

FABE

CURRICULUM



**Seven connected modules
from human creativity to
responsible AI use**

ERASMUS+ · VET · AI LITERACY

FABE

CURRICULUM

*Fantasy / Imagination and Artificial Intelligence for Business
Enhancement*

Seven connected modules - from human creativity to responsible AI use

EDITORS

Andrea D'Andrea

Dr. Kamila Franz

Diba Mokhtabad Amrei

AUTHORS

Massimo Gerardo Carrese, Andrea D'Andrea, Dr. Kamila Franz, Dr. Krisztina Keresztély, Diba Mokhtabad Amrei, Francesca Olivier, Luca Puglisi

LANGUAGE EDITING

Dr. Kamila Franz

AI USE DISCLAIMER

Artificial intelligence tools were used during the development of this curriculum to support text editing, language refinement, formatting, and overall presentation. Several different AI tools were employed at various stages of the process. The content, concepts, and final editorial decisions remain the responsibility of the project team.

PARTNERS

Almost There AB - Coordinator

Gothenburg, Sweden · www.almostthere.se

On the Edge

Salerno, Italy · oteproject.eu

Comparative Research Network EV

Berlin, Germany · www.crnonline.de

CONTACT

On the Edge

<https://oteproject.eu/>

info@oteproject.eu

Via Dionisio Martino, 6

84128 Salerno SA

EU DISCLAIMER

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

The external links are being provided as a convenience and for informational purposes only; they do not constitute an endorsement or an approval by the project partners of any of the products, services or opinions of the corporation or organization or individual. The project team bears no responsibility for the accuracy, legality, or content of the external site or for that of subsequent links. Contact the external site for answers to questions regarding its content. Images displayed in this booklet are the property of their respective photographers or creators. IMAGES MAY NOT BE REDISPLAYED WITHOUT PERMISSION OF THE PHOTOGRAPHER OR CREATOR. Contact the creator if you wish to obtain a reproduction of an image or if you wish to obtain permission to redisplay an image on another website. For information on the creators, please contact the project coordinator via the address above.

PUBLICATION & INTELLECTUAL PROPERTY

This curriculum was developed within the FABE project - Fantasy and Artificial Intelligence for Business Enhancement - co-funded by the Erasmus+ Programme of the European Union, under Small-scale partnerships in vocational education and training (KA210-VET), Form ID KA210-VET-9F46079C.

This work is published under a Creative Commons Attribution–NonCommercial–ShareAlike 4.0 International License (CC BY-NC-SA 4.0). You are free to share and adapt this material for non-commercial purposes, provided that appropriate credit is given to the authors and to the FABE project, and that any derivative works are distributed under the same license. Commercial use is not permitted without prior written authorisation from the authors and the project consortium.

The contribution of Massimo Gerardo Carrese (www.fantasiologo.com) to this curriculum consists of original texts written within the FABE project and released under the license above. The theories referred to in these texts are developed in full in his book **Il grande libro della fantasia** (Il Saggiatore, 2023, ISBN 9788842832997), whose rights remain with the author and the publisher and are not covered by this license.

The European Commission's support for the production of this publication does not constitute an endorsement of the contents, which reflect the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

CONTENTS

About the FABE Curriculum.....	8
Target group.....	8
Purpose and project context.....	8
FABE educational approach.....	9
Link to Erasmus+ priorities.....	9
How to use this curriculum.....	10
Overview of the Chapters.....	11
HUMAN CAPABILITIES IN AI-ENHANCED WORK.....	13
Introducing the Topic.....	13
Exploring Ideas and Examples.....	14
Human capabilities in AI-enhanced work.....	14
Begin with the need, not the tool.....	15
Efficiency and quality.....	16
Strengths, limitations, and human responsibility.....	16
Practice, Reflection, and Discussion.....	17
FANTASIOLOGY FOUNDATIONS: IMAGINATION, IMAGERY, AND CREATIVITY.....	24
Introducing the Topic.....	24
Exploring Ideas and Examples.....	26
What is Fantasiology?.....	26
Creativity as an observable and trainable process.....	27
Cross-disciplinary perspectives.....	28
The logic of “Also” and the transformation of information.....	30
Human imagination and AI-generated products.....	31
Practice, Reflection, and Discussion.....	33
FOCUS ON CREATIVITY: EXPLORATION AND EXPLOITATION.....	51
Introducing the Topic.....	51
Exploring Ideas and Examples.....	52
Exploration: opening the field.....	52
Exploitation: developing what matters.....	53

Managing the polarity.....	54
AI across exploration and exploitation.....	55
Practice, Reflection, and Discussion.....	56
LATENT SPACES IN AI.....	63
Introducing the Topic.....	63
Exploring Ideas and Examples.....	65
Understanding latent space as a computational representation.....	65
Semantic proximity.....	66
Prompting as navigation through latent space.....	68
Seeds, variation, and controlled comparison.....	70
Unexpected outputs as positional feedback.....	70
From output consumption to active navigation.....	72
Latent space as geography: archetypes.....	72
Overlap zones.....	73
Unexplored territories.....	74
Practice, Reflection, and Discussion.....	74
INTERCONNECTIONS AMONG IMAGINATION AND LATENT SPACES..	81
Introducing the Topic.....	81
Exploring Ideas and Examples.....	84
Two different spaces of possibility.....	84
From human imagination to latent-space exploration.....	86
From AI output to human reinterpretation.....	87
Fantasiological functions as a bridge.....	88
Designing a hybrid creative workflow.....	90
Practice, Reflection, and Discussion.....	92
BUILDING APPLICATIONS WITH AI.....	97
Introducing the Topic.....	97
Exploring Ideas and Examples.....	99
From idea to tool: building apps with AI.....	99
Building a tool step by step.....	100
Testing, refining, and taking responsibility.....	102
Practice, Reflection, and Discussion.....	104

ETHICS AND ECOLOGY OF AI.....	109
Introducing the Topic.....	109
Exploring Ideas and Examples.....	111
Assistance, authority, and human accountability.....	111
Bias, knowledge, and the concentration of power.....	112
Privacy, authorship, and transparency.....	114
Work, education, creativity, and the value of difficulty.....	115
The ecology of an apparently immaterial technology.....	116
Practice, Reflection, and Discussion.....	118
Final reflection: Keeping the question open.....	124
GLOSSARY.....	125

ABOUT THE FABE CURRICULUM

The FABE Curriculum - Fantasy/imagination and Artificial Intelligence for Business Enhancement supports trainers, professionals, and VET learners in developing creative, digital, and critical skills for working with Artificial Intelligence.

FABE approaches AI as a co-creative tool that can support idea generation, learning, communication, problem-solving, and professional innovation. At the same time, it emphasises that human imagination, judgement, context, and responsibility remain essential.

Within the curriculum, fantasy/imagination is understood as the capacity to generate possibilities, create connections, and transform ideas into meaningful outcomes. The modules combine conceptual knowledge with practical methods that can be adapted to different sectors, learning environments, and professional needs.

The curriculum is organised into seven connected modules, moving from the foundations of FABE and human creativity to latent spaces, human-AI collaboration, writing applications, and responsible AI use. The separate FABE Exercises document provides practical activities that can be used alongside the relevant lessons.

| Target group

The curriculum is primarily designed for trainers and educators in vocational education and adult learning, professionals working in creative, cultural, educational, digital, entrepreneurial, and innovative fields, and VET students interested in strengthening their digital, creative, and transversal skills.

| Purpose and project context

Artificial Intelligence is transforming professional and educational environments, creating new opportunities while also increasing the need for digital readiness, adaptability, and creative problem-solving. FABE was developed in response to the growing gap between the

potential of AI technologies and the confidence and skills professionals, trainers, and students need to integrate them effectively into their work.

The FABE approach combines fantasy/imagination, understood as a creative and imaginative cognitive capacity, with Artificial Intelligence as a co-creative tool, supporting idea generation, learning processes, decision-making, and innovation. Rather than replacing human capabilities, AI is approached as a system that can strengthen fantasy/imagination, creativity, and professional development when used critically and intentionally.

This curriculum aims to help participants develop digital and creative competencies while understanding how imagination and AI can work together to improve professional practice and prepare for evolving labour market demands.

| FABE educational approach

Fantasy/Imagination and Artificial Intelligence for Business Enhancement (FABE) is an educational approach that combines fantasy/imagination, imagery, creativity, and Artificial Intelligence to strengthen professional skills in evolving work environments. The project responds to the growing need for digital readiness and AI literacy while emphasising that technology should support, not replace, human creative capacities. FABE promotes AI as a co-creative system that can enhance problem-solving, innovation, and professional development through structured methodologies that integrate fantasy/imagination, critical thinking, and digital tools.

| Link to Erasmus+ priorities

FABE was developed within the Erasmus+ framework to respond to current educational and labour market challenges linked to digital transformation, skills development, and vocational education. Understanding how the project aligns with Erasmus+ priorities helps participants recognise the broader relevance of FABE and its contribution to preparing professionals and learners for evolving technological and professional environments.

- **Digital readiness**

FABE strengthens digital readiness by helping participants understand, evaluate, and apply AI and other digital technologies in professional contexts. The project promotes practical AI knowledge, digital resilience, critical thinking, responsible technology use, and the meaningful integration of digital tools into professional and educational practice.

- **Reskilling and upskilling**

FABE supports upskilling and reskilling by equipping professionals, trainers, and learners with creative and AI-related competencies. Participants develop both technical and transferable skills that enhance adaptability, continuous learning, and creative problem-solving in evolving labour markets.

- **VET adaptation to labour market needs**

FABE helps vocational education and training respond to technological change by introducing AI-enhanced methodologies and innovative teaching approaches. By combining creativity, imagination, and AI literacy, the project better prepares learners and professionals for changing labour market demands.

| How to use this curriculum

Each chapter follows the same learning rhythm: an introduction establishes the topic, a deeper exploration develops its ideas and examples, and a final section gathers exercises, reflection, discussion, case studies, and assessment. Facilitators may adapt the activities to their learners, sector, available tools, and time.

AI systems, interfaces, and access conditions change. Examples of tools should therefore be treated as replaceable illustrations rather than endorsements. Participants should work with tools that are accessible and suitable for the task while considering privacy, reliability, environmental impact, and organisational policy.

| Overview of the Chapters

Chapter 1 - Introduction to FABE

Introduces the FABE philosophy, AI as augmentation rather than replacement, tool selection, professional needs, and the difference between efficiency and quality.

Chapter 2 - Imagination, Imagery, and Creativity

Explains the foundations of Fantasiology and distinguishes the processes of generating possibilities, forming representations, and organising ideas into meaningful outcomes.

Chapter 3 - Focus on Creativity

Explores creativity as movement between exploration and exploitation, including divergent thinking, connection-building, refinement, and implementation.

Chapter 4 - Latent Spaces in AI

Introduces latent space, semantic proximity, iterative prompting, the idea of using AI outputs as feedback during creative navigation, and the geography of archetypes, overlaps, and unexplored territories.

Chapter 5 - Interconnections Among Imagination and Latent Spaces

Compares human imaginative processes with AI generation and develops a collaborative workflow that combines human meaning-making with computational exploration.

Chapter 6 - Writing Applications in AI

Apply FABE principles to create useful tools and small apps to streamline work routines or help define work goals: from design and creation to adaptation and correction of features.

Chapter 7 - Ethics and Ecology of AI

Examines responsible AI use, including bias, accountability, privacy, transparency, environmental impact, and sustainable digital practice.

A note from the authors

Artificial Intelligence is evolving very quickly. By the time you read this curriculum, some of the tools, examples, or information may already be outdated. They reflect the state of AI at the time of writing.

This curriculum is not meant to give fixed answers or step-by-step instructions. Instead, it offers ideas, examples, and questions to help you think critically and make informed decisions. While technologies will change, the bigger questions about creativity, responsibility, ethics, and the impact of AI will remain relevant.

If you come across information that has already changed, we encourage you to update it, question it, and continue the conversation because that is the purpose of this curriculum.

MODULE 1**HUMAN CAPABILITIES IN
AI-ENHANCED WORK**

Working creatively and critically with AI

| Introducing the Topic

Artificial Intelligence is becoming part of many professional and educational activities. Its presence changes not only how tasks are completed, but also which forms of human attention, imagination, interpretation, and judgement are required. Understanding this relationship is a necessary foundation for the practical and creative work developed in the following chapters.

In this chapter, AI is approached as a collaborative technology that may support research, communication, idea development, comparison, and provisional production. Collaboration does not mean that the system and the person contribute in the same way. People provide purposes, lived and professional context, ethical judgement, interpretation, and responsibility for consequences.

The central question is therefore not whether AI should replace a complete task, but how a workflow can be designed so that technology supports valuable human capabilities. This requires beginning with a real need, distinguishing efficiency from quality, recognising limitations, and deciding where human involvement is essential.

Learning objectives

- Understand FABE philosophy
- Understand AI as augmentation rather than replacement
- Reflect on professional needs and expectations
- Identify suitable AI tools based on purpose
- Understand quality vs efficiency in AI work
- Build awareness of AI limitations and strengths

By the end of this chapter, participants will understand the foundations of FABE and recognise AI not simply as a productivity tool but as a technology that can support creative thinking, professional innovation, and lifelong learning when combined with human imagination and critical judgement.

| Exploring Ideas and Examples

| Human capabilities in AI-enhanced work

Fantasy/imagination as cognitive capability

Fantasy/imagination is understood in FABE not as escapism or entertainment but as a cognitive ability that supports idea generation, problem-solving, and the exploration of possibilities.

Within FABE, fantasy/imagination becomes a technical and professional skill. It enables professionals and learners to think beyond existing patterns, approach challenges from new perspectives, and develop innovative solutions in collaboration with digital technologies.

AI as collaborative technology

Artificial Intelligence is introduced in FABE as a collaborative technology rather than a replacement for human thinking. AI systems can support research, idea development, automation, communication, and decision-making, but human judgement, creativity, and critical evaluation remain essential.

FABE promotes the understanding of AI as a co-creative tool that expands human capabilities. Professionals are encouraged to use AI intentionally, selecting technologies according to their needs while maintaining responsibility for interpretation, quality, and ethical decision-making.

Creativity in future labour markets

Rapid technological transformation is changing professional environments and creating demand for new combinations of digital and creative competences. FABE responds to the growing need for professionals who can combine technical understanding with imagination, adaptability, and innovative thinking.

Creativity is increasingly becoming a strategic professional competence. The ability to generate ideas, solve complex problems, adapt to emerging technologies, and work effectively alongside AI systems contributes to long-term employability and professional competitiveness.

AI-enhanced workflows

AI-enhanced workflows integrate Artificial Intelligence into professional activities to support efficiency, learning, research, communication, and creative production. FABE encourages participants to understand where AI can strengthen existing processes while recognising where human expertise remains essential.

This training material emphasises that AI should not only be used to save time but also to improve the quality of outcomes. Professionals benefit most when technology supports thoughtful decision-making, critical reflection, and creative development rather than simple automation.

| Begin with the need, not the tool

AI tools differ in their capabilities, access conditions, interfaces, data practices, and costs. Moving among several systems may not be affordable or useful for every participant. Tool selection should therefore begin with a clear account of the task, the intended audience, the sensitivity of the information, the required level of reliability, and the resources available.

Before selecting a tool, ask:

- What am I trying to understand, create, improve, compare, or automate?
- What knowledge and judgement must a person contribute?
- What information can safely and lawfully be entered into the system?
- How accurate, current, explainable, or original must the result be?
- How will the output be checked, revised, and used?
- Is the tool accessible to the intended participants in terms of cost, language, disability access, and technical requirements?

A tool is suitable only in relation to a purpose and context. Familiarity, popularity, or the number of available features is not sufficient evidence that it is appropriate for a particular task.

| Efficiency and quality

Using AI may reduce the time required for some tasks, but speed and quality are different objectives. Producing a first draft quickly may create more time for analysis and revision; it does not remove those stages. In unfamiliar, sensitive, or complex work, checking and correcting an AI-assisted output may require substantial effort.

The relevant question is therefore not only “Did this save time?” but also “What kind of work did it change?” AI may shift effort from initial production to framing, comparison, verification, editing, interpretation, and decision-making. A useful workflow makes this redistributed human work visible.

| Strengths, limitations, and human responsibility

AI systems can help generate alternatives, reorganise information, adapt language, identify patterns, or produce provisional drafts. They may also generate **incorrect information**, reproduce **stereotypes**, **overlook context**, present **generic material confidently**, or encourage users to accept a fluent result **without sufficient examination**.

Human responsibility includes deciding whether AI should be used, defining the purpose, protecting sensitive information, checking important claims, recognising missing perspectives, adapting material to the situation, and taking responsibility for the final outcome. Later chapters develop the creative, practical, and ethical dimensions of these responsibilities.

| Practice, Reflection, and Discussion

EXERCISE

1: Professional needs and expectations

Purpose

To identify where AI might support a real professional or educational need before considering a particular tool.

Instructions

1. Choose one recurring task from your work or learning context.
2. Describe the current process, including its purpose, audience, inputs, decisions, difficulties, and desired quality.
3. Mark which parts require human experience, relationship-building, accountability, contextual judgement, or confidential information.
4. Identify one part that AI might support and one part that should remain primarily human.
5. State what improvement would count as success. Examples include clearer communication, more alternatives, better accessibility, stronger analysis, or reduced repetitive work.

Expected outcome

A one-page task profile that distinguishes the professional need from a possible technological response.

Reflection questions

- Is AI necessary for this task, or would another method be more appropriate?
- What could be gained, and what could be weakened, by introducing AI?
- Who is affected by the resulting decision or output?
- What evidence would show that the workflow improved?

EXERCISE**2: Choosing a suitable tool****Purpose**

To compare possible tools according to a defined need rather than selecting one because it is familiar or fashionable.

Instructions

1. Use the task profile from Exercise 1 and identify two or three possible tools or non-AI alternatives.
2. Agree on comparison criteria before testing. Include purpose, output type, reliability, privacy, accessibility, cost, language support, and the amount of human review required.
3. Test only with non-sensitive material. Record what each option does well, what it does poorly, and what remains unknown.
4. Select an option-or decide not to use AI-and justify the decision in relation to the original need.

Expected outcome

A short comparison and a reasoned selection that makes constraints and uncertainties visible.

Reflection questions

- Which criterion changed the decision most?
- Did the most capable option also appear to be the most appropriate?
- Which access or cost assumptions might exclude participants?
- What would need to change before the selected option could be used responsibly?

EXERCISE**3: Efficiency versus quality****Purpose**

To examine how AI changes the distribution of work rather than assuming that faster production automatically creates a better result.

Instructions

1. Choose a manageable task such as drafting a short announcement, reorganising notes, or adapting an explanation for a different audience.
2. Complete the task once without AI and once with AI support, using comparable objectives.
3. Record time spent on understanding the task, producing material, checking claims, revising language, correcting problems, and making the final decision.

4. Compare both outcomes using agreed quality criteria rather than personal preference alone.
5. Decide whether AI improved speed, quality, both, or neither and explain the conditions that influenced the result.

Expected outcome

A process comparison showing where time and judgement were used in both workflows.

Reflection questions

- Which stages became faster, and which required additional work?
- Did fluency make weaknesses more difficult to notice?
- What type of quality could not be measured by time?
- Would the conclusion change for a more sensitive or specialised task?

EXERCISE

4: Augmentation or replacement?

Purpose

To distinguish collaboration that strengthens human capability from automation that removes necessary judgement, learning, or responsibility.

Instructions

1. Select a scenario: preparing learner feedback, developing a project idea, communicating with clients, analysing survey responses, or creating educational material.
2. Divide the workflow into decisions and actions.

3. Assign each element provisionally to the person, the AI system, or a collaborative stage.
4. Review the allocation for risks to learning, professional expertise, privacy, inclusion, accountability, and quality.
5. Redesign the workflow so that the role of human judgement is explicit at the beginning, during review, and in the final decision.

Expected outcome

A workflow map that explains what is delegated, what is retained, and why.

Reflection questions

- Which human capability is the workflow intended to strengthen?
- Which delegation could reduce learning or professional understanding?
- Where could responsibility become unclear?
- What should happen when the human and AI-supported conclusions differ?

CASE STUDY

Facilitators may choose one case according to the participants' professional context.

Option 1: The trainer's lesson plan

A trainer considers using AI to prepare a complete lesson plan. The first result is fast and well organised but does not reflect the learners' prior knowledge, accessibility needs, or local context. Participants decide which parts should be retained, revised, investigated, or created by the trainer.

- Identify the professional need, the possible role of AI, the required human contribution, and the criteria for a responsible decision.

Option 2: The small organisation's communication task

A small organisation needs multilingual public information but cannot subscribe to several AI services. Participants compare affordable AI and non-AI options while considering language quality, privacy, staff time, accessibility, and human review.

- Identify the professional need, the possible role of AI, the required human contribution, and the criteria for a responsible decision.

Option 3: The creative professional's first draft

A creative professional uses AI to generate many initial concepts. The volume is useful, but the suggestions feel generic and begin to shape the professional's choices. Participants redesign the workflow so that AI expands possibilities without replacing personal direction.

- Identify the professional need, the possible role of AI, the required human contribution, and the criteria for a responsible decision.

Option 4: The apparently efficient report

A team uses AI to produce a report quickly. Verification reveals invented details and missing contextual knowledge, and revision takes longer than expected. Participants examine whether the workflow failed because of the tool, the task, the process, or unrealistic expectations about efficiency.

- Identify the professional need, the possible role of AI, the required human contribution, and the criteria for a responsible decision.

ASSESSMENT TOOL

Assessment should focus on whether participants can make a reasoned decision about AI use in context. A participant may submit one completed exercise together with a short explanation of the chosen workflow.

- Explains the central FABE principle of augmentation rather than replacement.
- Defines a professional or educational need before choosing a tool.
- Compares options using purpose, quality, access, privacy, reliability, and human-review criteria.
- Distinguishes efficiency from quality and makes hidden review work visible.
- Identifies AI strengths and limitations without anthropomorphising the system.
- Keeps human judgement, accountability, and final responsibility explicit.

MODULE 2**FANTASIOLOGY FOUNDATIONS:
IMAGINATION, IMAGERY, AND
CREATIVITY**

| Introducing the Topic

This chapter introduces participants to the foundations of Fantasiology, the framework developed by Massimo Gerardo Carrese. Fantasiology approaches imagination/fantasy, imagery, and creativity not as mysterious talents that some individuals possess and others do not, but as human faculties that can be observed, studied, developed, and intentionally trained.

The module builds a conceptual foundation for later work with Artificial Intelligence and latent spaces by helping participants first understand human imaginative processes. Before exploring how AI systems generate associations and support creativity, participants are invited to reflect on how imagination operates within their own thinking, learning, and professional practice.

Fantasiology encourages participants to recognise imagination/fantasy as an active cognitive process that influences communication, problem-solving, innovation, and creative production across professional and educational contexts. Through this perspective, creativity becomes a practice that can be strengthened through awareness, experimentation, and structured methods rather than being understood as an innate gift.

Learning objectives

- Distinguish imagination/fantasy, imagery, and creativity

Participants learn the differences between generating possibilities, constructing internal representations, and transforming ideas into meaningful outcomes.

- Understand Fantasiology principles

Participants become familiar with Fantasiology as an interdisciplinary framework for studying imagination, imagery, and creativity.

- Understand imagination as a trainable competency

Participants recognise imagination not as an abstract or fixed characteristic but as a cognitive capacity that can be strengthened through practice and structured methods.

- Reflect on personal creative processes

Participants analyse how they think, generate ideas, solve problems, and engage with creative work in professional and educational contexts.

- Apply Fantasiology methods to objects, words, narratives, patterns, and visual material.
- Distinguish human imaginative experience from AI-generated products and processes.

Methodological premise

This Module proposes a methodological framework for exploring and organising creative processes, with particular attention to the interaction between human cognition and artificial intelligence. Its originality does not lie in introducing new cognitive functions. Functions such as inversion, combination, multiplication, and divergent questioning, such as “What else could it be?”, were already present in

intellectual, artistic, rhetorical, and cultural practices long before creativity emerged as a distinct field of systematic study in the twentieth century. The methodological contribution lies instead in the systematic recombination of these operations within the framework of fantasiology, and in their application to the comparison of human imagination, imagery, and creativity with the generative, associative, and probabilistic processes observed in AI systems.

Originality therefore concerns the overall theoretical, practical, and methodological architecture: the conceptual framework such as the “Also” conjunction, specific tools such as the “Hupp” application, and situated exercises such as “The Chair Problem” or “Joan of Arc’s Theory”.

The exercises are not merely devices for generating new ideas, but instruments for observation and comparison, enabling the investigation of differences between embodied experience, intentionality, and human symbolic reasoning, on the one hand, and processes observed in generative AI models, on the other.

Creativity is therefore not reduced to divergence: generated possibilities undergo selection, feasibility assessment, organisation, and design, potentially leading to concrete configurations.

The term “functions” (or “fantasiological functions”) is used here, and more specifically in Module 5, to designate operations such as combination, separation, and multiplication. It does not imply that these functions were invented by the proposed methodology; rather, the term identifies their specific role within the framework presented here.

| Exploring Ideas and Examples

| What is Fantasiology?

Fantasiology is an interdisciplinary field of study that investigates imagination, imagery, and creativity as observable and trainable human faculties. Rather than treating creativity as spontaneous inspiration or

exceptional talent, Fantasiology explores the mechanisms and processes through which human beings imagine, associate ideas, construct meaning, and create new possibilities.

Fantasiology draws knowledge from multiple disciplines to understand imagination more comprehensively and to demonstrate that creative thinking can be strengthened through conscious practice and reflection.

Fantasiology studies three connected faculties:

- **Imagination/fantasy** - how people generate possibilities and alternatives
- **Imagery** - how mental representations and symbolic constructions are formed
- **Creativity** - how ideas are organised, transformed, and developed into meaningful outcomes

The distinctions help participants observe a process that may otherwise appear instantaneous. Imagination/fantasy opens possibilities; imagery gives those possibilities perceptible or symbolic form; creativity organises and transforms them in relation to meaning, purpose, or use. In practice, the three faculties may overlap and influence one another.

| Creativity as an observable and trainable process

Creativity is often perceived as spontaneous inspiration or as a talent that only some individuals naturally possess. This discussion activity invites participants to explore an alternative perspective by examining creativity as a process that develops through observation, experimentation, reflection, association, and practice.

Participants reflect on their own experiences to identify how ideas emerge and how creative outcomes often result from identifiable cognitive processes rather than sudden moments of inspiration.

This perspective does not reduce creativity to a formula or guarantee that the same method will produce the same result for everyone. It creates a language for noticing how observation, association,

experimentation, memory, revision, and context contribute to creative work.

| Cross-disciplinary perspectives

Fantasiology approaches imagination/fantasy, imagery, and creativity as human faculties that operate across multiple disciplines and forms of knowledge. Understanding creativity from different perspectives helps participants recognise that imaginative processes influence not only artistic work but also learning, communication, reasoning, innovation, and problem-solving.

Philosophy

Philosophy explores how human beings think, perceive reality, construct meaning, and imagine possibilities. Philosophical reflection helps examine questions such as how ideas emerge, how knowledge develops, and how imagination influences understanding.

Imagining future possibilities before making decisions is not only a creative process but also a reflective and philosophical one.

Literature

Literature demonstrates how imagination shapes narratives, characters, symbolic meanings, and alternative realities. Stories often combine imagination, imagery, and creativity to communicate ideas and emotions.

A writer constructing a fictional world uses imagination to generate possibilities, imagery to build representations, and creativity to organise them into a meaningful narrative.

Design

Design applies creativity directly to problem-solving and innovation. Designers combine imagination with practical constraints to develop products, services, systems, and experiences.

Creating a new user experience for a digital application requires imagining possibilities, testing alternatives, and refining solutions through creative processes.

Mathematics

Mathematics also involves imagination and creativity. Mathematical thinking often requires recognising patterns, exploring abstract relationships, and developing conceptual models.

Solving a complex problem may involve mentally visualising structures, testing possibilities, and finding unexpected connections between concepts.

Language

Language shapes how people organise thoughts, communicate ideas, and construct meaning. Words, metaphors, symbols, and associations influence imaginative and creative processes.

Changing the language used to describe a challenge may generate different interpretations and open new possibilities for solutions.

Visual culture

Visual culture examines how images, symbols, media, and visual representations influence understanding and communication. People constantly interpret visual information and create meaning through visual experiences.

Advertising, digital media, diagrams, or visual storytelling often rely on imagery and symbolic construction to communicate ideas effectively.

These perspectives are not separate versions of Fantasiology. Together, they show that imagination, imagery, and creativity operate across reflection, storytelling, problem-solving, pattern recognition, language, design, and interpretation.

| The logic of “Also” and the transformation of information

A practical fantasiologic example: Redesigning a Coffee Mug The "What else could it be?" approach and the logic of "Also" applied to an everyday object.

For Massimo Gerardo Carrese, redesigning an everyday object does not mean, in this specific context, erasing its primary identity, but expanding its possibilities through the concept of Also: it is not just a mug, but also a multisensory, modular, and interactive desk tool capable of integrating function, perception, and meaning.

HUPP is an experimental app designed to compose and decompose information related to an object, a word, a sound, or any other informational element. Through this process, the application demonstrates how a single piece of information can be fragmented, analyzed, transformed, and recombined in order to stimulate the production of new ideas.

The purpose of HUPP is not only to identify possible uses of an object or concept, but also to explore the technical and imaginative aspects connected to Fantasiology. The application encourages the generation of specific project proposals (from initial sketches to possible realizations) by providing structured indications and evaluation parameters such as:

- originality;
- usefulness;
- method;
- technique of realization;
- materials;
- feasibility;
- development process;

HUPP shows that ideas can exist in different dimensions: they can remain active only in the mind, they can develop within the dimension of making-believe and play, or they can evolve into concrete, realizable

projects. It works with Carrese's methodology on imagination, imagery and creativity.

This approach does not simply encourage creativity. It provides a structured method for transforming an initial intuition into an educational project that can demonstrate innovation, impact, sustainability, transferability, and feasibility: the elements generally required by European funding programmes.

The practical value lies in the ability to move from a simple idea to a structured design process.

The detailed application of this method belongs in practice. Part 3 therefore invites participants to decompose and recompose an object rather than repeating a long model prompt in the explanatory text. Prompt construction itself is developed in Chapter 4.

| Human imagination and AI-generated products

However, this should not be understood as a human-like concept representation. The model's internal structures encode patterns of association and correlation rather than necessarily constituting representations grounded in perception, bodily experience, or subjective understanding.

Human imagination often involves internally generated simulations that can guide perception, reasoning, and action. In many cases, a person can mentally explore a possible object before producing it externally. However, human imagination is not always a fully formed image that precedes action: creative processes frequently involve cycles of generation, evaluation, and revision.

In current AI systems, there is no evidence of an equivalent phenomenological stage: no indication that the model internally "sees" or experiences an image before producing the output. The generation process consists of transformations through learned computational spaces that ultimately result in pixels, rather than the externalization of a previously experienced mental image.

The human imagination of a seawater chair is constrained by embodied knowledge. A person implicitly understands that water has weight, lacks structural rigidity, has a particular temperature, behaves under gravity, and interacts with the human body in specific ways.

The model can reproduce visual cues of physical plausibility because it has learned regularities from data, not because it has interacted with water, gravity, balance, or material resistance. Its understanding of physical properties is therefore derived from patterns in information rather than from embodied experience.

Human imagination involves a first-person dimension: there is a subject for whom the imagined content appears. Even when fleeting or unconscious, mental imagery belongs to an ongoing stream of experience.

Current AI models provide no evidence of subjective experience, awareness, or a perspective from which the generated image is present. The image exists as an output of a computational process rather than as the content of an experience for a subject.

Human creative imagination is typically guided by internal goals, expectations, and judgments about what should be produced. A person can compare an imagined outcome with an intended one and revise the process accordingly.

A model can be optimized toward producing certain outputs and can incorporate external evaluation mechanisms, but there is no evidence that it possesses an internally represented desired state or a subjective sense of success or failure comparable to human intentionality.

The result can be functionally indistinguishable from an object produced through human imagination: it depicts a novel possibility that did not previously exist. However, similarity of the final image does not imply identity of the underlying process.

The distinction between product and process allows participants to use AI-generated material without describing the system as if it possessed bodily experience, subjective imagery, or personal intention. The

exercises below use AI only where comparison can clarify this distinction or support reflection on a human creative path.

| Practice, Reflection, and Discussion

EXERCISE

1: Creativity is not magic

Purpose

Creativity is often perceived as spontaneous inspiration or as a talent that only some individuals naturally possess. This discussion activity invites participants to explore an alternative perspective by examining creativity as a process that develops through observation, experimentation, reflection, association, and practice.

Instructions

1. Individually, recall a situation in which you created something, solved a problem, adapted an idea, or found an unexpected response to a challenge.
2. Write down how the idea emerged, what influenced it, and whether experimentation, mistakes, testing, or revision changed it.
3. In small groups, compare the processes without ranking the outcomes. Identify recurring influences and meaningful differences.
4. As a group, create a short map of the factors that supported creativity.

Expected outcome

A process map showing how ideas developed through identifiable experiences, actions, and decisions.

Reflection questions

- Did the solution appear immediately or develop step by step?
- Which experiences, observations, associations, or influences contributed to the outcome?
- Did experimentation, revision, mistakes, or testing influence the result?
- What could be practised or changed to support creativity in future work?

EXERCISE**2: The Chair Problem****Purpose**

To compare an imaginative product with the human and computational processes that may produce it.

Exercise example: "The Chair**Instructions**

1. Observe the image of a chair made of seawater and describe only what is visible before interpreting how it was produced.
2. Identify which structural characteristics still make the object recognisable as a chair and which characteristics communicate seawater.
3. Discuss how a person might imagine this object through embodied knowledge, memory, purpose, and internal evaluation.
4. Compare this with an AI model that combines learned visual and linguistic regularities without evidence of subjective experience.
5. Test one additional impossible-material object if an AI image tool is available. Suggested starting points are listed below; the purpose is comparison, not prompt optimisation.

- "A spiral staircase made of smoke rising upward, keeping its shape while dissolving at the edges."
- "A violin built entirely of interwoven tree roots, still playable in form."
- "A wall clock whose numbers and hands are made of slowly melting ice, with water dripping down to form the missing hours."
- "A wall bookshelf whose shelves and books are carved from sand, with titles visible on the spines."
- "A pair of running shoes made of blown glass, suspended mid-air as if running, with cracks following the motion."

Expected outcome

A short comparison that distinguishes similarity in the visible product from differences in embodiment, experience, intention, and generation.

Reflection questions

- What makes the impossible object visually coherent?
- Which physical qualities can the image imitate without the system experiencing them?
- What knowledge or intention would a human bring to the same task?
- Why does similarity of output not demonstrate similarity of process?

EXERCISE**3: The Impossible World Builder****Purpose**

Objective: To explore how imagination transforms reality into symbolic and meaningful fictional worlds. The person will use imagination to create possibilities, imagery to represent them, and creativity to organise them into a narrative.

Instructions

1. Choose an ordinary place: a classroom, street, house, park, or library.
2. Imagination/fantasy: change one or more rules of reality and create at least three impossible elements. You may alter time, communication, natural laws, visibility, or the function of an ordinary object.
3. Imagery: describe the world through colours, shapes, sounds, smells, sensations, creatures, objects, and emotions.
4. Creativity: organise the material into a 200-300 word narrative containing a protagonist, a symbolic element, a problem or discovery, and a hidden message or idea.

Expected outcome

A short narrative that makes the movement from possibility to representation and organised meaning visible.

Reflection questions

- How did imagination help you create new possibilities?
- Which invented element changed the meaning of the story most?
- How did your imaginary world express an emotion or idea?
- What does your symbolic object or creature represent?
- How does your fictional world invite the reader to see reality differently?

EXERCISE**4: Redesign an everyday object through “Also”****Purpose**

To decompose an object’s existing information and recompose it into imaginative, playful, and potentially realisable possibilities.

Instructions

1. Choose an everyday object and record its material, shape, primary purpose, context of use, mechanisms, and sensory qualities.
2. Ask “What else could it be?” Generate at least three added identities or functions without erasing the original identity.
3. Develop one abstract or impossible version, one making-believe or performative version, and one demonstrable or functional version.
4. Select one proposal and examine its users, usefulness, materials, method of realisation, feasibility, limitations, and ethical implications.
5. Present the transformation as: original information → decomposition → new connections → proposal → possible prototype.

Expected outcome

A one-page concept sheet or simple prototype showing how the logic of “Also” expanded the object’s identity.

Reflection questions

- Which component of the object generated the most possibilities?
- Which proposal is most imaginative, and which is most useful? Must these be the same?

- What was gained or lost when the idea became more feasible?
- Does the proposal add meaningful value or only additional features?

EXERCISE

5: Discovering a pattern in a situated date

Purpose

This exercise originates from the perspective of Massimo Carrese's fantasiology, an approach that invites us to observe reality not only for what immediately appears, but also for the hidden possibilities it may contain. The aim is not simply to solve a numerical problem, but to train the ability to imagine connections, identify invisible structures, and transform everyday information into an opportunity for creative exploration.

Instructions

1. Choose a date connected to the place or moment of the learning activity. Write it in a consistent numerical format.
2. Observe and divide the date in several ways. Look for repetition, reversal, symmetry, rhythm, or another numerical or visual relationship.
3. State the possible rule clearly enough that another group could test it.
4. Search for other dates that follow the same rule within an agreed period. Participants may calculate manually, use a spreadsheet, write simple code, or compare with an AI system after developing their own reasoning.

- Record both successful and unsuccessful rules and explain what each attempt revealed.

Expected outcome

A clearly stated pattern, a method for testing it, and examples that meet or challenge the proposed rule.

Reflection questions

- What changed when the data was treated as a visual structure rather than only as information?
- Which observations were mathematical, and which were imaginative?
- Could another rule explain the same pattern?
- What did a digital or AI comparison add after the human investigation?

EXERCISE

6: Word Association Chain

Purpose

This exercise explores the hidden connections between apparently distant elements. It is a game about associations, creativity, and the way the human mind constructs meanings.

Instructions

- Choose two words that appear distant from one another.
- Individually create a chain that connects the first word to the second through memories, sounds, images, emotions, stories, cultural associations, or personal experiences.
- Write every intermediate step. Do not remove an association merely because it cannot be universally proved.

4. Explain the chain to a partner, who identifies where the path becomes personal, unexpected, or difficult to follow.
5. Compare several chains created from the same pair of words. If AI is used, compare its likely or conventional chain with the situated human paths.

Expected outcome

An association chain accompanied by an explanation of the personal or situated logic connecting its steps.

Reflection questions

- Which link was logical or functional, and which was personal or symbolic?
- Where did memory, language, sound, emotion, or culture influence the path?
- Can a connection be meaningful without being universally demonstrable?
- How did the different participants' chains change the meaning of the same words?

EXERCISE

7: Situated visual interpretation - The Joan of Arc's Theory

Purpose

To explore how observation, imagery, association, and situated knowledge contribute to different interpretations of visual material.

LUDIC

Day 2 AFTERNOON

JOAN OF ARC

GETTING INSPIRED BY A LOCAL NEWSPAPER AND TRYING TO DRAW A PART OF JOAN OF ARC'S SWOON

LUDIC

NON SENSE?

NON SENSE?




FIRST STEP

MY FIRST STEP WAS TO FIND
A MEANINGFUL CONNECTION
BETWEEN THE CITY OF **SALETERNO**
AND JOAN OF ARC. I REMEMBERED
THAT ONE OF HER MOST IMPORTANT
BATTLES WAS FOUGHT IN **ORLEANS**.
I NOTICED THAT THE TWO WORDS
FORM A SURPRISING ANAGRAM.
A STROKE OF LUCK? YES
BUT LUCK ONLY MATTERS
IF YOU KNOW HOW TO SEIZE IT.

SECOND

T
E
R

MY SECOND STEP WAS TO UNDERSTAND
WHICH OTHER CITIES COULD EMERGE
BY ANAGRAMMING THE SAME LETTERS,
AND I REALIZED THAT BY RETROING
THE ACCENT FROM ORLEANS I WOULD
OBTAIN **ORLEANS**, THAT IS
NEW ORLEANS!


TSHIERP

I THEN LOOKED AT THESE
PLACES ON A MAP AND
CONNECTED THEM TO FORM A
GEOMETRIC SHAPE.

WHAT WOULD EMERGE FROM THAT?

IMAGINE!

ONCE THIS TRIANGLE HAS BEEN OBTAINED,
 I BEGAN TO ASK "WHAT ELSE COULD IT BE?"
 FROM THERE, KNOWING THE STORY OF JOAN
 OF ARC, I RECOGNIZE THAT THIS SHAPE
 CAN BECOME A PART OF HER SWORD.
 SO I TRIED TO
 SYSTEMATIZE MY IDEA!



THE CREATION OF A HANDLE




I CREATE THE HANDLE OF
 THE SWORD BY TRACING THE
 OUTLINE OF THE BUILDING WHERE
 THE ASSOCIATION'S MEETINGS
 ARE HELD

AND REFLECT
 I RECREATE THE FIGURE
 AND JOIN EACH TO
 FORM A HANDLE



FROM THE TRIANGLE
 I OBTAIN SPACE AND CURVES





WHAT IS THIS KING OF GAME GOOD FOR?

THE JOAN OF ARC'S THEORY

MEANING CAN EMERGE
FROM CONSTRAINTS

→ WORKING WITH FIXED
ELEMENTS (E.G. LETTERS...)
STIMULATE IMAGINATION,
IMAGERY, CREATIVITY

ABSTINCTION AS A
LEARNING TOOL

TO THINK BEYOND LITERAL CONTENT

PLAY GENERATES DEPTH
AND (LUCKY) CASUALTIES

ENCOURAGE CURIOSITY
INSTEAD OF COMPLACENCE


LOCAL NEWSPAPER AS NARRATIVE

LOCAL AS
AN ACTIVE MEANING-
MAKER



AN OPERATIONAL METAPHOR FOR EMERGENT + PROJECT DESIGN

- WHAT LOCAL RESOURCES DO YOU HAVE (TERRITORY COMPETENCES, HISTORY) THAT COULD BECOME THE "BLADE" OF A PROJECT?
- HOW COULD YOU FIND EUROPEAN PARTNERS NOT BY CONVENIENCE BUT BY AUTHENTIC CONNECTION?
- WHAT "SWORD" (FINAL PRODUCT?) DO YOU WANT TO BUILD TOGETHER WITH PARTNERS?



Instructions

1. Observe the visual material silently. List visible forms, repetitions, contrasts, directions, and relationships without assigning meaning.
2. Write one personal interpretation and identify the memories, cultural references, emotions, or contextual knowledge that influenced it.

3. Compare interpretations in small groups. Look for shared readings as well as differences that cannot be resolved into one answer.
4. Develop one interpretation into a short title, narrative, image sequence, or performative idea.

Expected outcome

A documented interpretation that separates observation from association and explains how the interpretation became a creative proposal.

Reflection questions

- Which parts of the interpretation came from the image, and which came from the observer?
- How did the place, context, or cultural reference shape meaning?
- Did another participant's interpretation reveal something you had not noticed?
- What makes an interpretation coherent without making it the only correct reading?

CASE STUDY

OPTIONS

Facilitators may select one of the following cases according to the group's interests and available time. The options are alternatives, not compulsory sequential activities.

Option 1: An ordinary object with several identities

A training group applies the logic of "Also" to an object used in its own learning environment. Participants compare an

impossible version, a playful version, and a functional version before deciding whether one should become a prototype.

- Discuss what the case reveals about imagination/fantasy, imagery, creativity, and the transformation of information.
- Identify which parts of the process could be supported by AI and which require situated human interpretation or judgement.

Option 2: A place transformed into a fantasiological world

Participants begin with a familiar educational or public place and change one rule of reality. They explore how sensory imagery and symbolic elements can transform the place into a narrative or learning experience.

- Discuss what the case reveals about imagination/fantasy, imagery, creativity, and the transformation of information.
- Identify which parts of the process could be supported by AI and which require situated human interpretation or judgement.

Option 3: A situated pattern hidden in everyday information

A date, timetable, street number, local sign, or repeated sound becomes the starting point for discovering patterns and generating a site-specific creative activity.

- Discuss what the case reveals about imagination/fantasy, imagery, creativity, and the transformation of information.
- Identify which parts of the process could be supported by AI and which require situated human interpretation or judgement.

Option 4: One pair of words, several personal paths

Different participants connect the same two distant words. Their chains reveal how memory, language, culture, emotion, and experience produce different but explainable forms of meaning.

- Discuss what the case reveals about imagination/fantasy, imagery, creativity, and the transformation of information.
- Identify which parts of the process could be supported by AI and which require situated human interpretation or judgement.

Option 5: Similar products, different processes

Participants compare an AI-generated impossible object with one imagined and developed by a person. They focus on embodiment, subjective experience, intention, evaluation, and responsibility rather than deciding which product looks more creative.

- Discuss what the case reveals about imagination/fantasy, imagery, creativity, and the transformation of information.
- Identify which parts of the process could be supported by AI and which require situated human interpretation or judgement.

ASSESSMENT TOOL

Assessment should make the creative process visible rather than reward only a polished final product. Participants may submit an exercise outcome together with a short process record.

- Distinguishes imagination/fantasy, imagery, and creativity in the chosen process.

- Explains how information was observed, decomposed, associated, transformed, or reorganised.
- Makes personal, cultural, disciplinary, or situated influences visible.
- Develops an idea coherently without treating feasibility as the only measure of value.
- Distinguishes AI-generated products from embodied, subjective, and intentional human processes.
- Reflects on choices, limitations, alternative interpretations, and possible next developments.

MODULE 3**FOCUS ON CREATIVITY:
EXPLORATION AND EXPLOITATION**

| Introducing the Topic**Balancing discovery and development**

Creativity is often described as the generation of new ideas, yet ideas become meaningful only when they are examined, developed, and connected to a purpose. Creative work therefore moves between two complementary modes. Exploration opens the field by producing alternatives, making unusual connections, and tolerating uncertainty. Exploitation narrows the field by selecting, testing, refining, and implementing what appears promising. Neither mode is inherently more creative than the other; creativity depends on the relationship between them.

This chapter presents exploration and exploitation as a polarity rather than a sequence with a fixed beginning and end. A person or team may explore, select an idea, discover a limitation during refinement, and return to exploration. The movement can be rapid or gradual, individual or collaborative. What matters is recognising the mode that the work currently needs instead of remaining indefinitely in expansion or closing down possibilities too early.

The chapter also treats creativity as connection-making. New possibilities often emerge by linking experiences, concepts, stories, materials, or needs that initially seem unrelated. These connections are then shaped through judgement and revision. AI tools can contribute alternatives, comparisons, and organisational support, but they do not determine which connections are meaningful, appropriate, feasible, or worth pursuing. Those decisions remain situated human responsibilities.

Learning objectives

- Explain exploration and exploitation as complementary modes within a creative process.
- Use divergent thinking and connection-making to expand a field of possibilities.
- Select, test, refine, and implement ideas in relation to purpose and constraints.
- Recognise cognitive habits that can keep a person or team in one mode for too long.
- Move deliberately between exploration and exploitation as new information emerges.
- Use AI as provisional support while retaining responsibility for evaluation and decisions.

| Exploring Ideas and Examples

| Exploration: opening the field

Exploration is the mode in which people suspend premature judgement and investigate what else might be possible. It includes experimentation, observation, divergent thinking, and the willingness to remain with uncertainty. The aim is not to produce novelty for its own sake, nor to avoid all constraints. It is to prevent familiar assumptions from defining the solution before the problem has been understood from several angles.

Experimentation makes uncertainty observable. Changing one ingredient in a recipe, testing an unfamiliar material, reframing a service from another user's perspective, or combining two functions in one object can reveal opportunities that planning alone would not uncover. A result that fails may still provide information about the problem, the material, the audience, or the assumptions guiding the attempt. Exploration is productive when participants can notice what each attempt teaches them.

Divergent thinking expands the number and variety of possibilities before selection. Asked how a community might reduce food waste, a group could consider shared storage, smaller packaging, neighbourhood exchanges, restaurant redistribution, preservation workshops, or changes to purchasing habits. The examples do not yet need to be equally feasible. Their value at this stage lies in opening different frames: technology, education, infrastructure, policy, social practice, and individual behaviour.

Creativity through connections

Creative connections may form between apparently distant observations. Two colleagues with contrasting working styles might be interpreted only as being in conflict, or their difference might suggest a complementary system in which speed and careful analysis correct one another. A broken table leg and a branch left from gardening might lead not only to a repair but to a support that also holds plants. In both examples, the creative act lies in changing the relationship among existing elements, not in producing something from nothing.

Connections are not automatically valuable because they are unusual. They gain significance through context: Whose need do they address? What assumption do they interrupt? What new function or interpretation do they make possible? Exploration generates candidates for meaning; it does not remove the need to evaluate them.

| Exploitation: developing what matters

Exploitation is the mode in which selected possibilities are made clearer, stronger, and more workable. It includes refinement, optimisation, testing, and implementation. The word does not imply unethical use in this context; it refers to developing the value of an idea that has already emerged. If exploration asks, "What could exist?", exploitation asks, "What deserves further work, and how could it function in practice?"

Refinement improves an idea without necessarily changing its central purpose. A workplace relaxation room that employees rarely use might

be reconsidered through interviews and observation. Flexible furniture, quiet work options, clearer access, or short guided activities could make the original idea more relevant. The creative contribution is not limited to the initial concept; it continues through attention to evidence, weaknesses, and use.

Optimisation examines how resources, effort, time, and quality relate. A reporting process that requires three meetings might be redesigned around a shared document and one decision-focused meeting. Implementation then translates the developed idea into responsibilities, timing, resources, a small pilot, and criteria for learning. These stages should not be treated as purely administrative: each practical constraint may require another creative choice.

Questions that support development

- What purpose should remain stable as the idea changes?
- Which evidence suggests that the idea is useful to the intended people?
- What is unclear, unnecessary, inaccessible, or difficult to maintain?
- Which constraint requires adaptation, and which constraint should be challenged?
- What is the smallest, responsible test that could produce useful feedback?

| Managing the polarity

Exploration and exploitation create a productive tension. Too much exploration can produce an expanding collection of ideas without commitment or learning. Too much exploitation can perfect a familiar solution while excluding alternatives that would address the problem more effectively. Polarity management means paying attention to both risks and deciding when the process needs expansion, selection, development, or a return to uncertainty.

The movement is rarely linear. A team may explore several approaches, select one, refine it, encounter a limitation, and reopen a specific part of

the problem. This is different from restarting without direction: the new exploration is informed by what the attempted implementation revealed. A useful rhythm might be represented as explore → compare → select → test → refine → reconsider. The arrows indicate learning, not an automatic procedure.

Individual and group habits influence this balance. Attachment to the first plausible idea, preference for consensus, fear of failure, endless searching for a more original answer, or pressure to demonstrate quick progress can all distort the process. Naming the current mode can help a group discuss its needs more precisely: Are we still trying to understand the possibility space, or are we avoiding a decision? Are we refining because evidence supports this direction, or because it is already familiar?

| AI across exploration and exploitation

AI can increase the speed and volume of both modes. During exploration, it may suggest alternative perspectives, variations, analogies, or categories that a participant can investigate. During exploitation, it may help organise material, compare versions, identify omissions, or draft a possible implementation sequence. These outputs are starting material rather than independent decisions, and speed can create the appearance of progress even when the underlying purpose remains unclear.

The participant therefore needs to ask what role the tool is playing. Does it broaden the field or merely reproduce dominant patterns? Does it help articulate an idea or make an untested idea appear finished? Whose experience, knowledge, or perspective is absent? What needs to be checked through observation, expertise, consultation, or evidence outside the system? Detailed prompting techniques are developed in Chapter 4; here the emphasis remains on deciding when AI support contributes to exploration and when it supports development.

Human judgement connects the two modes. People define the purpose, interpret relevance, recognise cultural and professional context, decide

what consequences matter, and take responsibility for what is implemented. An AI-generated possibility may be surprising without being valuable, and a well-structured plan may be coherent without being appropriate. The creative process remains incomplete until these distinctions are made.

| Practice, Reflection, and Discussion

EXERCISE

1: Expand the possibility space

Choose a concrete problem connected to the participants' work or community. In small groups, generate at least twenty possible responses. Begin with obvious ideas, then deliberately change perspective, scale, material, user, location, or time frame. If an AI tool is used, keep its suggestions distinguishable from the group's own contributions so that both sources can be examined.

- Which ideas repeat familiar assumptions about the problem?
- Which connection opened a genuinely different direction?
- Did the number of ideas increase their variety, or only create minor variations?
- What knowledge or lived experience changed the group's understanding?

EXERCISE

2: Select and strengthen

From the ideas generated in Exercise 1, choose three that represent different directions. Agree on criteria relevant to the

context—for example usefulness, inclusion, feasibility, originality, environmental impact, or capacity for learning—before comparing them. Select one idea for further development and record why the others are being paused rather than simply discarded.

Refine the selected idea by clarifying its purpose, intended users, required resources, likely obstacles, and smallest test. Identify one aspect that should remain stable and one that may need to change after feedback. The group should be able to explain how its criteria influenced the decision and what uncertainty remains.

EXERCISE

3: Map movement between the modes

- Choose a past or current creative project and draw its development as a timeline. Mark moments of exploration, selection, testing, refinement, implementation, and renewed exploration. Pay attention to transitions: What information caused the movement? Who made the decision? Was the shift deliberate, imposed by circumstances, or recognised only afterwards?
- Where did the process remain in one mode for too long?
- Was a promising possibility closed too early, or was a weak idea refined for too long?
- Which failure or constraint reopened exploration productively?
- What would help the team recognise the need to shift modes next time?

EXERCISE**4: Compare human and AI-supported contributions**

Give two groups the same practical challenge. One group begins without AI; the other may use an AI system to generate alternatives or organise possible next steps. After a limited period, exchange the work and continue developing it. Do not judge success by the number or polish of outputs. Compare the kinds of connections, assumptions, omissions, and decisions produced in each phase.

- Where did AI broaden the possibilities, and where did it narrow them toward familiar patterns?
- Which contributions required situated knowledge or interpretation from participants?
- Did a fluent output encourage premature selection?
- Which claims or assumptions would need independent verification before implementation?
- How did responsibility for the final choice remain visible?

CASE STUDY**OPTIONS**

The following cases offer different ways of exploring creativity, fantasiology, and movement between exploration and exploitation. A facilitator may select one case according to the group's interests and available time. The options are alternatives rather than a sequence of compulsory activities.

Option 1: The impossible object

Participants design an object that combines two apparently incompatible functions—for example, a clock that makes time

feel slower, a chair that encourages movement, or a window designed for a room without walls. The group begins by generating possibilities without evaluating them. It then selects one concept and develops its function, materials, intended user, and possible form. During refinement, participants observe whether practical questions strengthen the concept or remove the strangeness that made it interesting.

1. At what point should exploration give way to development?
2. Which element of the original idea should remain unchanged?
3. When does refinement strengthen an imaginative idea, and when does it make it conventional?
4. Could the object be valuable even if it cannot be produced literally?

Option 2: A fantasiological transformation

Choose an ordinary object such as a key, spoon, umbrella, staircase, or mirror. Explore what the object might become if one of its usual characteristics or purposes were transformed. A key could open conversations rather than doors, an umbrella could collect weather instead of protecting someone from it, or a staircase could lead through stages of a memory. Generate several transformations, select one, and develop it as an image, story, performance, object, or learning activity.

1. Which associations were obvious, and which created a new direction?
2. What connection gives the transformation coherence?
3. How much explanation does the idea require?
4. Does development reveal the idea's potential or reduce its ambiguity too quickly?

Option 3: The imaginary profession

Invent a profession that does not currently exist, such as a collector of forgotten sounds, a designer of coincidences, a translator of shadows, or a repairer of unfinished dreams. First explore what this person might do, who might need the work, which tools they would use, and what knowledge the profession would require. Then select one proposal and develop it into a coherent professional profile, including its purpose, activities, users, and possible social value.

1. Which ideas emerged through unexpected connections?
2. What made the invented profession feel meaningful rather than random?
3. Which practical details strengthened the concept?
4. At what point did feasibility begin to restrict imagination?
5. Could an apparently impossible profession reveal a real social need?

Option 4: Reimagining an everyday place

Choose a familiar place—a library, classroom, waiting room, bus stop, or workplace kitchen—and imagine what else it could represent before proposing practical improvements. A library might become a garden of unfinished questions, a waiting room a place for exchanging knowledge, or a bus stop a temporary stage. Generate several interpretations, select one, and consider how it could be translated into an actual experience without losing its imaginative meaning.

1. Which interpretations changed how the group understood the place?
2. Which proposal combined distant ideas most successfully?
3. What was lost when the concept was adapted to practical constraints?

4. How could the group preserve imaginative meaning during implementation?

Option 5: One object, many realities

Give every participant the same ordinary object and ask each person to place it within a different imaginative system. A cup might become an archaeological discovery, a musical instrument, a shelter, a measurement of silence, or evidence from an unknown civilisation. Compare the interpretations, identify the associations behind them, and select one reality to develop collaboratively-or explain why several interpretations should remain open.

1. How did the same object lead to different meanings?
2. Which interpretations depended on personal or cultural experience?
3. Did group discussion expand the possibilities or create pressure to agree?
4. Which idea should be developed, and according to what criteria?
5. Can several interpretations remain valuable without selecting a single winner?

ASSESSMENT TOOL

Assessment should focus on the participant's ability to explain and manage the creative process, not only on the novelty or polish of the final result. Participants may submit a short process record containing the initial challenge, examples of divergent connections, selection criteria, a developed proposal, feedback or test evidence, and a reflection on movement between the two modes.

- Distinguishes exploration from exploitation without treating them as isolated stages.
- Generates varied possibilities and explains the connections behind them.
- Selects and refines an idea using explicit, context-relevant criteria.
- Recognises when evidence or constraints justify returning to exploration.
- Evaluates AI contributions critically and identifies missing perspectives or verification needs.
- Explains final decisions and maintains human responsibility for implementation.

A strong submission makes the reasoning visible: it shows not only what was chosen, but how alternatives were opened, compared, developed, revised, or rejected.

MODULE 4

LATENT SPACES IN AI

| Introducing the Topic

Understanding latent space without anthropomorphizing AI

In previous modules, participants explored imagination, imagery, and creativity as human cognitive processes. They learned how people generate possibilities, build associations, and create meaning through experience, reflection, and connection-making.

This chapter introduces the concept of latent space, one of the mechanisms used by many generative Artificial Intelligence systems. While humans use imagination, memories, emotions, experiences, and cultural understanding to create new ideas, AI systems operate differently. They learn relationships, patterns, and structures from large amounts of data and represent these relationships within complex computational spaces known as latent spaces.

Understanding latent spaces helps explain why AI can generate images, texts, ideas, and associations that sometimes appear creative. At the same time, it helps participants recognise the important distinction between human imagination and computational representation.

Rather than viewing AI as a system that "thinks" or "imagines" like humans, this chapter encourages participants to understand AI as a system that navigates relationships between learned concepts and generates outputs based on patterns and probabilities.

This understanding creates an important bridge between creativity and Artificial Intelligence and prepares participants for more effective and responsible use of AI tools.

This chapter uses spatial language-coordinates, regions, movement, distance, and geography-as a teaching metaphor. These terms do not describe a place that users can see or enter, and the model does not consciously navigate. The metaphor is useful only when it remains

connected to the underlying idea: generated outputs reflect learned numerical relationships and probabilistic processes.

Learning objectives

- This chapter introduces participants to the concept of latent spaces and their role in Artificial Intelligence systems. Building on the previous exploration of imagination and creativity, participants learn how AI represents relationships between concepts, generates associations, and produces outputs. The objective is not to compare AI directly to human thinking, but to understand how computational systems process and organise information differently from human cognition.

Participants will:

- Understand latent spaces conceptually. Participants develop a foundational understanding of latent spaces as computational environments where relationships between concepts, patterns, and representations are organised. They learn how these structures enable AI systems to generate text, images, and other outputs.
- Distinguish imagination from computational representation. Participants explore the differences between human imagination and AI-generated processes. They learn that imagination involves personal experiences, emotions, context, and meaning-making, while AI systems operate through learned patterns and statistical relationships.
- Understand semantic proximity. Participants learn how concepts can be positioned closer together or further apart within a latent space. They explore how relationships between ideas influence AI outputs and how seemingly distant concepts can become connected through shared characteristics or associations.

- Use prompting strategically. Participants learn how prompts influence AI-generated outputs. They develop an understanding of how wording, context, constraints, and instructions can guide exploration within latent spaces and produce different results.
- Interpret variation and unexpected outputs. Participants learn to distinguish computational variation, ambiguity, generic patterns, factual errors, and potentially productive creative deviation.
- Navigate without surrendering judgement. Participants practise using outputs as provisional feedback while retaining responsibility for purpose, interpretation, verification, and final selection.

| Exploring Ideas and Examples

| Understanding latent space as a computational representation

A latent space is a mathematical representation used by many Artificial Intelligence systems to organise patterns learned from data. The word *latent* refers to structures that are present within the model but are not directly visible to the user. Instead of storing concepts in the same way that a dictionary stores separate definitions, the model represents them through numerical relationships distributed across many dimensions.

These dimensions do not correspond to simple categories such as "colour," "emotion," or "style." They are complex combinations of features learned during training. A concept may therefore be represented through relationships involving shape, language, context, tone, texture, function, and many other characteristics at the same time.

For teaching purposes, latent space can be imagined as a dense, volumetric cloud. The image conveys density, because many concepts

coexist and overlap; fluidity, because boundaries are blurred rather than sharp; and multidimensionality, because relationships cannot be reduced to a flat map. Concepts that share characteristics may be represented closer together, while concepts with fewer shared characteristics may be further apart. The cloud is only a metaphor: it is not a mind, a conscious archive, or the model's imagination.

For example, an image-generation model may learn relationships among:

An image-generation model may learn relationships such as concrete → architecture → monumentality → brutalism, or stone → ornament → historical architecture → baroque. Combining “brutalist architecture” with “baroque ornament” draws on learned relationships among visual forms, materials, styles, descriptions, and examples. The model does not understand their cultural history in the human sense.

The same distinction applies to an apparently imaginative request such as “a chair made of seawater”. The system connects learned patterns associated with chairs, water, transparency, waves, and liquid forms. The result may support human imagination, but its mechanism is computational representation rather than memory, emotion, or intention.

| Semantic proximity

Semantic proximity refers to how closely concepts are related within a model's learned representation. Concepts may be close because they frequently appear in similar contexts, share visual or linguistic characteristics, or are connected through patterns in the training data.

For example:

forest → trees → leaves → green → nature

These concepts are likely to have relatively strong relationships because they frequently appear together. However, proximity can also develop through less obvious features:

forest → density → repetition → vertical lines → architecture

The connection between a forest and architecture may not be immediate from a literal perspective. Nevertheless, the two can share characteristics such as repeated vertical forms, layered structures, shadows, scale, or spatial density.

Semantic proximity helps explain why AI can produce unexpected associations. When a model connects two concepts, it may be responding to characteristics that are not immediately visible to the user.

Obvious and distant associations

Some semantic connections are predictable:

coffee → cup → café → morning

Others appear more distant:

coffee → ritual → concentration → solitude → library

The second path involves more interpretative movement. An AI model may generate such a connection because these concepts appear in related textual or cultural contexts. However, the system does not experience morning rituals, concentration, or solitude. It recognises statistical relationships among how these concepts are represented in data.

Different proximity logics

Human and computational proximity should not be treated as identical.

For a person, two ideas may be close because of:

- personal experience;
- emotion;
- cultural identity;
- professional knowledge;
- memory;
- values;
- individual associations.

For an AI system, concepts are connected through learned patterns. The model does not possess personal meaning or lived experience.

For example:

brutalism → childhood neighbourhood → family memory → nostalgia

This may be a meaningful human association based on personal experience.

An AI system may instead connect:

brutalism → concrete → monumentality → geometric forms → dramatic shadows

The first path is biographical and emotional. The second is based on patterns in data.

The Latent Spaces describes these as two forms of navigation that may appear similar while remaining fundamentally unequal: people navigate culturally and personally meaningful associations, whereas AI operates through statistical, visual, linguistic, and pattern-based relationships.

The difference becomes clear with a word such as “home”. A model may return common associations-family, warmth, house, security, kitchen, comfort-while a participant may connect home with migration, two languages, distance, uncertainty, or a grandmother. Semantic proximity can support exploration, but personal meaning depends on a life and context that the model does not possess.

| Prompting as navigation through latent space

Users cannot directly observe or control a model’s latent space. They influence generation indirectly through descriptions, examples, constraints, selections, parameters, and repeated interaction. A prompt can therefore be understood as an approximate coordinate: it establishes a direction within a range of possible outputs, not an exact destination or guaranteed command.

A useful prompt usually identifies the subject, context, purpose, relevant characteristics, boundaries, and desired format. It may also indicate whether the task calls for a conventional answer, several alternatives, a distant association, or deliberate movement beyond an obvious pattern. These components are options rather than a formula; unnecessary detail can create contradictions or close down useful exploration.

- Subject: what the output concerns.
- Context and audience: where the output will be used and by whom.
- Purpose: what the output should help someone understand, decide, imagine, or do.
- Characteristics and constraints: what should be emphasised, limited, included, or avoided.
- Format: how the result should be organised or presented.
- Direction of exploration: whether to stay familiar, compare alternatives, combine distant concepts, or move away from a dominant pattern.

Prompting becomes navigation when the first output is treated as feedback rather than a finished answer. The cycle is: starting point → output → observation → one directional change → new output. The participant asks what became dominant, what disappeared, what remained generic, and what unexpected association appeared. The next movement should follow from that observation.

Changing one important variable at a time makes comparison clearer. If a community-library concept feels too generic, the participant might change only its atmosphere, material, cultural reference, social function, or central contrast. Comparing the results does not reveal the model's entire internal structure, but it provides practical evidence about how selected semantic changes influence generation.

Three simple movements are especially useful: adding introduces a new concept or characteristic; removing excludes a dominant or unwanted pattern; and replacing exchanges one characteristic for another. The purpose is not to perfect a prompt in isolation, but to make

each change deliberate enough that the participant can learn from the difference it produces.

| Seeds, variation, and controlled comparison

In some generative image systems, a *seed* is the numerical starting point used for generation. When the prompt and seed remain the same, the system will often produce the same or a very similar result. When the prompt remains the same but the seed changes, the composition or details may change because the generation begins from another numerical starting condition.

Seeds help participants understand that variation does not necessarily indicate intention or a new idea formed by the system. Different outputs can emerge from the same prompt because generative systems include probabilistic processes.

For example:

same prompt + same seed → similar result
same prompt + different seed → different variation
changed prompt + same seed → clearer comparison of semantic influence

Not every AI tool allows users to view or control seeds. Where this function is unavailable, participants can still compare several responses generated from the same prompt.

The purpose is not to teach advanced technical control. It is to demonstrate that outputs are shaped by both semantic instructions and computational variation.

| Unexpected outputs as positional feedback

Unexpected results can have different causes. The prompt may be ambiguous, the model may rely on dominant patterns, the concept may be poorly represented in the training data, or several instructions may compete with each other.

In factual and high-stakes contexts, errors must be identified and corrected. An inaccurate legal, medical, historical, or professional statement should never be treated as valid merely because it is interesting.

In exploratory creative work, an unexpected output can function as positional feedback: it shows how the system interpreted the current direction and clarifies where the participant has not yet arrived. This does not make every error creative. The appropriate response depends on the type of mismatch and the consequences of using it.

- Archetypal result: the output is predictable and generic, suggesting that the interaction is close to a high-density pattern. A useful response is to shift one variable, add an unexpected element, or remove the dominant one.
- Misunderstood direction: the output is incoherent with the intention, suggesting ambiguity or competing instructions. A useful response is to identify the essential core and reformulate rather than simply adding more words.
- Adjacent possibility: the output misses the intention but contains a useful phrase, shape, contrast, or association. A useful response is to identify that element and decide whether it offers a direction worth following.

This distinction is explored in the Museum of Hallucination activity. The word “hallucination” is used there as a creative frame for examining deviation and plausible explanation. In factual, legal, medical, historical, or other high-stakes contexts, fabricated or inaccurate content must instead be identified, corrected, and rejected as evidence.

- Is this output accurate and reliable?
- Does this output contain an interesting association that can support exploration?

These questions are not interchangeable.

An inaccurate statement cannot become fact because it inspires an idea. At the same time, a failed creative output may contain material

that supports a new direction. Accuracy and creative usefulness are separate questions, and the participant remains responsible for deciding which question applies.

| From output consumption to active navigation

Understanding latent space changes the role of the user.

A passive user asks the AI to produce a finished answer.

A navigator:

- establishes a starting point;
- observes the result;
- identifies dominant associations;
- decides which direction to follow;
- changes selected variables;
- compares alternatives;
- introduces personal and professional context;
- verifies the final result;
- accepts responsibility for selection and use.

This process does not mean that the user controls every element. AI interaction always involves uncertainty. The objective is not total control but conscious orientation.

The central working principle is:

A prompt establishes a direction; an output provides evidence of the current position; reflection determines the next movement.

| Latent space as geography: archetypes

A space is not simply a set of equidistant points. It has a structure. Some zones are crossed continuously, others remain almost unexplored. Understanding this structure changes how participants move within it.

The first geographic feature is the archetype. In this chapter, an archetype is not a Jungian symbol but a high-density zone where many recurring patterns from training data converge, making a familiar response statistically likely.

For example, if a visual model is asked to generate "a charismatic leader," the result will almost certainly contain recurring physical features, colours, and posture. Not because the model holds an opinion about what leadership means, but because that combination of elements has such high statistical density that the output becomes nearly inevitable.

Recognising archetypes helps participants notice where a model returns to dominant patterns and where a deliberate shift may be needed to produce a less predictable result.

| Overlap zones

The second geographic feature is the overlap zone: an area where apparently distinct concepts share visual, linguistic, or structural patterns.

A useful example is the relationship between brutalism and baroque: culturally distant, yet close within latent space because they share features such as complex geometry, monumental scale, and a strong use of contrast between light and shadow. The model holds no knowledge of art history: it recognises that certain visual patterns recur in both contexts and places them close together.

These zones are counter-intuitive for human imagination, which tends to work through culturally constructed and often rigid categories. Overlap zones are also the points where the model produces its most statistically original outputs - and this is where a participant's intentionality can turn that statistical originality into a meaningful result.

| Unexplored territories

The third geographic feature is the unexplored territory: a region of possible associations that users rarely seek because typical requests follow familiar paths.

These territories may be approached through unusual combinations, lateral requests, or systematic changes to one element at a time. They are not literally empty or undiscovered spaces; the term describes less frequently activated relationships within the model's learned representation.

Archetypes, overlap zones, and unexplored territories are not three isolated categories, but three different densities within the same space: from maximum predictability (archetype) to maximum novelty (unexplored territory), passing through counter-intuitive hybridisation (overlap). Recognising which zone one currently occupies is part of the navigation competence introduced in 4.8.

This geography connects with the exploration-exploitation polarity introduced in Chapter 3. Archetypes correspond to maximum exploitation: dense, familiar, and predictable. Unexplored territories correspond to stronger exploration: rare, uncertain, and less frequently activated. Overlap zones sit between them, offering combinations that may balance novelty and control. Navigation therefore involves deciding when familiarity serves the purpose and when a deliberate movement away from it is needed.

| Practice, Reflection, and Discussion

EXERCISE

1: Building a semantic proximity map

Choose one starting concept connected to the group's work, such as "leadership", "learning", "community", or "sustainability". Individually, write five concepts that feel close and three that feel

distant. Mark which connections come from personal experience, professional knowledge, culture, emotion, or memory.

Ask an AI system for associations around the same starting concept and organise its responses into an approximate map. Compare the two maps. Which relationships overlap? Which human connections are biographical or situated? Which AI-generated links appear conventional, unexpected, stereotypical, or difficult to explain? The objective is not to discover the model's real internal map, but to compare different logics of association.

EXERCISE

2: Prompt navigation - one movement at a time

Begin with a broad request relevant to a professional or educational task. Generate one output and treat it as positional feedback. Identify what the system understood, what became dominant, what disappeared, and what remains too conventional or unclear.

Choose one directional change only-add, remove, or replace one important element-and generate again. Repeat once more if useful. Compare the versions and explain which change produced which movement. End by deciding whether the latest output should be developed, redirected, or rejected.

EXERCISE**3: The Museum of Hallucination**

The Museum of Hallucination, conceived by Massimo Gerardo Carrese as a fantasiological application, explores low-probability semantic connections. A participant enters a word and the system produces a distant image together with an explanation of the path connecting them. The explanation may describe an actual generative relationship or may be a plausible story constructed after the result; the uncertainty is part of the activity.

For example, “morning” might become an automated factory through the shared idea of activation; “bread” might become a geographical atlas through distribution and networks; and “silence” might become a power plant through invisible energy. The educational value lies less in the aesthetic output than in examining how deviation is made coherent and how a viewer interprets it.

Creative deviation must remain distinct from factual reliability. In this exercise, participants are invited to explore an improbable association. They should not interpret the system’s explanation as verified access to its internal reasoning or treat fabricated factual content as valid.

Activity

Choose an ordinary word. Ask the system to move away from its immediate meaning and create a distant image or concept. Then ask for a short chain of associations connecting the starting word with the result. Participants may work with their own words or choose from: library, river, key, night, or mirror.

- What is the semantic chain between the starting point and the result?
- Which step appears conventional, surprising, culturally specific, or weak?
- Does the explanation reveal a useful relationship or merely make the result sound inevitable?
- What human interpretation gives the deviation meaning?
- Would this way of treating an error be appropriate in a factual or high-stakes task? Why or why not?

Groups select one deviation to “archive” with its starting word, result, explanation, and a short critical note. The note should identify both its creative potential and the limits of the system’s explanation.

EXERCISE

4: Reverse generation - reconstructing a possible creative path

Participants receive an AI-generated image or text without information about how it was produced. The objective is not to recover the original prompt, which may be impossible, but to reconstruct one plausible creative logic behind the visible result. This is an exercise in interpretation rather than technical prompt recovery.

Step 1 - Observe before interpreting

Describe the visible elements: objects, forms, colours, characters, spaces, repetitions, transformations, and contradictions. Keep observation separate from explanation during this first pass.

Step 2 - Identify relationships

Move from isolated elements to connections. Ask what changes into what, which distant elements appear linked, whether a hidden rule organises the image, and how an impossible situation becomes meaningful.

Step 3 - Formulate a possible generative idea

Write one sentence that describes the hidden logic rather than merely listing objects. For example, “architecture gradually transforms into living forms” expresses a relationship, while “a house, a tree, and a bird” only inventories the surface.

Step 4 - Compare through controlled generation

Use the reconstructed idea to generate two or three new outputs. Multiple versions help distinguish interpretive variance-whether the group understood the relationship-from generative variance-the differences produced by a stochastic system even when the direction remains the same.

- What remains stable across the generated variants?
- Which differences are likely to be generative variation rather than a failure of interpretation?
- Does the reconstructed idea capture the original relationship or only its visible objects?
- What alternative interpretation could another group reasonably propose?

The answer is not unique. The exercise develops visual literacy, critical interpretation, idea formulation, awareness of stochastic variation, and collaborative creativity. Its final competence is not reproducing a hidden prompt but learning to read and question the structures behind a visible result.

CASE STUDY

A vocational training team asks an image system for a visual identity for a new programme on green construction. The first outputs repeatedly use leaves, green circles, hard hats, and generic modern buildings. Instead of choosing the most polished version, the team reads these outputs as evidence that the interaction is located inside a strong archetype.

The team removes literal ecological symbols, replaces “green” with “repair and long-term care”, and introduces a connection between traditional craft and future infrastructure. The next outputs are less predictable but also less immediately recognisable. Discuss which movement improved the concept, which introduced new problems, and how the team should balance accessibility, originality, professional purpose, and the risk of visual stereotypes.

ASSESSMENT TOOL

Assessment should focus on navigation and interpretation rather than on producing a visually impressive result. Participants can submit a short process record containing a starting point, two or three controlled movements, selected outputs, rejected outputs, and a reflection on what each change revealed.

- Explains latent space as a computational representation without anthropomorphising the model.
- Distinguishes human personal association from model-based semantic proximity.
- Uses prompt components selectively and changes variables deliberately.
- Interprets variation and unexpected outputs without confusing creative usefulness with factual accuracy.

- Can explain why an output was selected, redirected, transformed, or rejected.
- Maintains responsibility for context, verification, and final use.

A strong result is not the one that reaches a perfect output fastest. It is the one in which the participant can explain the movement through alternatives, recognise the model's dominant patterns, and make a reasoned final decision.

MODULE 5**INTERCONNECTIONS AMONG
IMAGINATION AND LATENT
SPACES**

| Introducing the Topic

Two spaces of possibility: connecting human imaginative processes with AI-supported exploration

Human imagination is connected to lived experience, emotion, intention, culture, memory, embodiment, and personal meaning. Latent space, by contrast, is a computational representation formed through statistical relationships learned from data. Nevertheless, both can be approached as spaces of possibility in which connections, variations, and transformations can be explored.

The interconnection appears during human-AI interaction. A person begins with a question, intention, problem, image, or possibility and expresses only part of it to an AI system. The system returns a phrase, structure, visual form, interpretation, or unexpected association. The person evaluates this material through their own knowledge and intentions and decides whether to accept, reject, modify, or redirect it. The creative process develops through repeated movement between human imagination and computational generation.

Fantasiological functions provide a practical bridge between these two spaces. Operations such as combination, separation, inversion, multiplication, reduction, and transformation can first be applied by the participant as imaginative techniques and then used to structure AI-supported exploration. For example, a participant may imagine combining a library with a forest and ask an AI system for different forms of that relationship. The system does not share the participant's intention or personal understanding, but it can expose additional patterns and associations for the participant to examine and develop.

The relationship is therefore collaborative but asymmetrical. Human beings establish purpose, interpret meaning, introduce context, and evaluate utility, originality, and ethical consequences. AI systems provide computational variation and access to a broad field of learned associations. Creativity emerges through the interaction between these different capacities rather than from either side operating alone.

The purpose of this chapter is not to prove that human imagination and AI work in the same way. It is to help participants design a relationship between two unlike processes. Human beings bring situated experience, intention, values, and the capacity to care about an outcome. AI contributes rapid pattern-based variation. The educational question is how to combine these contributions without confusing them or allowing the faster process to define the purpose of the slower one.

The word “collaboration” should therefore be used carefully. It describes an organised interaction, not equal agency or shared responsibility. The AI system does not become a colleague with intentions of its own; the human participant remains responsible for framing, interpretation, selection, and consequences.

Learning objectives

This chapter develops participants’ ability to connect human imaginative processes with AI-supported exploration while maintaining a clear distinction between human cognition and computational representation. Participants learn how Fantasiological functions can be translated into practical strategies for navigating latent spaces and organising human-AI creative processes.

Participants will:

- Understand imagination and latent space as different spaces of possibility. Participants recognise that both human imagination and AI systems can generate associations and

alternatives, while understanding that they operate through fundamentally different mechanisms.

- Identify similarities without treating imagination and AI as equivalent. Participants compare processes such as association, combination, variation, transformation, and iteration while avoiding the assumption that AI possesses consciousness, personal experience, or imagination.
- Apply Fantasiological functions to latent-space navigation. Participants learn how combination, separation, inversion, multiplication, reduction, and transformation can give a deliberate direction to AI-supported exploration.
- Understand the human-AI feedback loop. Participants examine how a human intention is selectively expressed to an AI system and how generated material returns to human imagination for evaluation, reinterpretation, and further development.
- Use AI to expand rather than replace human imagination. Participants practise using AI to encounter alternatives, question assumptions, and explore unfamiliar directions while maintaining responsibility for purpose, meaning, and final decisions.
- Recognise creative friction as a productive element. Participants understand that differences between human intention and AI output can create useful tension. Unexpected or incomplete results may reveal new directions when examined critically.
- Maintain human agency throughout AI-supported creativity. Participants learn to define goals, select relevant material, verify information, reject unsuitable outputs, and evaluate outcomes according to context, originality, utility, and ethics.
- Make human and AI contributions visible. Participants learn to describe which parts of a process were human-framed,

AI-generated, human-selected, and human-revised, rather than attributing the outcome vaguely to "collaboration".

- Reflect on dependence and interruption. Participants examine whether AI expands their imaginative process or begins to determine its starting point, pace, language, and range of possibilities.

| Exploring Ideas and Examples

Curriculum note: Chapter 4 contains the main explanation and practice of prompting. Chapter 5 refers to prompting only as the point where human intention becomes a partial instruction for AI-supported exploration. Chapter 6 applies clear descriptions to building tools, without reopening general prompt engineering.

| Two different spaces of possibility

Human imagination and AI latent space can both support the exploration of possibilities, but they do so in fundamentally different ways.

Human imagination develops through lived experience, memory, emotion, culture, knowledge, intention, and personal meaning. A person may connect two ideas because of a childhood memory, a professional challenge, a cultural reference, or an emotional response.

AI latent space is different. It is a computational representation of patterns learned from data. Concepts are connected through statistical relationships, recurring contexts, similarities, and differences encoded numerically within the model.

The comparison can be introduced through the following framework:

Human imagination	AI latent space
Personal	Statistical
Cultural	Computational
Context-rich	Pattern-based
Emotional	Numerical
Dynamic	Learned representation

The framework should not be interpreted as a complete definition of either side. Its purpose is to help participants recognise the different foundations of human and computational association.

Human imagination is personal because it develops through the individual's life, but it is also influenced by shared cultures, communities, language, and social environments. Latent space is statistical because it represents relationships learned from large datasets, but these datasets also contain cultural patterns, values, conventions, stereotypes, and inequalities produced by people.

The distinction is therefore not between a completely human world and a completely neutral machine world. AI systems are trained on human-produced material, but they do not experience or understand that material as people do.

Discussion: Where the analogy works

The analogy between imagination and latent space is useful because both can involve:

1. starting from existing information;
2. building connections among concepts;
3. exploring alternatives;

4. combining distant ideas;
5. producing variations;
6. moving through repeated experimentation;
7. encountering unexpected results.

For example, both a person and an AI system may connect the ideas of "forest" and "architecture." A person may imagine a building inspired by memories of walking among trees. An AI system may connect forests and architecture through patterns such as vertical lines, repetition, density, natural materials, light, and shadow.

Discussion: Where the analogy breaks

The same caution introduced in 4.1 applies here: latent space has no personal memory, emotions, intentions, cultural identity, conscious understanding, independent imagination, or responsibility for the final outcome. The comparison with human imagination should always be read against this background.

AI can generate a relationship between concepts, but it does not experience why that relationship matters. The human participant provides purpose, meaning, context, and interpretation.

A useful synthesis is to think of human imagination as a compass and latent space as a map. The compass represents direction: intention, desire, lived experience, and the question of why a journey matters. The map represents a patterned territory of learned associations, but it does not choose a destination. A compass can point beyond the available map, while a map can offer many routes without knowing which one is worth taking. Used together, they support navigation; they do not become the same instrument.

| From human imagination to latent-space exploration

Chapter 4 explains prompts as approximate coordinates and develops practical techniques for navigating latent space. This chapter does not repeat those techniques. It focuses on the more important transition that occurs around the prompt: how a situated human intention is made

partial and explicit, and how generated material returns for interpretation.

The movement can be represented simply as:

human experience and intention → selective description → AI-generated variation → human interpretation and decision.

The description can never transfer the participant's full imagination into the system. Personal memory, embodied experience, cultural references, values, and tacit professional knowledge are only partly expressible.

For example, a participant may imagine a “welcoming learning space”. Before using AI, the participant should identify what welcoming means in this context: accessibility, dignity, multilingual communication, informal conversation, quiet individual work, or something else. AI can then offer variations around the selected qualities, but the participant must recognise when common visual conventions replace the more specific intention.

The central question is not whether the participant wrote a technically perfect prompt. It is what survived the translation, what was flattened or omitted, what the generated alternatives revealed, and what the participant chooses to restore, reject, or develop.

| From AI output to human reinterpretation

This lesson reverses the direction of the previous process. Instead of beginning with a developed human concept, participants begin with AI-generated material and examine how human interpretation can give it purpose and meaning.

An AI system can rapidly produce combinations, variations, images, phrases, structures, or scenarios, but the output does not arrive with an inherent professional purpose. For example, the phrase “a library that behaves like a garden” could suggest a space that grows through community participation, changes with the seasons, supports ecological learning, or offers different forms of quietness and diversity.

The generated phrase provides material; the human participant decides what the metaphor means and whether it is useful in a real context.

Human reinterpretation may involve:

1. Connecting the output with lived experience.
2. Introducing cultural meaning.
3. Identifying an unplanned metaphor.
4. Placing the material within a real professional challenge.
5. Rejecting stereotypical or irrelevant elements.
6. Changing the function of the original output.
7. Combining it with existing knowledge.
8. Transforming it into a practical concept.

AI outputs do not always correspond with the user's expectations. The gap between intention and result can create creative friction: a surprising output may suggest an unforeseen direction, an incoherent result may show that the intention was not represented, and an almost-useful result may invite targeted iteration. Friction is not automatically valuable, and not every error deserves interpretation.

Creative friction becomes productive only through human judgement. The participant must be able to explain why an unexpected element matters and decide whether to select, transform, develop, or reject it in relation to the purpose of the work.

Participants should distinguish among three responses:

1. Acceptance: the output is useful and can be developed further.
2. Rejection: the output is inaccurate, irrelevant, stereotypical, unsuitable, or ethically problematic.
3. Reinterpretation: the output is not useful in its original form, but one element can be transformed into a different direction.

| Fantasiological functions as a bridge

Fantasiological functions provide structured ways to transform information. They begin as human cognitive operations and can then give a deliberate direction to AI-supported generation. The functions do

not provide direct access to latent space; they organise the participant's movement through possible associations and reduce unstructured generation.

Combination

Combination brings two or more elements together. A participant might ask what happens when a library and a garden are combined, then use AI-supported exploration to examine different spatial or social forms of that relationship. The system contributes learned associations, while the participant decides whether the combination carries meaning and practical value.

Separation

Separation removes, isolates, or fragments an element that is normally treated as fixed. Imagining a school without classrooms, for example, can open alternatives based on workshops, community spaces, mobile units, or workplaces. AI may expand the alternatives, but the human purpose of the separation is to expose and question an assumption.

Inversion

Inversion reverses a familiar relationship, order, direction, or function. Reversing "teacher teaches learner" into "learner teaches teacher" could lead to an activity in which learners share their digital practices and trainers build the lesson from those practices. The value lies in deliberately moving away from a conventional pattern, not merely reversing words.

Multiplication

Multiplication increases the number, scale, frequency, or variety of an element. Moving from one solution to many parallel solutions supports divergence and can reveal alternatives that would otherwise remain unexplored. AI can generate variations quickly, but the participant must ensure that they are genuinely different and later organise and evaluate them.

Reduction

Reduction removes elements to identify what is essential. A complete training programme might be reduced to one objective, one activity, and one outcome, or a two-hour workshop might become a 20-minute activity without losing its central purpose. This function is useful when generated material is too complex, repetitive, or unfocused.

Transformation

Transformation changes form, function, audience, material, medium, or context. A written policy, for example, might become a participatory role-playing activity for VET learners. The change of form is useful only if the participant preserves what matters while adapting the material to real users and conditions.

| Designing a hybrid creative workflow

A hybrid workflow does not mean alternating randomly between human and AI activity. It requires a clear understanding of which tasks should remain human-led, which can be supported by AI, and where interpretation and evaluation must occur. One possible workflow contains five connected phases.

Phase 1 - Human framing

The participant defines the problem, purpose, audience, context, values, constraints, and criteria for a useful result. This phase should remain human-led because an AI system cannot independently decide what matters in a specific professional or educational context.

Phase 2 - Imaginative expