning and problem-solving without replacing independent reasoning, authorship, and personal responsibility;
- analyse realistic academic and professional cases involving AI and propose ethically justified and responsible solutions;
- use the EmpowerAI chatbot and project Guidelines as supporting resources when addressing ethical questions and practical cases;
- reflect critically on their own use of AI and make informed decisions about when, why, and how AI should be used.

V. EXPECTED LEARNING OUTCOMES

Upon successful completion of the micro-credential, learners will have developed the following knowledge, skills, and competences:

V.1. Knowledge

Learners will be able to:

- explain the basic concepts, capabilities, and limitations of Artificial Intelligence, generative AI, and Large Language Models;
- describe key opportunities and risks associated with the use of AI in academic, innovation, business, and professional contexts;

- identify the main principles of ethical and responsible AI use, including transparency, fairness, accountability, privacy, inclusivity, and human oversight;
- explain the relevance of academic integrity, authorship, disclosure, data protection, confidentiality, and copyright when using AI-supported tools;
- recognise common limitations and risks of AI-generated content, including hallucinations, bias, fabricated information, and misleading outputs;
- describe the basic relevance of the EU AI Act and related ethical requirements to the use of AI in education and professional settings.

V.2. Skills

Learners will be able to:

- critically evaluate AI-generated outputs for reliability, accuracy, bias, and ethical implications;
- distinguish between acceptable, questionable, and inappropriate uses of AI in academic work and other practical contexts;
- identify ethical, academic, privacy, and accountability risks in realistic AI-use scenarios;
- apply the EmpowerAI Ethical Guidelines when analysing practical cases and making decisions about AI use;
- use the EmpowerAI chatbot as a supporting resource for exploring ethical questions and solving case-based assignments;
- formulate and justify appropriate mitigation measures and responsible courses of action in cases involving AI;
- communicate and justify decisions concerning the responsible use of AI.

V.3. Competences

Learners will be able to:

- use AI in a transparent, critical, and responsible manner while maintaining personal responsibility for the final outcome;

- exercise human judgement and oversight when using AI-supported tools;
- preserve academic integrity and independent reasoning when AI is used to support learning, research, or problem-solving;
- apply ethical reasoning to AI-related situations across different academic disciplines and professional or business contexts;
- make informed decisions about when, why, and under what conditions AI should or should not be used;
- critically reflect on their own use of AI and adapt their practices in response to ethical, academic, and professional considerations.

VI. STRUCTURE AND LEARNING WORKLOAD

The micro-credential is designed as a 10-academic-hour learning pathway combining guided instruction, practical interaction with the EmpowerAI chatbot, independent case-based work, final assessment, and reflection.

The programme is structured around the three thematic areas defined in the EmpowerAI project: Foundations of AI and Ethics, AI in Academia, and AI in Business across Industries. These three areas constitute the core content modules of the micro-credential. They are followed by an applied learning component in which students use the EmpowerAI chatbot and the project Guidelines to work on a realistic case, and by a final assessment and reflection activity.

The distribution of the learner workload is presented below.

Learning component	Format	Workload (academic hours)	Main purpose
Module 1 – Foundations of AI and Ethics	Guided instruction	1	Introduce basic AI concepts, generative AI and LLMs, their capabilities and limitations, and the fundamental principles of ethical and responsible AI use.
Module 2 – AI in Academia	Guided instruction	1	Examine opportunities, risks and ethical challenges of AI use in learning,

Learning component	Format	Workload (academic hours)	Main purpose
			assessment and academic work, with particular attention to academic integrity, transparency, reliability, privacy and human responsibility.
Module 3 – AI in Business across Industries	Guided instruction	1	Explore the use of AI in innovation, business and professional contexts and examine the need to balance efficiency and innovation with ethical responsibility, human oversight and risk awareness.
EmpowerAI Chatbot – Guided Practical Activity	Guided/self-directed practical activity	1	Familiarise learners with the EmpowerAI chatbot and enable them to use it as a supporting resource for exploring ethical questions and preparing for the case-based assignment.
Independent Case-Based Assignment	Self-directed learning	5	Apply the acquired knowledge to a realistic academic or professional case by identifying ethical issues and risks, consulting relevant project resources, evaluating possible actions and proposing a justified responsible solution.
Final Assessment and Reflection	Individual assessment	1	Demonstrate understanding of the core concepts and principles through a knowledge assessment and reflect critically on the learner's own use of AI.

Learning component	Format	Workload (academic hours)	Main purpose
Total learner workload	10 academic hours		

VI.1. Guided Instruction

The first three academic hours provide the common knowledge base required for all participants, regardless of their field of study or previous experience with AI.

The guided component is intentionally concise and focuses on essential concepts rather than advanced technical knowledge. Its purpose is to establish a common understanding of AI, its opportunities and limitations, and the ethical principles that should guide its use in academic and professional contexts.

The three modules are interconnected. Module 1 establishes the conceptual and ethical foundation; Module 2 applies this foundation to university learning and academic work; and Module 3 extends the same principles to innovation, business and professional environments. Guided instruction may include short presentations, examples, discussion questions, mini-cases and interaction with learners. Depending on the participating institution, it may be delivered face-to-face, online or in hybrid format.

VI.2. EmpowerAI Chatbot – Guided Practical Activity

Following the guided instruction, learners complete a practical activity using the EmpowerAI chatbot.

The purpose of this activity is not simply to demonstrate the use of a chatbot, but to introduce students to a structured way of using AI-supported guidance when addressing ethical questions. Learners explore selected questions related to responsible AI use, compare the chatbot's guidance with the principles introduced during the training, and become familiar with the project resources on which the chatbot is based.

This activity also prepares learners for the independent case-based assignment. Students are encouraged to use the chatbot as a supporting resource for reasoning and reflection, rather than as a tool that provides a ready-made final answer.

VI.3. Independent Case-Based Assignment

The independent case-based assignment constitutes the main applied component of the micro-credential and accounts for five academic hours of the total workload.

Each learner is provided with a realistic case involving the use of AI in an academic, innovation, business or professional context. The cases are designed to present situations in which the use of AI creates both opportunities and potential ethical or practical risks.

Using the knowledge acquired during the guided learning component, the EmpowerAI Guidelines for Ethical Use of AI in Academia, and the EmpowerAI chatbot, learners are expected to analyse the assigned case and develop a reasoned solution.

The case-based work should require learners to:

1. identify the main problem and relevant context;
2. identify the role of AI in the situation;
3. recognise the main ethical, academic, legal or professional risks;
4. relate the identified issues to relevant principles of responsible AI use;
5. consider possible alternative actions;
6. propose appropriate mitigation measures;
7. recommend and justify a responsible course of action;
8. indicate where human judgement, oversight and responsibility are required;
9. reflect on how the case relates to their own academic or future professional practice.

Learners submit their solution using a common case-study template provided within the programme. This ensures comparability of outputs across participating institutions while still allowing cases to reflect different disciplinary contexts.

VI.4. Final Assessment and Reflection

The final session is dedicated to assessment and reflection (regarding the case study work) concerning the core concepts covered in the program, including the responsible use of AI, hallucinations and the reliability of AI-generated information, bias and transparency, academic integrity, privacy and data protection, as well as human oversight and accountability.

The case study demonstrates the participants' ability to apply ethical judgment in a specific situation, while the knowledge check confirms their understanding of the key concepts and principles underpinning the micro-credential.

Learners also complete a short reflective activity focused on their own use of AI. The reflection encourages them to consider how their approach to AI may change as a result of the programme and what principles they intend to apply in future academic and professional situations.

VI.5. Integration of Learning Components

The structure of the micro-credential moves from introductory knowledge and discussion to practical application and individual reflection. Students first become familiar with the main concepts and ethical principles, then apply them through the EmpowerAI chatbot and the case-based assignment, and finally demonstrate their understanding through assessment and reflection.

This structure is intended to ensure that learners do not only receive information about ethical AI use, but actively apply it to realistic situations and demonstrate their ability to make responsible and justified decisions.

The relative emphasis on independent and applied learning reflects the overall philosophy of the programme: responsible AI use cannot be developed through theoretical instruction alone. Students need opportunities to analyse uncertainty, recognise competing considerations, consult appropriate resources and take responsibility for their final decisions.

VII. MODULE DESCRIPTIONS

VII.1. Module 1 – Foundations of AI and Ethics

Duration: 1 academic hour

Format: Guided instruction; face-to-face, online or hybrid

Module overview

This module provides the common conceptual and ethical foundation for the micro-credential. It introduces learners to the basic characteristics of Artificial Intelligence, generative AI and Large Language Models and explains why apparently convincing AI-generated outputs may still be incomplete, inaccurate or misleading.

The module also introduces the main ethical principles that underpin responsible AI use.

Particular attention is given to the fact that AI systems do not assume moral or professional

responsibility for their outputs. Human judgement, verification and responsibility therefore remain essential whenever AI is used to support learning, decision-making or professional activities.

Learning objectives

By the end of the module, learners should be able to:

- explain in general terms what Artificial Intelligence, generative AI and Large Language Models are;
- recognise the principal capabilities and limitations of generative AI;
- understand why AI-generated information may contain hallucinations, errors or fabricated information;
- identify the main ethical principles relevant to responsible AI use;
- explain why human oversight and responsibility remain necessary when AI tools are used.

Indicative content

1. Basic concepts

Introduction to Artificial Intelligence and the distinction between AI, generative AI and Large Language Models. The discussion focuses on concepts relevant to users rather than technical architecture.

Learners are introduced to the types of outputs AI systems may produce, including text, code, images, recommendations and other generated content.

2. Capabilities and limitations of generative AI

The module discusses what generative AI can support effectively, while emphasising that fluent or convincing output should not automatically be interpreted as reliable knowledge.

Key issues include:

- generation of plausible rather than necessarily factual responses;
- hallucinations and fabricated information;
- incomplete or outdated information;
- dependence on the quality and context of the input;
- possible reproduction of bias;
- the need for independent verification.

3. Core principles of ethical and responsible AI

Learners are introduced to the ethical principles used throughout the EmpowerAI project:

- fairness;

- transparency and explainability;
- inclusivity;
- privacy and data protection;
- human agency and oversight;
- accountability.

These principles are presented as practical considerations that should influence how individuals select, use and evaluate AI tools rather than as abstract ethical concepts.

4. Human responsibility

The module introduces a central principle of the programme: AI may support a task or decision, but responsibility for the final outcome remains with the human user.

Learners consider why AI-generated outputs should be checked and why an error cannot simply be attributed to the tool when the output has been accepted and used by a person.

5. Basic regulatory context

Learners receive a short introduction to the risk-based approach of the European AI regulatory framework and its relevance to the responsible use of AI. The purpose is to provide awareness of the broader regulatory context rather than detailed legal training.

Learning activities

The trainer may use short questions, polls or examples to stimulate discussion. Learners may be asked to compare a reliable-looking AI response with an output containing an error or fabricated information and identify why independent verification is needed.

A short discussion may also address the question:

If an AI system produces an incorrect or harmful result, who is responsible for deciding whether that result should be used?

Connection to assessment

The concepts introduced in Module 1 are assessed through questions in the final knowledge assessment and are expected to be applied when learners identify risks and responsibilities in the case-based assignment.

Suggested learning resources

- *EmpowerAI Guidelines for Ethical Use of AI in Academia*, particularly the sections on foundations of ethical AI use and core ethical principles;
- EmpowerAI student training materials;
- relevant European resources on ethical and responsible AI and the EU AI Act.

VII.2. Module 2 – AI in Academia

Duration: 1 academic hour

Format: Guided instruction; face-to-face, online or hybrid

Module overview

This module examines the use of AI from the perspective of university students. It addresses both the opportunities AI creates for learning and academic work and the risks that arise when AI replaces rather than supports the learning process.

The module focuses particularly on academic integrity, transparency, reliability of information, privacy, fairness and the distinction between legitimate assistance and inappropriate substitution of student work.

Learning objectives

By the end of the module, learners should be able to:

- recognise appropriate ways in which AI can support university learning;
- identify risks associated with excessive or inappropriate reliance on AI;
- distinguish between acceptable, questionable and inappropriate uses of AI in academic work;
- understand the importance of transparency and disclosure when AI has contributed to academic work;
- recognise the learner's responsibility for verifying AI-generated information;
- apply basic ethical questions before deciding whether and how to use AI for an academic task.

Indicative content

1. Opportunities for AI-supported learning

Examples of legitimate learning support may include:

- explanation of difficult concepts;
- generation and organisation of ideas;
- support for text structuring;
- translation and language support;
- exam preparation;
- generation of examples for practice;
- programming and debugging support;
- support for accessibility and different learning needs.

The emphasis is placed on AI as a learning assistant rather than a replacement for the intellectual work expected from the student.

2. Risks in academic use

Learners discuss risks including:

- loss of critical thinking and independent problem-solving;
- inaccurate or misleading information;
- fabricated sources and quotations;
- bias and unfairness;
- academic dishonesty;
- privacy and disclosure of personal or confidential information;
- copyright and authorship concerns;
- unequal access to AI tools;
- overdependence on automated support.

These risks are consistent with those identified in the EmpowerAI Ethical Guidelines.

3. Academic integrity and AI

The module addresses the changing boundaries of academic integrity in the context of generative AI.

Learners consider the difference between using AI to:

- support understanding;
- brainstorm ideas;
- improve language;
- suggest examples or code that the learner understands;

and using AI to:

- produce substantial work that is submitted as the learner's own;
- generate complete assignments without meaningful student contribution;
- fabricate or introduce unchecked references;
- provide answers during assessments where AI use is not permitted.

The module emphasises that the relevant question is not only whether AI has been used, but also what intellectual contribution remains with the student and whether the use is consistent with the requirements of the course or assessment.

4. Transparency and disclosure

Where required by institutional, course or assessment rules, learners should be able to explain:

- which AI tool was used;
- why it was used;
- how it contributed to the work;
- what parts of the result were independently checked;
- what responsibility the learner assumes for the submitted work.

5. A practical ethical self-check

Before using AI for academic work, learners are encouraged to consider questions such as:

- Is the use fair to other students?
- Is the use of AI sufficiently transparent?
- Have I independently checked the result?
- Can I personally explain and defend the final work?
- Could the use disadvantage or exclude others?
- Are personal, academic or confidential data protected?

This self-check is drawn from the practical approach already developed in the EmpowerAI student training materials.

Learning activities

The module may include a short scenario-based discussion in which learners classify different forms of AI use as acceptable, questionable or inappropriate, taking into account the purpose of the assignment and the applicable rules.

A short case may also be used, for example a student who submits AI-generated code or written work that appears correct but cannot adequately explain or defend it.

The purpose of the discussion is not to establish one universal rule for every discipline, but to train students to identify the factors that should be considered before AI is used.

Connection to assessment

The final knowledge assessment includes questions related to academic integrity, transparency, hallucinations, privacy and responsible AI use. The same principles are expected to be applied in the independent case-based assignment.

Suggested learning resources

- *EmpowerAI Guidelines for Ethical Use of AI in Academia*, especially the sections on ethical use cases, academic integrity and the ethical checklist;
- EmpowerAI student training materials;
- applicable institutional or course-level rules on AI use.

VII.3. Module 3 – AI in Business across Industries

Duration: 1 academic hour

Format: Guided instruction; face-to-face, online or hybrid; may include contribution from an invited business expert

Module overview

This module extends the discussion beyond the university context and introduces students to the role of AI in contemporary organisations, innovation and professional practice.

The purpose is not to provide detailed technical training in enterprise AI systems or to present a catalogue of commercial tools. Instead, learners examine how AI can support common organisational activities, what value it may create, and what risks need to be considered before AI-supported outputs are used in professional decisions.

The module also helps students understand that the ethical principles discussed in the previous modules remain relevant when AI is used in business and professional environments.

Learning objectives

By the end of the module, learners should be able to:

- recognise common ways in which AI is used in contemporary organisations;
- explain how AI can support rather than automatically replace human professional judgement;
- identify potential benefits and risks of AI-supported business processes and decisions;
- recognise risks related to bias, confidentiality, privacy, intellectual property and unreliable AI-generated information;
- understand the importance of human review, transparency and accountability in professional AI use;
- apply responsible AI principles to a simple business or organisational scenario.

Indicative content

1. AI in contemporary organisations

Learners are introduced to common areas in which AI may support organisational activities, including:

- content creation and communication;
- information search and analysis;
- customer support;
- marketing and communications;
- human resources;
- project and operations management;
- data analysis and reporting;
- decision support.

The examples are used to illustrate functions rather than to teach particular commercial platforms.

2. AI, innovation and human–AI collaboration

The module considers AI as a possible support for innovation, entrepreneurship and productivity.

Students discuss how AI may assist with idea development, analysis, communication, automation and decision support, while recognising that a technological solution should respond to a genuine problem or need rather than being adopted simply because AI is available.

3. Responsible use of AI in professional contexts

The same ethical principles introduced earlier are applied to business and organisational settings.

Particular attention is given to:

- verifying AI-generated information before professional use;
- protecting confidential organisational and personal data;
- identifying possible bias and discrimination;
- ensuring human review of decisions with important consequences;
- transparency concerning AI-supported work;
- responsibility for final decisions;
- copyright and intellectual property considerations;
- avoiding excessive dependence on AI-generated recommendations.

These themes are already present in the working material developed for the business component of WP5.

4. AI-supported decisions

Students consider the difference between AI supporting a professional decision and AI effectively making the decision without sufficient human scrutiny.

Examples may include:

- AI-supported recruitment;
- market analysis;
- customer recommendations;
- automated communication;
- financial or operational recommendations.

The discussion emphasises that greater potential impact requires greater attention to verification, oversight and accountability.

5. Business experts and real-world perspective

In line with the design of WP5, pilot activities are enriched by contributions from business experts in the field of AI. Their role is to provide students with practical perspectives on current opportunities, challenges and responsible uses of AI in professional environments. The approved project foresees nine invited business experts across the participating countries. The contribution may take the form of a short guest input, practical example, discussion or business case, depending on the organisation of the local pilot.

Learning activities

Learners may work briefly with a business scenario. For example:

A manager receives an AI-generated market analysis before making an important business decision.

Students consider:

- Which parts of the analysis should be verified?
- What additional information would be needed?
- What risks could arise from accepting the recommendation without further review?
- Who remains responsible for the final decision?

Alternative scenarios may address recruitment, customer service, project management or another professional field relevant to the participating students.

Connection to assessment

Business and professional contexts may be included in both the final knowledge assessment and the pool of cases used for the independent assignment. Learners are expected to transfer the same principles of transparency, verification, fairness, privacy, human oversight and accountability from academic situations to professional contexts.

Suggested learning resources

- EmpowerAI Guidelines and project learning materials relevant to responsible AI use;
- WP5 business-oriented training materials;
- selected EU resources on responsible and trustworthy AI;
- examples and practical input provided by participating business experts.

VIII. PRACTICAL LEARNING COMPONENT

The practical learning component enables students to apply the concepts and ethical principles introduced in the three content modules to realistic situations involving the use of Artificial Intelligence. It consists of two connected activities: a guided practical activity with the EmpowerAI chatbot and an independent case-based assignment, followed by reflection.

Together, these activities account for six academic hours of the total learner workload. They are designed to strengthen students' ability to identify ethical risks, consult appropriate resources, evaluate alternative courses of action, and justify responsible decisions concerning the use of AI.

VIII.1. EmpowerAI Chatbot – Guided Practical Activity

Duration: 1 academic hour

Format: Guided or self-directed practical activity

The EmpowerAI chatbot is used as a learning support tool that helps students explore questions related to ethical and responsible AI use. The purpose of the activity is not to obtain ready-made answers, but to familiarise learners with a structured way of asking questions, examining ethical concerns and consulting the knowledge resources developed within the EmpowerAI project.

During the activity, students interact with the chatbot on selected topics connected to the content of the micro-credential. These may include academic integrity, transparency of AI use, hallucinations and reliability, bias and fairness, privacy and data protection, human oversight and accountability.

Students are encouraged to compare the chatbot's responses with the principles introduced during the guided instruction and with the EmpowerAI Guidelines for Ethical Use of AI in Academia. Where appropriate, they should question the response, ask for clarification, explore alternative interpretations and consider whether additional verification is necessary.

The activity also prepares students for the independent case-based assignment. It demonstrates how the chatbot can be used as a source of guidance and reflection while leaving the analysis, judgement and final decision to the learner.

The trainer may provide several example prompts or ethical questions to help students begin the activity. Students should also formulate at least one question of their own related to an academic or professional situation in which the use of AI may create an ethical dilemma.

At the end of the activity, learners should be able to use the EmpowerAI chatbot purposefully while maintaining a critical attitude towards AI-supported guidance.

VIII.2. Independent Case-Based Assignment

Duration: 5 academic hours

Format: Independent learning

The independent case-based assignment is the main applied learning activity of the micro-credential. It gives learners the opportunity to demonstrate that they can transfer the knowledge and principles introduced during the programme to a realistic situation involving the use of AI.

Each learner receives or selects a case from a common pool of scenarios. Cases may address situations arising in academic study, research, innovation, business or professional practice. They should present a genuine dilemma or problem in which the use of AI offers potential benefits but also creates ethical, academic, legal or professional risks.

Students analyse the case using the knowledge acquired during the guided learning component, the EmpowerAI Guidelines for Ethical Use of AI in Academia, and the EmpowerAI chatbot as supporting resources.

The analysis should address the following elements:

1. **Description of the problem and context** – identify the main issue and explain how AI is being used in the situation.
2. **Identification of risks and ethical concerns** – identify relevant issues such as reliability, bias, transparency, privacy, academic integrity, confidentiality, intellectual property or excessive automation.
3. **Relevant ethical principles** – relate the identified concerns to the principles of responsible AI use discussed during the programme.
4. **Use of the EmpowerAI chatbot and other learning resources** – indicate how the project resources were used to support the analysis.
5. **Alternative courses of action** – consider possible responses to the situation and their likely consequences.
6. **Proposed mitigation measures** – explain how the identified risks could be reduced or managed.
7. **Recommended solution** – propose and justify an appropriate and responsible course of action.
8. **Human oversight and responsibility** – indicate where human judgement, verification and accountability are required.
9. **Case-specific reflection** – briefly explain what the learner has learned from the case, which aspects of the situation were most challenging, and how the analysis may influence their approach to similar AI-related situations in the future.

Learners submit their work using a **common case-study template** developed for the micro-credential. The use of a common template ensures that the same core elements are addressed by students across the participating institutions while allowing the cases themselves to reflect different academic disciplines and professional contexts.

The assignment is not intended to test technical expertise in AI. Assessment focuses on the learner's ability to recognise relevant issues, apply ethical principles, reason critically and justify a responsible solution.

VIII.3. Role of the Learner and the AI Tool

Throughout the practical component, the learner remains responsible for the analysis and final submitted work. Use of the EmpowerAI chatbot or another permitted AI tool does not replace the learner's own reasoning.

Students should therefore be able to explain how AI was used during the activity, evaluate the usefulness and limitations of the support received, and take responsibility for the accuracy and quality of the final submission.

IX. ASSESSMENT AND COMPLETION REQUIREMENTS

Assessment within the micro-credential is designed to determine whether learners have understood the main principles introduced during the programme and are able to apply them to realistic situations involving the use of Artificial Intelligence.

In accordance with the design of Work Package 5 of the EmpowerAI project, the assessment combines formative and summative elements. Formative assessment supports learning throughout the programme, while summative assessment verifies achievement of the expected learning outcomes at the end of the learning process.

The assessment focuses on ethical reasoning, critical evaluation and responsible decision-making rather than on advanced technical knowledge of AI systems.

IX.1. Formative Assessment

Formative assessment takes place throughout the guided and practical learning activities. Its purpose is to provide learners with opportunities to check their understanding, identify misconceptions and improve their reasoning before completing the final assessment.

Formative assessment may include:

- responses to questions and short discussions during the three guided modules;
- analysis of short scenarios and ethical dilemmas;
- identification of risks in examples of AI use;
- interaction with the EmpowerAI chatbot and critical consideration of the guidance provided;
- comparison of alternative courses of action;
- trainer feedback during guided activities.

Formative activities are not graded and do not contribute directly to the final result. Their purpose is to prepare learners for the independent case-based assignment and final knowledge assessment.

IX.2. Summative Assessment

Summative assessment consists of two components:

A. Independent Case-Based Assignment

The case-based assignment is the principal applied assessment activity. Learners analyse a realistic situation involving the use of AI and submit their response using the common case-study template described in Section VIII.

The assignment assesses the learner's ability to:

- identify the main problem and the role of AI in the situation;
- recognise relevant ethical, academic, legal or professional risks;
- apply appropriate principles of ethical and responsible AI use;
- distinguish between appropriate and inappropriate courses of action;
- propose realistic mitigation measures;
- provide and justify a responsible solution;
- identify where human judgement, verification and oversight are required;
- take responsibility for the final recommendation.

The case is evaluated using a common assessment rubric to ensure consistency across the participating institutions.

B. Knowledge Assessment

Learners also complete a short individual knowledge assessment addressing the main concepts covered by the micro-credential.

The assessment should include questions related to topics such as:

- basic characteristics and limitations of generative AI;
- hallucinations and reliability of AI-generated information;
- bias and fairness;
- transparency and disclosure of AI use;
- academic integrity;
- privacy, data protection and confidentiality;
- human oversight and accountability;
- appropriate and inappropriate uses of AI in academic and professional contexts.

The knowledge assessment may include multiple-choice, true/false and short scenario-based questions. It is completed individually.

A minimum score of 60% is required to pass the knowledge assessment.

IX.3. Assessment of the Case-Based Assignment

The case-based assignment is evaluated against the following criteria:

Criterion	What is assessed
Identification of the problem and risks	The learner correctly identifies the main issue, the role of AI and the relevant risks.
Application of ethical principles	Relevant principles such as transparency, fairness, privacy, accountability, academic integrity and human oversight are appropriately applied.
Critical reasoning	The learner considers the consequences of different possible actions and demonstrates independent judgement rather than simply reproducing AI-generated guidance.
Proposed solution and mitigation measures	The recommended response is realistic, responsible and supported by appropriate measures for reducing the identified risks.
Human oversight and responsibility	The learner clearly identifies where human verification, judgement and accountability are required.
Justification of the final decision	The proposed solution is clearly explained and supported by reasoned arguments.

The assignment is assessed as:

Satisfactory / Revision Required

A **Satisfactory** result requires the learner to demonstrate adequate achievement across all core assessment criteria. Where substantial elements are missing or the analysis does not demonstrate the expected learning outcomes, the assignment is returned for revision.

IX.4. Final Reflection

At the end of the programme, learners complete a short individual reflection on their overall learning experience and future use of AI.

The reflection focuses on how their understanding of responsible AI use has developed during the micro-credential, which ethical principles they consider most relevant to their own academic or professional context, and how they intend to apply these principles in future situations involving AI.

The final reflection is a completion requirement but is not graded.

IX.5. Requirements for Successful Completion

To successfully complete the micro-credential, a learner must:

1. participate in the guided learning component;
2. complete the practical activity with the EmpowerAI chatbot;
3. submit the independent case-based assignment and receive a Satisfactory result;
4. achieve at least 60% in the final knowledge assessment;
5. complete the final reflection.

Learners who successfully meet these requirements receive a Certificate of Completion and an EmpowerAI digital badge certifying completion of the 10-academic-hour transdisciplinary micro-credential.

IX.6. Pilot Evaluation and Participant Feedback

Evaluation of the learning programme itself is conducted separately from learner assessment.

During the pilot implementation, participants are invited to complete a feedback questionnaire concerning the clarity, relevance, usefulness and overall quality of the micro-credential, including the guided training, practical activities, EmpowerAI chatbot and case-based assignment. The responses do not influence the learner's assessment result or eligibility for the Certificate of Completion and digital badge.

Feedback may also be collected from participating trainers and business experts as part of the overall evaluation of the pilot implementation.

X. REFERENCES AND KEY RESOURCES

EmpowerAI. (2025). Guidelines for Ethical Use of AI in Academia.

<https://empower-ai.eu/download/469/?tmstv=1789305370>

European Commission. Artificial Intelligence Act – Regulatory Framework for AI.

<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).

<https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32024R1689>

European Commission. (2022). Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators.

<https://education.ec.europa.eu/focus-topics/digital-education/action-plan/action-6-ethical-guidelines-for-educators-on-the-use-of-ai-and-data-in-teaching-and-learning-for-educators>

UNESCO. (2023). Guidance for Generative AI in Education and Research.

<https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>

UNESCO. (2024). AI Competency Framework for Students.

<https://www.unesco.org/en/articles/ai-competency-framework-students>

European Commission. AI Literacy, Skills and Talent.

<https://digital-strategy.ec.europa.eu/en/policies/ai-talent-skills-and-literacy>

APPENDIX 1 – CASE STUDIES

Case Study: 1	
Title: AN AI GRADING TOOL APPEARS TO GIVE LOWER SCORES TO NON-NATIVE SPEAKERS	
Description:	Your course uses an AI grading platform for assignments. You appreciate the fast feedback, but after comparing results with classmates you notice a worrying pattern: students who are non-native speakers often receive lower scores even when their answers demonstrate strong understanding of the subject. Some comments from the system focus heavily on grammar and style rather than the quality of the ideas. You are directly affected by a grading process whose criteria you cannot fully see. You must decide how to respond responsibly: whether and how to challenge the grade, what evidence to collect, how to raise the issue without making unsupported accusations, and what safeguards you believe the university should provide when AI contributes to assessment.
Key points for consideration:	• Algorithmic bias and possible discrimination. • Fairness and equal treatment in assessment. • Transparency about the role and criteria of AI grading. • The right to meaningful human review. • Responsible ways for students to question an automated decision.
Task:	Resolve the case from the student perspective. 1. Identify the main ethical and academic problems in this case. 2. What would you do immediately after receiving the grade? Give concrete steps. 3. What evidence would you collect before challenging the result? 4. How would you communicate your concern to the instructor or department? 5. Should you request human re-grading? Explain why and under what conditions. 6. What responsibilities do the instructor and university have when AI is used in grading? 7. Propose safeguards that would make AI-supported grading fairer and more transparent.
Main principles:	fairness, transparency, accountability, and non-discrimination

Case Study: 2	
Title: YOUR INSTRUCTOR SUSPECTS THAT YOUR ESSAY WAS AI-GENERATED, BUT THERE IS NO CLEAR PROOF	
Description:	You submit an essay that is substantially more polished than your earlier work. You used an AI tool during preparation, but the extent and purpose of that use may not be obvious from the final submission. Your instructor tells

	you that the essay appears AI-generated and questions whether the work represents your own learning. There is no clear course rule defining which uses of generative AI are permitted, and AI-detection tools cannot provide certain proof. You now need to respond in a way that protects academic integrity while also ensuring that you are treated fairly. You must be able to explain how the work was produced, distinguish your own contribution from AI assistance, and decide what evidence or documentation you should provide.
Key points for consideration:	• Uncertainty and limitations of AI-detection methods. • Academic honesty and disclosure of AI use. • Fair treatment when evidence is inconclusive. • Student responsibility for submitted work. • Importance of clear course policies.
Task:	Resolve the case from the student perspective. 1. Explain what you would say to the instructor about how the essay was produced. 2. What evidence could demonstrate your authorship and understanding (drafts, notes, sources, version history, oral explanation, etc.)? 3. If you used AI, what should you disclose and how should you describe its contribution? 4. Would you agree to an oral defense or a supervised rewriting task? Explain. 5. How should the case be handled if the course had no explicit AI-use policy? 6. Where should the boundary lie between legitimate AI assistance and academic misconduct? 7. Propose a fair resolution for both the student and instructor. 8. Formulate a short AI-use declaration that you would attach to future assignments.
Main principles:	authorship, disclosure, verification, learning outcomes, responsibility

Case Study: 3	
Title: YOU RELY ON A UNIVERSITY AI CHATBOT FOR ACADEMIC ADVISING, BUT SOME OF ITS GUIDANCE MAY BE WRONG	
Description:	Your university provides an AI chatbot for course registration, degree requirements, deadlines, and career questions. You use it because it is available at any time and gives immediate answers. Based on its advice, you begin planning your courses. Later, classmates tell you that the chatbot has given outdated or inaccurate information about graduation requirements. You also realize that you have shared personal academic information in the conversation and do not know how those data are stored or used. You must decide how much trust to place in the chatbot, how to verify advice that could affect your studies, and when a human advisor must become

	involved. You also need to consider your privacy and the consequences of acting on incorrect automated guidance.
Key points for consideration:	• Accuracy and currency of AI-generated advice. • Student responsibility to verify high-impact information. • Privacy and handling of personal data. • Appropriate division of roles between AI and human advisors. • Accountability when automated advice causes harm.
Task:	Resolve the case from the student perspective. 1. Which parts of the chatbot's advice would you verify before acting on them? 2. How and with whom would you verify information affecting graduation, registration, or other important decisions? 3. What personal information would you avoid sharing with an advising chatbot? 4. What would you do if you had already acted on incorrect AI advice? 5. Who should be accountable when an official university chatbot provides harmful misinformation? 6. Identify questions that should always be referred to a human advisor. 7. Propose student-facing rules for safe and responsible use of AI advising systems. 8. Give your final resolution: how should you use the chatbot without becoming over-reliant on it?
Main principles:	reliability, transparency and explainability, privacy and data protection, accountability, human oversight

Case Study: 4	
Title: AN AI PLATFORM PERSONALIZES YOUR LEARNING PATH, BUT ITS RECOMMENDATIONS MAY LIMIT YOUR CHOICES	
Description:	In an introductory course, an AI platform analyzes your test scores, assignments, and survey responses and creates a personalized learning path. It recommends readings, videos, and exercises that it predicts will suit your needs. The recommendations are useful, but after several weeks you notice that the system repeatedly directs you toward the same types of material and rarely suggests topics outside your predicted strengths and weaknesses. You begin to wonder whether the platform is helping you learn or quietly narrowing what you are encouraged to explore. You are also unsure what personal and academic data the system collects, how long they are stored, and whether they could influence future academic decisions.
Key points for consideration:	• Benefits and limits of personalized learning. • Risk of profiling or “pigeonholing” students. • Student autonomy, curiosity, and freedom to explore. • Transparency of recommendations. • Privacy and responsible use of learner data.
Task:	Resolve the case from the student perspective.

	1. What benefits would you accept from personalized AI recommendations? 2. What signs would make you concerned that the system is limiting your learning choices? 3. Would you follow, ignore, or challenge a recommendation that conflicts with your interests? Why? 4. What information should the university disclose about how your learning profile is created? 5. What control should you have over the data used for personalization? 6. Should students be able to opt out or reset their profile? Defend your answer. 7. Propose a model in which personalization supports learning without restricting autonomy. 8. State your recommended resolution and the ethical principles supporting it.
Main principles:	autonomy, fairness, transparency and explainability, privacy and data protection, human agency and oversight

Case Study: 5	
Title: YOU USE GENERATIVE AI TO HELP WITH ESSAYS, CODE, OR ASSIGNMENTS: WHERE SHOULD YOU DRAW THE LINE?	
Description:	You regularly use generative AI when studying. It can brainstorm ideas, explain difficult concepts, suggest essay structures, debug code, and sometimes produce complete answers. The course rules are not always clear, and your classmates use AI in very different ways. You know that these tools will probably be part of your future professional life, but you also know that submitting work you do not understand could undermine your learning and violate academic integrity. You must decide how to use AI in a way that genuinely supports your education rather than replacing your own intellectual work. The difficult part is defining the boundary between assistance, collaboration, and unacceptable outsourcing of assessed work.
Key points for consideration:	• Changing concepts of authorship and plagiarism. • AI as a legitimate learning assistant. • Student ownership and responsibility for submitted work. • Need to demonstrate understanding and original contribution. • Importance of discipline- and assignment-specific rules.
Task:	Resolve the case from the student perspective. 1. Classify several possible uses of AI as acceptable, conditionally acceptable, or unacceptable in your discipline. 2. Explain where you would draw the line between assistance and outsourcing your work. 3. What parts of an assessed task must remain demonstrably your own? 4. How would you verify AI-generated explanations, code, calculations, or references before using them? 5. When and how should AI use be declared? 6. If the instructor has provided no explicit rule, what principles should guide your decision? 7. Propose a practical personal code of conduct for using AI in assignments.

	8. Give your final resolution and explain how it protects both learning and academic integrity.
Main principles:	academic integrity, authorship, transparency and disclosure, responsibility, human agency and oversight

Case Study: 6	
Title: AI-GENERATED CAPTIONS AND LECTURE SUMMARIES HELP YOU FOLLOW THE COURSE, BUT THEY SOMETIMES CONTAIN ERRORS	
Description:	Your course provides real-time AI captions and automatically generated lecture summaries. You find them extremely useful, especially when the lecture is fast, highly technical, or delivered in a language that is not your first language. Some students depend on the captions for accessibility. However, you notice that technical terms are sometimes transcribed incorrectly and that summaries occasionally omit important qualifications or oversimplify complex ideas. You must decide how to benefit from this accessibility tool without treating its output as an authoritative record of the lecture. The case also raises a broader question: students should have equal access to learning, but accessibility support must be reliable enough not to create new disadvantages.
Key points for consideration:	• Accessibility and inclusion. • Risk of technical errors and oversimplification. • Verification of AI-generated learning materials. • Equal access to accurate educational content. • Shared responsibility among students, instructors, and the institution.
Task:	Resolve the case from the student perspective. 1. How would you use AI captions and summaries without relying on them uncritically? 2. What would you do when a caption or summary conflicts with the lecture, textbook, or your notes? 3. How should students report recurring errors in technical terminology? 4. What additional support should be available to students who depend on these tools for accessibility? 5. Should AI-generated summaries be treated as official course material? Why or why not? 6. What responsibilities do students have for checking AI-generated study materials? 7. Propose safeguards that preserve accessibility while reducing the risk of misinformation. 8. Present your recommended resolution for responsible student use of these tools.
Main principles:	inclusion and accessibility, reliability and accuracy, fairness, transparency, human oversight

Case Study: 7	
Title: THE INVISIBLE 'ASSISTANT' IN THE LABORATORY: HOW MUCH OF YOUR PROJECT CAN AI DO?	
Description:	You must submit an individual laboratory project containing data analysis, code, graphs, interpretation, and conclusions. You use an AI assistant for ideas, parts of the code, wording, and possibly the creation of some figures. The final report looks polished, but some functions produce results you cannot fully explain and parts of the writing resemble material found online. During the project defense, you realize that you cannot justify several analytical choices with confidence. The course has no clearly formulated AI-use policy. You must decide how to handle the situation before or during assessment. The central issue is not simply whether AI was used, but whether the submitted project truthfully represents your own understanding and contribution.
Key points for consideration:	• Ownership of code, analysis, text, and figures. • Ability to explain and reproduce submitted work. • Verification of AI-generated code and outputs. • Disclosure and citation of AI assistance. • Fairness when course rules are unclear.
Task:	Resolve the case from the student perspective. 1. Identify which uses of the AI assistant are ethically problematic and which may be legitimate. 2. Before submission, what parts of the project should you personally verify, reproduce, or rewrite? 3. How would you document the AI tools used, including what they contributed? 4. What should you do if you cannot explain a piece of AI-generated code or analysis? 5. Should you voluntarily disclose the extent of AI assistance even without an explicit policy? Explain. 6. How should your individual contribution be demonstrated during assessment? 7. Draft an AI-use declaration appropriate for this laboratory project. 8. Propose a final resolution that preserves both academic integrity and the legitimate educational use of AI.
Main principles:	academic integrity, authorship and disclosure, verification, accountability, human oversight

Case Study: 8	
Title: THE EXPERIMENT THAT NEVER WAS: CAN AI-GENERATED DATA REPLACE REAL EMPIRICAL DATA?	

Description:	Your assignment requires you to plan a small experiment, collect real empirical data, and analyze the results. You struggle to recruit enough participants and worry that you will miss the deadline. You therefore use generative AI to create a clean synthetic dataset and complete the statistical analysis. The results look convincing, the graphs are well presented, and the discussion is your own. However, the report does not initially make clear that the data were simulated rather than collected from a real experiment. You must decide whether this is an acceptable adaptation to practical difficulties or a fundamental violation of the assignment. The key issue is whether simulated data can ever substitute for required empirical evidence and, if simulation is used, how it must be disclosed and interpreted.
Key points for consideration:	Difference between empirical and synthetic data. • Truthfulness about how evidence was produced. • Assignment learning outcomes: conducting a real experiment and handling imperfect data. • Appropriate scientific uses of simulation. • Risk of presenting artificial results as scientific reality.
Task:	Resolve the case from the student perspective. 1. Is generating synthetic data acceptable when the assignment explicitly requires real data? Explain. 2. What should you do when practical difficulties prevent completion of the planned experiment? 3. If synthetic data are used for practice or demonstration, how must they be labelled and documented? 4. What is scientifically misleading about presenting simulated data as empirical observations? 5. Should you inform the instructor before changing from real to simulated data? Why? 6. How could you redesign the project honestly if participant recruitment fails? 7. Propose rules distinguishing acceptable simulation from falsification or fabrication. 8. Give your final resolution and explain what you would submit.
Main principles:	research integrity, transparency, scientific validity, verification, accountability

Case Study: 9	
Title: THE BRIDGE DESIGNED BY THE INVISIBLE ENGINEER: AI-GENERATED CALCULATIONS MAY CREATE REAL SAFETY RISKS	
Description:	For a structural-design project, you must calculate loads and design a small pedestrian bridge according to specified materials, dimensions, safety factors, and engineering standards. To save time, you use generative AI to produce calculations, tables, and CAD-related suggestions. The project looks professional, but during review you discover that some safety factors are wrong and several cited “EN standards” do not exist. You had assumed the AI output was reliable and did not independently verify every result.

	You now face a problem that goes beyond academic honesty. In engineering and other safety-critical disciplines, unverified AI output can lead to decisions that endanger people. You must determine what you should do before presenting or submitting the project and what level of verification is required when AI contributes to professional calculations.
Key points for consideration:	Verification of AI-generated calculations and standards. Professional competence and responsibility. Academic integrity and disclosure. Safety-critical consequences of hallucinations and errors. Human accountability for AI-assisted engineering work.
Task:	Resolve the case from the student perspective. 1. What should you do immediately after discovering the incorrect calculations and fictitious standards? 2. Which parts of a safety-critical project should never be accepted from AI without independent verification? 3. How would you verify standards, formulas, safety factors, and calculations using authoritative sources? 4. Should AI-generated engineering content be allowed in student projects? Under what conditions? 5. What must you disclose about AI use in the project? 6. How would you demonstrate that you personally understand and take responsibility for the design? 7. Propose an AI-use protocol for safety-critical academic work. 8. Give your final resolution, prioritizing both academic integrity and public safety.
Main principles:	safety, reliability and accuracy, verification, professional responsibility, human oversight, accountability

Case Study: 10	
Title: AI SUGGESTS QUOTATIONS AND SOURCES FOR YOUR ESSAY, BUT SOME OF THEM DO NOT EXIST	
Description:	For a critical essay, you use a generative AI assistant to help structure your argument, improve wording, suggest interpretations, and identify possible quotations and literary sources. The final essay is fluent and persuasive. While checking some references, however, you discover that several quotations do not match the original texts and some cited sources may be entirely fabricated. The central ideas are yours, but parts of the supporting evidence came from the AI assistant. You must decide what to do before submission – or, if the problem is discovered afterwards, how to respond. A polished text cannot be academically valid if its evidence is false. The case requires you to distinguish legitimate assistance with expression from the responsibility to verify every quotation, citation, and scholarly claim.
Key points for consideration:	• Accuracy and authenticity of scholarly sources. • AI hallucinations and fabricated citations. • Student responsibility for all submitted references and quotations.

	• Difference between assistance with writing and outsourcing scholarly verification. • Transparency about AI use.
Task:	Resolve the case from the student perspective. 1. What should you do as soon as you discover that AI-suggested quotations or sources may be false? 2. How would you verify every quotation, citation, and bibliographic reference? 3. Is it acceptable to use AI to discover possible sources? Under what conditions? 4. Who is responsible for a fabricated citation that originated from an AI tool but appears in your submitted essay? 5. How should you disclose AI assistance with structure, wording, or source discovery? 6. If the essay has already been submitted, how should you address the problem with the instructor? 7. Develop a verification checklist for using generative AI in academic writing. 8. Give your final resolution and explain how it protects scholarly and academic integrity.
Main principles:	academic integrity, source verification, reliability and accuracy, transparency, responsibility

Case Study: 11	
Title: REVERSE ENGINEERING AN AI-GENERATED ESSAY: CAN YOU CRITICALLY ANALYSE CHATGPT RATHER THAN LET IT WRITE FOR YOU?	
Description:	In an advanced English Studies course on modern literature, you are given an analytical essay generated by ChatGPT about a literary work. Your task is not to accept the essay as a model answer or use it as a substitute for your own writing. Instead, you must reverse engineer it by identifying its thesis, arguments, themes, perspectives, textual evidence, rhetorical strategies, and the way its ideas are connected. As you analyse the essay, you notice that some interpretations appear persuasive while others are superficial, unsupported, or different from your own reading. You must decide which elements are convincing, which require verification or criticism, and how your own interpretation differs. You then use this critical analysis to construct an original essay outline or concept map supported by examples from the literary work. The challenge is to use generative AI as an object of critical analysis rather than allowing it to replace your own reading, reasoning, and authorship.
Key points for consideration:	Critical analysis of AI-generated academic writing. Identification of thesis, arguments, evidence, and rhetorical strategies. Verification of interpretations and textual evidence. Comparison between AI-generated and student-generated interpretations. Originality, authorship, and academic integrity.

	Responsible use of AI to support higher-order learning.
Task:	Resolve the case from the student perspective. 1. Deconstruct the AI-generated essay: identify its thesis, principal arguments, evidence, assumptions, and rhetorical strategies. 2. Which parts of the AI analysis are strong and which are weak, superficial, inaccurate, or insufficiently supported? 3. Compare the AI's interpretation with your own reading of the literary work. Where do you agree and disagree? 4. How would you verify quotations, textual evidence, and interpretive claims produced by the AI? 5. At what point would using the AI-generated essay stop being a learning aid and become inappropriate dependence on AI? 6. Develop an original essay outline or concept map that reflects your own interpretation and uses independently selected evidence. 7. Explain how you would acknowledge the role of AI while preserving your own authorship. 8. Give your final assessment of whether reverse engineering an AI-generated essay can strengthen analytical skills and academic integrity, and justify your position.
Main principles:	critical thinking, verification, academic integrity, authorship, human agency and responsibility

Case Study: 12	
Title: BUILDING A FIVE-YEAR BUSINESS STRATEGY WITH AI: WHEN SHOULD YOU TRUST AI INSIGHTS AND WHEN SHOULD YOU TRUST YOUR OWN ANALYSIS?	
Description:	In a business course, you work in a team that must develop a five-year strategy for a hypothetical startup. Your instructor provides the company's complete financial figures for Year 1. Your team must use these data as the starting point for projections for Years 2–5 while also conducting product research, identifying market opportunities, analysing competitors, determining market positioning, and developing a marketing strategy. Your team may use AI tools to analyse information, suggest strategic options, explore projections, and challenge or validate parts of the proposed strategy. The AI produces convincing recommendations and forecasts, but some team members begin to treat its suggestions as authoritative. Others argue that actual market research, financial reasoning, and human judgement should determine the final decisions. Your challenge is to develop and defend a coherent business strategy in which AI genuinely supports your work without replacing critical analysis. You must explain how Year 1 data informed your projections, what evidence supports your market assumptions, which AI recommendations you accepted or rejected, and why.
Key points for consideration:	• Critical use of AI-generated business insights. • Integration of financial data with product and market research. • Verification of AI recommendations against actual market evidence.

	- Competitive analysis and strategic differentiation. - Team responsibility for AI-assisted decisions. - Ability to justify and defend the final strategy independently of AI.
Task:	Resolve the case from the student perspective. - Analyse the Year 1 financial data and explain how it should constrain or inform the strategy for Years 2 – 5. - Identify what product, market, and competitor research your team must conduct independently of AI. - Which parts of the strategy could appropriately be supported by AI, and which require direct evidence and human judgement? - How would you verify AI-generated forecasts, market claims, competitor information, and strategic recommendations? - Describe a situation in which your team should reject an apparently convincing AI recommendation. - How should your team document the contribution of AI and distinguish it from your own analysis and decisions? - Prepare the main arguments you would use to defend your integrated financial, product, marketing, and competitive strategy. - Give your final resolution: propose a responsible workflow for combining AI insights, real-world market research, financial analysis, and team judgement.
Main principles:	reliability and verification, transparency, accountability, human oversight, critical judgement and responsibility

Case Study: 13	
Title: WHEN AN AI SYSTEM TREATS GROUPS DIFFERENTLY	
Description:	A university uses an AI-based assessment system to support the evaluation of students' academic performance. After several weeks of use, students begin to notice a consistent pattern: non-native speakers appear to receive lower scores than other students, even when their demonstrated subject knowledge seems comparable. The university maintains that the system improves the efficiency and consistency of assessment. However, there is limited information about which features or characteristics the AI system considers when generating scores, making it difficult to understand why these differences occur.
Key points for consideration:	- The assessment outcomes appear to differ systematically between groups of students. - Language background may be influencing scores in ways that are unrelated to actual subject knowledge. - The AI system's decision-making process and the features used to generate scores are not sufficiently transparent. - It is unclear whether the observed differences result from biased data, biased system design, inappropriate assessment criteria, or another factor. - Students affected by the system may have a legitimate interest in understanding how their scores are determined.

	• Human oversight may be necessary to identify, investigate, and correct potentially unfair outcomes. • The university should balance the benefits of automated assessment with principles of fairness, accountability, transparency, and equal treatment.
Task:	1. Resolve the case from the student perspective. 2. Identify potential sources of bias that could cause non-native speakers to receive systematically lower scores. 3. Determine what evidence and data would be needed to investigate whether the AI system produces discriminatory or unfair outcomes. 4. Consider how the university could test whether differences in assessment results are statistically and educationally significant. 5. Assess what information about the AI system should be made available to students affected by its decisions. 6. Discuss what forms of transparency are necessary when an AI system is used to influence academic assessment. 7. Determine what role human reviewers should have in reviewing, challenging, or overriding AI-generated scores. 8. Consider what safeguards should be introduced to prevent or correct unfair outcomes. 9. Decide whether the university can reasonably consider the system fair and responsible in its current form, and justify your conclusion.
Main principles:	fairness, bias, transparency, explainability, accountability, human oversight.

Case Study: 14	
Title: PERSONALISATION OR PIGEONHOLING?	
Description:	An educational AI platform analyses students' academic performance, assignment history, and survey responses in order to recommend personalised learning materials. While some students find these recommendations useful and relevant, others notice that the system repeatedly directs them towards the same topics based on its predictions about their abilities and interests. As a result, they feel that the system may be limiting their opportunities to explore unfamiliar subjects or develop new areas of interest. Students are also uncertain about what personal data the platform collects, how long it is stored, how it is used, and whether it is shared with others.
Key points for consideration:	• AI-based personalisation may improve the relevance and efficiency of learning recommendations. • Personalisation can also create a risk of limiting students to topics that match their existing performance or predicted strengths. • The system may reinforce existing patterns or assumptions rather than supporting students' development and exploration. • Students should have a clear understanding of what personal data is collected and how it is used.

	• The collection and use of student data raise questions about consent, privacy, and data retention. • Students' autonomy may be affected if AI recommendations are presented as authoritative or difficult to challenge. • The recommendation system may contain biases resulting from the data, algorithms, or assumptions used to generate recommendations. • Transparency is important so that students can understand why particular materials are recommended. • Students may need the ability to reject, modify, or opt out of AI-generated recommendations. • Personalisation should support learning without unnecessarily restricting students' choices and opportunities.
Task:	1. Identify the potential benefits and risks of using AI to personalise students' learning experiences. 2. Determine what types of student data the system should be allowed to collect and analyse. 3. Consider what information students should receive about the collection, storage, use, and sharing of their personal data. 4. Discuss when students' consent should be required and what would constitute meaningful and informed consent. 5. Assess how AI-based recommendations could limit students' autonomy or create a "pigeonholing" effect. 6. Identify possible sources of bias in the recommendation system and consider how these could affect different groups of students. 7. Determine what level of transparency students should have regarding the reasons behind individual recommendations. 8. Discuss whether students should be able to reject, modify, or opt out of AI-generated recommendations. 9. Propose safeguards that would allow students to explore topics outside their predicted strengths or interests. 10. Decide under what conditions AI-based personalisation would be acceptable, fair, and responsible in an educational context.
Main principles:	privacy, autonomy, informed consent, fairness, transparency

Case Study: 15	
Title: AI-GENERATED ACCESSIBILITY CONTENT: SUPPORT OR SOURCE OF RISK?	
Description:	An educational institution uses an AI system to provide live captions and summaries of lectures and other learning activities in order to improve accessibility. The system is particularly valuable for students with hearing difficulties and for students who are learning in a second language, as it can make spoken content easier to follow and review. However, the system sometimes mistranscribes technical terms, omits important details, or removes nuance when generating summaries. These inaccuracies may lead

	students to misunderstand key concepts or rely on information that does not accurately reflect what was originally communicated.
Key points for consideration:	• AI-generated captions and summaries can significantly improve access to educational content. • The technology may provide particular benefits for students with hearing difficulties and students learning in a second language. • Inaccurate transcription of technical terminology can change the meaning of important academic content. • Summaries may omit context, qualifications, or nuance that is relevant to understanding the original material. • Students may not always be able to recognise when AI-generated content is inaccurate. • The level of verification required may depend on the importance and potential consequences of the information. • Users should be clearly informed that captions and summaries may contain errors and should understand their limitations. • Responsibility for inaccurate accessibility content should be clearly defined between the institution, educators, technology providers, and users. • Accessibility tools should support students without creating additional barriers or disadvantages. • Human review may be particularly important when the information is essential for learning, assessment, or decision-making.
Task:	1. Identify the main benefits of using AI-generated captions and summaries as accessibility tools. 2. Identify the potential risks associated with inaccurate transcriptions and summaries. 3. Determine in which situations AI-generated accessibility content can reasonably be relied upon. 4. Consider when human verification or correction should be required. 5. Identify what types of information, such as technical terminology, assessment instructions, or key academic concepts, require a higher level of accuracy. 6. Decide how users should be informed about the limitations and potential inaccuracies of AI-generated captions and summaries. 7. Consider whether students should have access to the original recording or other sources so they can verify potentially important information. 8. Determine whether students should be expected to identify and correct AI-generated errors themselves or whether the institution should provide adequate quality control. 9. Discuss what reasonable alternatives should be available when AI-generated accessibility content is inaccurate or unavailable. 10. Determine who should be responsible when inaccurate AI-generated content causes a student to misunderstand important information. 11. Consider how responsibility should be shared between the educational institution, educators, AI system provider, and other relevant parties.

	12. Propose safeguards and quality-control measures that could reduce the risk of harmful inaccuracies. 13. Consider whether the use of AI-generated accessibility tools could unintentionally disadvantage the very students they are intended to support. 14. Decide what standards of accuracy, transparency, human oversight, and accountability should apply to AI-generated accessibility content. 15. Decide under what conditions AI-generated accessibility content can be considered sufficiently reliable, fair, and responsible for use in education.
Main principles:	inclusion, reliability, transparency, proportionality, human oversight

Case Study: 16	
Title: RESPONSIBLE AI IN PRACTICE: ETHICAL ANALYSIS AND DECISION-MAKING	
Description:	An educational institution uses an AI chatbot to provide students with 24/7 information about institutional regulations, deadlines, procedures, and available support services. The chatbot helps reduce waiting times and allows students to access information outside normal working hours. However, some students receive outdated or inaccurate information, while others are concerned about sharing personal or sensitive information with the chatbot. It is also unclear to students when they are interacting with an AI system, whether their conversations are stored, and who is responsible if the chatbot provides incorrect advice.
Key points for consideration:	• An AI chatbot can provide students with fast and accessible information at any time. • AI assistance may reduce pressure on administrative staff and waiting times for routine enquiries. • Incorrect or outdated information can have serious consequences when it concerns deadlines, regulations, eligibility, or institutional procedures. • Not all student questions are equally appropriate for an AI system to answer. • Questions involving individual circumstances, sensitive information, or significant consequences may require human assistance. • Students should be clearly informed when they are interacting with an AI system. • Personal and sensitive student information should be collected only when necessary and handled securely. • Students should know what data is stored, why it is collected, how long it is retained, and who can access it. • Information provided by the chatbot should be based on reliable and regularly updated institutional sources.

	• There should be a clear process for reporting incorrect information and escalating complex cases to a human. • Responsibility and accountability for incorrect or harmful advice should be clearly established.
Task:	1. Identify the types of questions and requests that are appropriate for an AI student adviser. 2. Identify questions or situations that should automatically be escalated to a human adviser. 3. Consider whether questions involving personal, sensitive, or confidential information should be handled by an AI chatbot. 4. Determine what students should be told about the chatbot's use of AI and its limitations. 5. Decide how clearly the system should identify itself as an AI rather than a human adviser. 6. Determine what personal data the chatbot should be permitted to collect and under what conditions. 7. Consider how students' personal information and chat histories should be protected. 8. Decide how long information provided by the chatbot should be retained and who should have access to it. 9. Determine how the accuracy and currency of information provided by the chatbot should be verified. 10. Consider how frequently the chatbot's information sources should be reviewed and updated. 11. Propose a process for students to report inaccurate, outdated, or misleading information. 12. Determine when a human adviser should review or take over a conversation. 13. Consider what safeguards are needed when incorrect advice could affect a student's academic progress, rights, deadlines, or access to services. 14. Determine who should be accountable when the chatbot provides incorrect or harmful advice. 15. Consider how responsibility should be shared between the institution, human staff, and the AI system provider. 16. Decide what level of human oversight is necessary for the chatbot to be considered responsible and appropriate for student support. 17. Propose conditions under which the AI adviser could be used safely while preserving students' privacy, access to human support, and ability to make informed decisions.
Main principles:	transparency, privacy, reliability, accountability, appropriate human oversight

Case Study: 17

Title: **HOW MUCH OF THE PROJECT IS REALLY MINE?**

Description:	A student is preparing an individual university project that includes data analysis, computer code, graphs, and a written report. A generative AI assistant is capable of producing substantial parts of all four components. While some of the AI-generated code appears to work correctly, other parts produce results that the student does not fully understand or cannot independently explain. The student must decide how AI can be used without undermining the purpose of an individual assessment or misrepresenting AI-generated work as their own.
Key points for consideration:	• Generative AI can support students with coding, data analysis, visualisation, and academic writing. • AI-generated work may contain errors even when it appears to function correctly. • A student may be responsible for submitted work that they cannot fully explain or verify. • The use of AI may raise questions about authorship, academic integrity, and independent learning. • Different types of AI assistance may carry different levels of risk and may be subject to different university rules. • Students should understand and verify the methods, calculations, code, results, and conclusions included in their submission. • AI assistance may need to be disclosed to ensure transparency about how the work was produced. • The student should not rely on AI-generated content simply because it produces plausible results. • Academic responsibility ultimately remains with the student for the work they submit. • Clear institutional guidance is needed regarding acceptable and unacceptable uses of generative AI in individual assessments.
Task:	1. Identify which forms of generative AI assistance could reasonably be considered acceptable for an individual university project. 2. Distinguish between AI use that supports the student's own work and AI use that effectively completes the assessed work on the student's behalf. 3. Determine which parts of the project the student must personally understand and be able to explain. 4. Consider what level of understanding is necessary when using AI-generated code or data-analysis methods. 5. Determine how the student should verify that AI-generated code, calculations, graphs, and results are correct. 6. Consider what risks arise when AI-generated code produces apparently correct results that the student cannot explain. 7. Decide which uses of AI should be disclosed in the final submission. 8. Consider what information about the AI assistance should be disclosed, such as the type of tool used and the nature of its contribution. 9. Determine whether AI-generated text, code, analysis, or graphs should be treated differently for purposes of authorship and academic integrity.

	10. Consider what the student should do if they discover that an AI-generated part of the project cannot be explained or independently verified. 11. Decide whether the student should remove, replace, or substantially revise AI-generated material that they do not understand. 12. Consider who is responsible for errors, inaccurate results, or misleading conclusions contained in AI-assisted work. 13. Propose practical safeguards that students can use to maintain control over their work while benefiting from AI assistance. 14. Consider how university policies should define acceptable, restricted, and prohibited uses of generative AI in individual assessments. 15. Decide under what conditions the final project can fairly be considered the student's own work when generative AI has been used.
Main principles:	authorship, disclosure, verification, learning outcomes, responsibility

Case Study: 18	
Title: THE PERFECT ESSAY WITH FAKE SOURCES	
Description:	A student uses a generative AI system to help structure an academic essay and identify potentially relevant academic literature. The AI produces a well-written and convincing draft, including quotations, references, and claims that appear academically credible. However, when the student checks the references, they discover that several quotations do not accurately reflect the original sources and that some of the cited publications cannot be found or may not exist. The student must decide how to handle the AI-generated material before submitting the essay and how to ensure that the final work is based on reliable academic evidence.
Key points for consideration:	• Generative AI can be useful for brainstorming, structuring an essay, and identifying potentially relevant areas of research. • AI systems can generate plausible but inaccurate quotations, references, citations, and factual claims. • A source mentioned by an AI system should not automatically be assumed to exist or to support the claim attributed to it. • Academic sources must be independently verified using reliable and appropriate research databases or original publications. • Quotations should be checked against the original source for accuracy and context. • Students remain responsible for the accuracy and integrity of the sources cited in their work. • An apparently convincing AI-generated response does not constitute academic evidence in itself. • Using fabricated or inaccurate references without verification can undermine the credibility and academic integrity of the entire assignment. • AI assistance should not replace the student's own critical evaluation of academic sources. • The final submission should contain only sources and claims that the student has been able to verify.

Task:	1. Identify the main risks associated with using generative AI to identify academic literature and produce references. 2. Determine which parts of the AI-generated essay should be checked before they are considered for inclusion in the final submission. 3. Develop a responsible process for verifying every source suggested by the AI system. 4. Explain how the student should verify that a cited publication actually exists and is an appropriate academic source. 5. Determine how quotations should be checked against the original publication. 6. Consider how the student should verify that a source genuinely supports the claim for which it is cited. 7. Decide what should be done with sources that cannot be located, quotations that are inaccurate, or claims that cannot be verified. 8. Consider whether AI-generated references can ever be used without independent verification. 9. Explain why a convincing or authoritative-looking AI response cannot itself be treated as academic evidence. 10. Determine what responsibility the student has for errors or fabricated sources generated by an AI system. 11. Consider whether and how the student should disclose the use of generative AI in preparing the essay
Main principles:	hallucinations, source verification, academic integrity, critical evaluation

Case Study: 19	
Title: CAN AI-GENERATED DATA REPLACE A REAL EXPERIMENT?	
Description:	A university assignment requires students to conduct a small empirical study and collect real-world observations before analysing the results. However, the student is struggling to recruit enough participants and is concerned about completing the assignment on time. A generative AI system offers to create a realistic synthetic dataset that could be used to perform the required analysis. The student must decide whether using AI-generated data would be an acceptable solution or whether it would undermine the purpose of the assignment, particularly if the task is intended to assess the student's ability to collect and work with empirical evidence.
Key points for consideration:	• Synthetic data can be useful for practising data analysis, testing methods, or conducting simulations. • Synthetic data is not equivalent to observations collected from real participants or real-world events. • Using AI-generated data may undermine an assignment if collecting empirical data is itself part of the learning objective. • The distinction between an empirical study and a simulation should be clearly maintained.

	• Students should not present synthetic observations as if they were collected from real participants. • Any permitted use of synthetic or simulated data should be clearly disclosed. • The origin and nature of the dataset should be transparent to the instructor and readers of the work. • AI-generated data may contain unrealistic patterns, hidden assumptions, or biases that affect the analysis. • The acceptability of synthetic data depends partly on the purpose and requirements of the assignment. • Clear labelling is necessary if synthetic data is permitted so that it cannot be mistaken for empirical evidence.
Task:	1. Identify the purpose of the assignment and determine whether collecting real observations is an essential part of the learning objective. 2. Explain the difference between empirical data collected from real-world observations and synthetic data generated through AI or simulation. 3. Assess whether replacing the required real observations with AI-generated data would change the nature of the assignment. 4. Identify the potential academic integrity issues that could arise if synthetic data were presented as real observations. 5. Determine under what circumstances the use of synthetic data could be acceptable in an academic assignment. 6. Consider whether the student should seek permission from the instructor before using synthetic or simulated data. 7. Decide what information about the origin and generation of the dataset should be disclosed. 8. Determine how synthetic data should be labelled in the assignment so that it cannot be mistaken for empirical evidence. 9. Consider what limitations and risks should be acknowledged when analysing AI-generated data. 10. Assess whether the student should verify the properties, assumptions, and characteristics of the synthetic dataset before using it. 11. Consider what alternatives the student could pursue if recruiting participants proves difficult. 12. Decide whether completing the assignment on time justifies replacing real observations with synthetic data. 13. Propose how an instructor could distinguish between an empirical study and a simulation-based assignment
Main principles:	research integrity, fabrication, transparency, scientific validity

Case Study: 20	
Title: CAN AN AI DETECTOR PROVE CHEATING?	
Description:	A university student submits an essay that is flagged by an automated AI-detection system as likely to have been generated by AI. The essay's writing style is noticeably different from the student's previous work,

	which raises concerns about whether generative AI was used without permission. However, there is no direct evidence showing that the student used AI to produce the submitted content. The institution must decide whether the AI-detection result is sufficient to establish academic misconduct and what additional evidence and procedures should be considered before reaching a conclusion.
Key points for consideration:	• AI-detection systems provide an indication or probability, rather than definitive proof of AI use. • Detection tools can produce false positives and may incorrectly identify human-written work as AI-generated. • Differences in writing style do not necessarily prove that AI was used. • Students may have legitimate explanations for changes in writing style or language. • Relying heavily on automated detection may create risks of unfair treatment or discrimination. • Students should have an opportunity to understand and challenge allegations made against them. • Academic misconduct procedures should provide appropriate due process and a fair opportunity to respond. • Additional evidence may be necessary before concluding that a student has violated academic integrity rules. • Institutions should be transparent about how AI-detection tools are used and what weight their results carry. • A balance is needed between protecting academic integrity and avoiding unjustified accusations based solely on probabilistic technology.
Task:	1. Explain what an AI-detection score actually indicates and what it cannot establish. 2. Assess whether an AI-detection result alone can reasonably be treated as proof of academic misconduct. 3. Identify the potential causes and consequences of false-positive AI detections. 4. Consider whether differences between a student's current and previous writing styles constitute sufficient evidence of AI use. 5. Identify what additional evidence could reasonably be considered before reaching a conclusion. 6. Consider whether drafts, notes, version history, references, or records of the student's research process could provide relevant evidence. 7. Determine what opportunity the student should have to explain or challenge the AI-detection result. 8. Consider what due-process safeguards should apply when automated systems are used in academic misconduct investigations. 9. Decide how much weight an AI-detection score should have compared with other available evidence. 10. Consider whether students should be informed that their work will be assessed using an AI-detection system. 11. Determine what information the institution should disclose about the detector, including its limitations and expected accuracy.

	12. Identify potential fairness concerns when AI-detection systems are used across students with different linguistic, educational, or writing backgrounds. 13. Consider who should make the final decision about whether academic misconduct has occurred. 14. Propose a fair and transparent procedure for investigating suspected unauthorised AI use. 15. Decide what level and combination of evidence should be required before concluding that a student has committed academic misconduct.
Main principles:	fairness, transparency, evidence and verification, academic integrity, human oversight and accountability

Case Study: 21	
Title: AI ASSISTANT OR AI SUBSTITUTE?	
Description:	Generative AI tools can now support students in many different ways, including brainstorming essay ideas, solving mathematical equations, generating computer code, rewriting text, and explaining difficult concepts. A student is preparing an assessed university assignment and has access to all of these capabilities. The student must decide which forms of AI assistance can support their own learning and work, which uses should be disclosed, and when AI begins to replace the student's own contribution to the assessed task.
Key points for consideration:	• AI can provide different levels and types of assistance, ranging from explanations and brainstorming to producing complete answers. • The educational value of AI use depends on whether it supports the student's learning or replaces the work the assessment is intended to measure. • Some forms of AI assistance may be acceptable but still require disclosure. • Other uses may be inappropriate because they involve submitting AI-generated work as the student's own. • The same AI tool may be appropriate in one assignment but inappropriate in another, depending on the assessment requirements. • Students remain responsible for the accuracy, quality, and integrity of the work they submit. • Students should understand the boundaries established by their institution, instructor, or assessment instructions • Transparency about AI use can help distinguish legitimate assistance from unauthorised substitution. • A clear boundary is needed between using AI as a learning support and using it to complete assessed work on the student's behalf.
Task:	1. Identify different ways in which AI could be used when preparing an assessed assignment. 2. Classify each use as acceptable assistance, assistance requiring disclosure, or inappropriate substitution of the student's own work.

	3. Consider whether brainstorming ideas with AI should be treated differently from asking AI to write the complete assignment. 4. Determine whether using AI to explain a difficult concept is different from using AI to produce the answer that will be submitted for assessment. 5. Assess whether using AI to solve equations is acceptable when the student independently understands and can explain the solution. 6. Consider when AI-generated code constitutes legitimate assistance and when it replaces the programming work the student is expected to demonstrate. 7. Determine whether using AI to rewrite or improve the student's own text changes the authorship or originality of the final work. 8. Consider which forms of AI assistance should be disclosed and what such disclosure should include. 9. Identify factors that could change whether a particular use of AI is acceptable, such as the purpose of the assignment, learning outcomes, or assessment rules. 10. Consider what responsibility the student has for checking, correcting, and understanding AI-generated content before submitting it. 11. Identify the risks of relying on AI-generated answers, calculations, code, or explanations that the student cannot independently explain or verify. 12. Consider whether the amount of AI-generated content should affect whether its use is considered acceptable. 13. Discuss whether using AI to generate an initial draft and then substantially revising it is different from submitting largely AI-generated work. 14. Consider how the boundary between assistance and substitution may differ for different types of assignments and assessment methods. 15. Propose a clear principle for distinguishing AI as a learning assistant from AI as a substitute for the student's own work. 16. Develop practical criteria that a student could apply before using AI for an assessed task. 17. Decide what evidence of independent contribution or understanding a student should be able to provide if their use of AI is questioned. 18. Consider what responsibilities instructors and universities have to communicate clear and consistent rules about AI use. 19. Evaluate the consequences of treating all AI use as either acceptable or unacceptable without considering the purpose and context of the assignment. 20. Decide where you would draw the boundary between acceptable assistance, use requiring disclosure, and inappropriate substitution, and justify your decision.
Main principles:	appropriate use, disclosure, authorship, independent learning

Case Study: 22

Title: **THE CODE WORKS – BUT DO YOU UNDERSTAND IT?**

Description:	A student uses a large language model (LLM) to generate most of the code for a university programming assignment. The resulting program works as expected, and the accompanying documentation appears professional and well structured. However, while preparing for the oral defence, the student realises that they cannot adequately explain several important concepts used in the program or justify why certain implementation decisions were made. The student must decide whether the assignment is ready for submission and how the use of AI should be addressed.
Key points for consideration:	• Functioning code does not necessarily demonstrate that the student understands the programming concepts being assessed. • An individual programming assignment may be intended to assess problem-solving, programming ability, and conceptual understanding, not simply the final output. • AI-generated code may contain design choices that the student does not understand or would not be able to reproduce independently. • The student remains responsible for understanding, testing, and verifying code included in their submission. • AI assistance may be legitimate when it supports learning, debugging, explanation, or experimentation. • Using AI to produce most of the assessed work may cross the boundary between assistance and substitution. • The extent and nature of AI assistance may need to be disclosed according to institutional or course requirements. • The student should be able to explain the important concepts, algorithms, design choices, and functionality of the submitted code. • Unexplained AI-generated code may create problems during an oral defence or other assessment of individual understanding. • AI can be used more responsibly as a learning aid when the student actively analyses, tests, modifies, and learns from the generated code.
Task:	1. Identify the purpose and learning outcomes of the programming assignment. 2. Assess whether functioning code alone is sufficient evidence that the student has achieved the intended learning outcomes. 3. Identify which parts of the submitted program the student must personally understand and be able to explain. 4. Determine whether the student should be able to justify the main algorithms, design choices, libraries, functions, and implementation decisions used in the code. 5. Consider what the student should do with code that works but cannot currently be explained or understood. 6. Decide whether the student should rewrite, simplify, test, or otherwise modify AI-generated code before submission. 7. Consider how the student can use the AI system to learn the concepts behind the generated code rather than simply copying its output. 8. Identify appropriate uses of AI as a learning aid, such as requesting explanations, examples, debugging support, or alternative implementations.

	9. Determine which forms of AI assistance should be documented or disclosed in the assignment. 10. Consider what information an AI-use statement should contain, including the purpose and extent of AI assistance. 11. Assess the risks of submitting substantial amounts of AI-generated code that the student cannot independently explain. 12. Consider how the inability to explain parts of the code should affect the student's decision to submit the assignment. 13. Determine what evidence could demonstrate that the student has developed sufficient understanding of the submitted code. 14. Propose practical steps the student could take before the oral defence to ensure genuine understanding of the program. 15. Consider how instructors could distinguish between AI-assisted programming and inappropriate substitution of the student's own work. 16. Develop a principle for using AI in programming assignments that supports learning without replacing the skills and knowledge being assessed. 17. Decide what the student should do before submission and justify the decision.
Main principles:	conceptual understanding, transparency, responsibility for generated code

Case Study: 23	
Title: AI-GENERATED CURRICULUM ATTRIBUTION AND INTELLECTUAL PROPERTY DISPUTE	
Description:	A university department utilizes a Generative AI tool to draft a new microcredential module on "Sustainable Urban Planning." The AI produces a comprehensive syllabus, including learning outcomes, assessment rubrics, and reading lists, based on a prompt provided by a junior lecturer. The department publishes the microcredential without disclosing that the core structure was AI-generated. Later, a student challenges the validity of the credentials, claiming the content lacks human oversight and "authorship." Furthermore, it is discovered that the AI model was trained on copyrighted textbooks from a competing university without permission or compensation. The department now faces a dilemma: they must decide how to attribute the work, how to verify the accuracy of the AI-generated content, and how to address the ethical implications of using unlicensed data to build a public educational product.
Key points for consideration:	• Legal and moral authorship of AI-generated content. • The impact of non-disclosure of AI involvement on the perceived value and credibility of microcredentials. • Institutional responsibility for verifying the factual accuracy and pedagogical integrity of AI-produced materials. • Privacy, fairness, and copyright implications of using competitors' materials to train AI models.

	The impact of algorithmically derived curricula on student autonomy and agency.
Task:	Resolve the case from the student perspective. - Evaluate the situation based on the Ethical Guidelines. - Propose a new "Transparency Protocol" for the university that defines how AI tools can be used in curriculum development while ensuring clear authorship, informed consent for students, and fair use of intellectual property.
Main principles:	authorship, disclosure, verification, transparency, fairness, responsibility, informed consent.

Case Study: 24	
Title: SYNTHETIC DATA GENERATION AND THE PARADOX OF STUDENT PRIVACY	
Description:	A university research team is developing a microcredential for "Data Science Ethics." To provide students with high-quality practice, they decide to use a Large Language Model (LLM) to generate a massive synthetic dataset that mimics real patient records from a local hospital. The goal is to ensure that no actual patient identities are compromised, thereby upholding privacy. However, the team discovers that the LLM occasionally "hallucinates" realistic but non-existent medical histories that correlate with specific, real-world demographic anomalies. Students are tasked with using this synthetic data to build predictive models. A controversy arises when a student discovers that the "synthetic" patterns inadvertently reveal sensitive socio-economic trends of a specific marginalized neighborhood, leading to questions about whether the data is truly anonymous. Furthermore, because the data was AI-generated, the team struggles to define authorship of the dataset for academic credit. The department must decide if the students can provide informed consent to use "hallucinated" data that may still carry proxy biases, and who holds the responsibility if a student's model produces discriminatory outcomes based on this synthetic input.
Key points for consideration:	- Ownership and authorship of AI-generated datasets versus human-curated data. - The risk of "re-identification" or proxy bias in synthetic data. - Transparency regarding the limitations and "hallucination" risks of AI tools. - Balancing the need for realistic training data with the absolute right to privacy. - Informed consent when the source of the training data is computationally opaque.
Task:	Resolve the case from the student perspective. - Ownership and authorship of AI-generated datasets versus human-curated data. - The risk of "re-identification" or proxy bias in synthetic data.

	3. Transparency regarding the limitations and "hallucination" risks of AI tools. 4. Balancing the need for realistic training data with the absolute right to privacy. 5. Informed consent when the source of the training data is computationally opaque. 6. Analyze the tension between data utility and individual privacy. 7. Propose a framework for "Verified Synthetic Data" that includes a transparency report for students. 8. Determine who is responsible (the researchers, the AI provider, or the students) if a model trained on this data demonstrates unfair bias.
Main principles:	transparency, accountability, human oversight, accuracy and reliability, fairness and non-discrimination, privacy and data protection, authorship and intellectual property, informed consent, human autonomy

Case Study: 25	
Title: AI-GENERATED RESEARCH SYNTHESSES AND ACADEMIC INTEGRITY	
Description:	A postgraduate student is working on a microcredential project involving a literature review of complex ethical guidelines. To save time, the student uses a generative AI tool to synthesize thirty different research papers into a cohesive summary. The AI produces a high-quality summary that includes accurate data but attributes the overarching synthesis as the student's original work. The student submits this summary without disclosing the use of the AI tool. Furthermore, the AI tool used for this synthesis was trained on proprietary datasets without the original authors' explicit permission, raising questions about the provenance of the information provided to the student. The student's supervisor praises the synthesis but does not know that the core analytical structure was produced by an algorithm, not the human learner. This creates a conflict regarding who holds the responsibility for the accuracy of the content and whether the student's learning outcomes were truly met through independent critical thinking.
Key points for consideration:	• Disclosure of AI use in academic work • Authorship and attribution of AI-assisted content • Responsibility for the accuracy and integrity of the submitted work • Verification of AI-generated synthesis and interpretation • Transparency regarding the provenance and training data of AI tools • Copyright and intellectual property implications • Authenticity of the student's learning and critical-thinking process • Alignment between AI-assisted work and intended learning outcomes • Academic integrity and institutional expectations • The role of human oversight in AI-assisted academic work
Task:	Resolve the case from the student perspective.

	1. What are the ethical implications of using generative AI to synthesize academic literature? 2. Who should be responsible for the accuracy and integrity of AI-assisted academic work? 3. How should authorship and attribution be addressed when AI contributes to the analytical structure of the work? 4. Does the use of AI affect whether the student has achieved the intended learning outcomes, particularly independent critical thinking? 5. What level of disclosure should be required when AI is used in academic work? 6. How should AI-generated content be verified before it is submitted as academic work? 7. What are the implications of using AI tools trained on proprietary datasets without explicit author permission? 8. How can institutions balance the benefits of AI-assisted learning with academic integrity and human agency?
Main principles:	transparency, accountability, human oversight, academic integrity, authorship and attribution, intellectual property, data provenance, human autonomy, fairness and responsibility

Case Study: 26	
Title: AI-ASSISTED RESEARCH AND ACADEMIC AUTHORSHIP	
Description:	A graduate student in the Microcredentials program uses a sophisticated Large Language Model (LLM) to synthesize vast amounts of qualitative data from various peer-reviewed journals to produce a comprehensive literature review. The student prompts the AI to "identify key themes and summarize the findings in the style of a professional academic paper." The resulting output is highly polished and accurate, covering complex nuances that would have taken weeks to manually synthesize. The student submits the paper for a final assessment. The instructor notices that while the content is excellent, the "voice" of the writing feels inconsistent with the student's previous work. The student argues that because they provided the prompts, curated the source list, and edited the final output for clarity, they are the sole author. However, the AI's underlying training data is not disclosed, and it is unclear which specific insights were generated by the model versus the student's original critical thinking. The department must decide how to attribute authorship, the level of disclosure required for AI assistance, and how to ensure the student actually met the learning outcomes regarding critical synthesis.
Key points for consideration:	• When does AI assistance become a primary contributor or author rather than simply a supporting tool? • How does a lack of transparency about AI use affect the ability to verify a student's learning and achievement? • To what extent can prompt engineering be considered a sufficient academic contribution to establish authorship?

	What risks to academic integrity arise when the source and process of a synthesis are based on an opaque algorithm?
Task:	- Evaluate the student's submission based on the Ethical Guidelines. - Determine what specific disclosure statement should have been included and propose a grading rubric that distinguishes between "AI-assisted synthesis" and "student-led analysis."
Main principles:	authorship, disclosure, verification, learning outcomes, transparency, responsibility

Case Study: 27	
Title: ALGORITHMIC BIAS IN AUTOMATED STUDENT PROGRESS TRACKING	
Description:	A university adopts an AI-driven predictive analytics tool to monitor student engagement and predict "at-risk" students for early intervention. The system analyzes data points such as login frequency, time spent on modules, and forum participation. However, some students disadvantaged students notice that the algorithm frequently flags them as "disengaged" due to inconsistent internet access or non-traditional study schedules, while high-performing students with similar patterns are not flagged. The university has not disclosed the specific weights used by the algorithm, and students have no way to contest the automated "risk" status. The department wants to use this data to limit access to supplementary tutoring resources, which are currently first-come, first-served.
Key points for consideration:	- The "black box" nature of the algorithm prevents students from understanding why they were flagged. - The tool may perpetuate systemic inequalities by penalizing students with different socioeconomic realities. - Does the automated labeling infringe on a student's ability to manage their own learning path? - Were students aware that their behavioral metadata would be used for predictive profiling rather than just instructional support? - Who is accountable if a student is denied support based on a flawed algorithmic prediction?
Task:	- Develop a policy framework that ensures the AI tool operates with transparency and fairness. - Propose a mechanism for students to appeal an automated "at-risk" designation and outline how the university will balance data-driven efficiency with individual student privacy and rights.
Main principles:	transparency, fairness, informed consent, responsibility, privacy, autonomy

Case Study: 28	
Title: THE AI SHORTLIST SAYS THESE ARE THE BEST CANDIDATES: SHOULD YOU TRUST IT?	
Description:	You are completing an internship in the Human Resources department of a medium-sized company. The company has introduced an AI recruitment system that analyses CVs, cover letters and online assessment results and produces a ranked shortlist for each vacancy. The tool is fast and helps the team manage hundreds of applications. However, while reviewing the results, you notice that applicants with career gaps, foreign qualifications or less conventional career paths often appear near the bottom of the ranking, even when they seem to meet the stated job requirements. The recruiter is under time pressure and proposes interviewing only the highest-ranked candidates. The vendor provides only limited information about how the ranking is produced. You must decide how the AI-generated shortlist should be used, what evidence should be checked before it influences recruitment decisions, and what role human judgement, transparency and fairness should play in the process.
Key points for consideration:	• Potential bias and discrimination in AI-supported recruitment. • Job relevance and quality of the data used to rank candidates. • Transparency about the role and limitations of the AI system. • Meaningful human review rather than automatic acceptance of rankings. • Privacy, accountability and the ability to question or correct an AI-supported decision.
Task:	Resolve the case from the student perspective. 1. Identify the main benefits and ethical risks of using the AI-generated shortlist. 2. What information would you request about the data, criteria and validation of the recruitment system before relying on its rankings? 3. Which candidate characteristics should be considered relevant to the job, and which could create unfair or unjustified disadvantages? 4. How should the HR team review candidates who receive low AI scores but appear to meet the job requirements? 5. What should applicants be told about the role of AI in the recruitment process and the possibility of human review? 6. What evidence would you collect to investigate whether the system produces systematically different outcomes for particular groups? 7. What should the company do if the system is efficient overall but shows signs of unfair or discriminatory effects? 8. Give your final resolution: propose a responsible workflow in which AI supports recruitment without replacing accountable human judgement.
Main principles:	fairness and non-discrimination, transparency, privacy and data protection, human oversight, accountability

Case Study: 29	
Title: THE AI MARKETING CAMPAIGN KNOWS WHAT EACH CUSTOMER IS LIKELY TO BUY: HOW FAR SHOULD PERSONALISATION GO?	
Description:	You are completing a marketing internship at an online retailer. The company uses an AI system to analyse customers' purchase history, browsing behaviour, location, responses to previous campaigns and selected customer-service interactions. The system creates customer segments, predicts what each person is likely to buy and automatically drafts personalised emails, product recommendations and special offers. The campaign performs well and sales increase. As you review the outputs, you notice several concerns. Some messages use highly personalised information that customers may not expect to influence advertising. Certain segments receive stronger pressure-based messages and different offers, while some AI-generated product claims have not been checked against the company's official information. A customer also asks whether the message they received was written by a person or by AI. You must decide how to preserve the legitimate benefits of personalisation while protecting customers from intrusive, misleading or unfair practices.
Key points for consideration:	• Privacy, purpose limitation and data minimisation in customer profiling. • Fairness and the risk of intrusive or manipulative personalisation. • Accuracy and verification of AI-generated marketing claims and offers. • Transparency when AI is used to generate or personalise customer communications. • Human responsibility for campaigns, complaints and high-impact customer interactions.
Task:	Resolve the case from the student perspective. 1. Which customer data would you consider appropriate for personalisation, and which data uses would require stronger justification or should be avoided? 2. Where would you draw the line between useful personalisation and intrusive or manipulative targeting? 3. How should the company verify AI-generated product claims, prices, discounts and conditions before messages are sent? 4. Should customers be informed when AI is used to generate or personalise communications? Explain what information would be meaningful. 5. What should happen if an AI-generated message makes a false promise or gives a customer inaccurate information? 6. When should an automated customer interaction be escalated to a human employee? 7. How would you check whether customer segments or personalised offers produce unfair outcomes for particular groups? 8. Give your final resolution: propose a responsible workflow for AI-supported marketing and customer relations that preserves customer autonomy, trust and human accountability.
Main principles:	privacy and data protection, transparency, fairness, reliability and accuracy, customer autonomy, human oversight and accountability

Case Study: 30	
Title: THE AI MODEL RECOMMENDS DENYING A LOAN: CAN A SCORE MAKE THE DECISION?	
Description:	You are completing a placement in the credit-risk team of a bank or financial technology company. The organisation uses an AI system to assess loan applications and provide a risk score and recommendation to the human analyst. One applicant is a self-employed person seeking financing for their small business. The system gives a low score and recommends rejection. The model considers information such as repayment history, existing debt, recent cash flow and other financial indicators, but the explanation for the final score is limited. The applicant challenges the result and provides new information showing that a recent drop in income was temporary and that a signed contract is expected to improve future cash flow. You also learn that analysts rarely override the AI recommendation because the system has performed well on average. You must decide what meaningful human review should involve, how the applicant's new evidence should be considered, and whether a statistically accurate model can still create unfair or poorly justified individual decisions.
Key points for consideration:	• High-impact AI-supported financial decisions and the risk of automation bias. • Fairness, bias and the relevance of data used to assess creditworthiness. • Accuracy, explainability and the ability to contest an AI-supported recommendation. • Meaningful human review of individual circumstances and new evidence. • Accountability for the final lending decision.
Task:	Resolve the case from the student perspective. 1. Identify the main ethical and professional risks in relying on the AI recommendation. 2. Which input data should be checked for relevance, accuracy and currency before the score is used? 3. What information should the human analyst examine in addition to the AI score before making a final decision? 4. How should the applicant's new evidence be reviewed, and when should the AI recommendation be reconsidered or overridden? 5. What explanation should be provided to a person affected by an AI-supported credit decision? 6. How could the organisation test whether the model produces systematically unfair outcomes for particular groups? 7. What should the organisation do if the model is accurate overall but performs worse for some groups or fails to capture important individual circumstances? 8. Give your final resolution: propose a responsible decision process that combines AI-supported risk analysis with meaningful human judgement, review and accountability.
Main principles:	fairness, transparency and explainability, privacy and data quality, human oversight, accountability, contestability

Case Study: 31	
Title: WHEN AI RECOMMENDS A MAJOR BUSINESS DECISION: SHOULD MANAGEMENT TRUST THE FORECAST?	
Description:	You are working as a junior business analyst in a medium-sized company that is considering entering a new international market. The proposed expansion would require a substantial investment in new facilities, staff, marketing, logistics, and local partnerships. To support the decision, the management team uses an AI-based decision-support system. The system analyses the company's historical sales data, publicly available market information, customer trends, competitor activity, and economic indicators. It produces a highly positive five-year forecast and recommends proceeding with the investment immediately. The AI-generated report predicts strong demand, rapid revenue growth, and a relatively short period for recovering the initial investment. The presentation is detailed and convincing, and several senior managers believe that the recommendation provides an objective and data-driven basis for the decision. However, while reviewing the report, you notice that it does not clearly explain which data sources were used, how recent the market information is, or what assumptions were made about future economic conditions. The report also gives limited information about uncertainty, alternative scenarios, or the possibility that market conditions may change. Some important factors, such as new regulatory requirements, local consumer behaviour, supply-chain risks, and recent competitor developments, appear to receive little attention. You also discover that some of the company's internal data used by the AI system may be incomplete or inconsistent. Sales figures from different departments have been recorded using different categories, and information from one regional office has not been updated for several months. The management team is expected to make its final decision within a few days. You must decide how to present your concerns without rejecting the potential value of AI-supported analysis. You must also determine what information should be verified, what additional evidence is needed, and how AI-generated recommendations should be combined with professional expertise and human judgement before the company makes a high-impact investment decision.
Key points for consideration:	• AI can support management by analysing large volumes of financial, market, operational, and customer data. • AI-generated forecasts may appear objective and authoritative even when they are based on incomplete, outdated, inconsistent, or biased data. • The quality of an AI recommendation depends on the quality, relevance, and completeness of the information used by the system. • Forecasts may rely on hidden assumptions about market growth, customer behaviour, economic conditions, costs, competitors, and regulatory developments. • A single prediction may conceal uncertainty and fail to represent alternative scenarios, including less favourable outcomes.

	• Management should understand the main factors that influenced the recommendation and the limitations of the model before relying on its output. • Important qualitative factors, such as local knowledge, regulatory changes, stakeholder concerns, and emerging risks, may not be adequately represented in the system. • Commercially sensitive and confidential company data must be protected when they are processed by an external AI system. • AI-generated recommendations should support rather than replace professional expertise, critical analysis, and managerial responsibility. • Employees should be able to question an AI recommendation without being expected to prove that the entire system is unreliable. • High-impact business decisions require appropriate human review, independent verification, and clearly assigned accountability. • Final responsibility for the investment decision must remain with the authorised human decision-makers and cannot be transferred to the AI system.
Task:	Resolve the case from the perspective of a junior business analyst advising the management team. 1. Identify the potential benefits of using an AI-based decision-support system when evaluating a major business investment. 2. Identify the main risks associated with accepting the AI-generated recommendation without additional review. 3. Determine which data sources, assumptions, calculations, and forecasts should be verified before the recommendation is used. 4. Explain how you would assess whether the company's internal data are sufficiently accurate, complete, consistent, and up to date. 5. Identify what additional market research, financial analysis, regulatory information, and professional expertise should be considered independently of the AI system. 6. Explain how the management team should evaluate uncertainty and develop alternative scenarios, including best-case, expected, and worst-case outcomes. 7. Consider whether the company should request an explanation of the main factors that influenced the AI recommendation. What level of transparency would be necessary for responsible decision-making? 8. Explain how you would communicate your concerns to senior management without making unsupported claims or dismissing the potential value of the AI system. 9. Determine what information about the AI tool, its data sources, assumptions, limitations, and level of uncertainty should be documented before a final decision is made. 10. Identify the privacy, confidentiality, and security safeguards required when internal company data are processed by an AI system or an external technology provider. 11. Decide what role financial specialists, market experts, operational managers, legal advisers, and local stakeholders should have in reviewing the recommendation.

	12. Consider whether the decision should be postponed if important data cannot be verified or if the system's assumptions remain unclear. Justify your position. 13. Determine who should be accountable if the company follows the AI recommendation and the investment produces serious financial or operational losses. 14. Propose a responsible decision-making workflow that combines AI-generated analysis, verified data, independent research, expert assessment, alternative scenarios, and human judgement. 15. Give your final recommendation to the management team and explain the conditions that should be met before the company proceeds with the investment.
Main principles:	reliability and accuracy, transparency and explainability, accountability, data quality, privacy and confidentiality, proportionality, critical judgement, human oversight and responsibility.

APPENDIX 2 CASE STUDY RESPONSE TEMPLATE

Case Study No: _____ & Title: _____

1. Understanding the AI Use

Briefly describe how AI is being used in this case.

Question	Response
What task or problem is the AI system being used to support?	
Who is using the AI system or affected by its use?	
Is the proposed use of AI appropriate for this situation? Why or why not?	

2. Potential Benefits

Identify the main benefits that the use of AI could provide in this situation. Consider, where relevant.

Benefit	Response
efficiency or time savings	
improved access to information	
personalisation or accessibility	
decision support	
improved productivity or quality	
other potential benefits	

3. Risks and Ethical Concerns

Identify the main problems or risks associated with the use of AI in this case. Consider, where relevant.

Problems or risks	Response
inaccurate or fabricated information	
bias or unfair treatment	
privacy and data protection	
confidentiality and security	
lack of transparency	
over-reliance on AI	
academic or professional integrity	
safety	
inappropriate automation	
unclear responsibility or accountability	

4. Verification, Transparency and Human Oversight

Explain what should be checked or verified before the AI-generated output is accepted or used.

Checked or verified before the AI-generated output is accepted or used	Response
What information or output needs to be verified?	

• What reliable sources, data or expertise could be used for verification?	
• Should the use of AI be disclosed? If yes, to whom and how?	
• At what point should a human review, question or override the AI output?	
• Who should remain responsible for the final decision or outcome?	

5. Your Analysis

Based on your analysis, explain what should be done in this situation.

Point	Response
whether AI should be used	
under what conditions it should be used	
what safeguards or changes are necessary	
what role humans should retain	
why you consider this approach responsible	

6. EmpowerAI Chatbot Activity

Ask the EmpowerAI chatbot at least one question related to the ethical or responsible use of AI in this case.

Question asked to the chatbot	Summary of the chatbot's response

7. Your evaluation of the chatbot response:

Question	Your evaluation
Do you agree with the chatbot's recommendation? Why or why not?	
Is there anything in the response that should be verified?	
Is any important consideration missing?	
Would you change your original recommendation after consulting the chatbot?	

8. Final Reflection

What are the key takeaways, and what did you learn from this case regarding the responsible use of artificial intelligence? (Describe in no more than three sentences.)

Response: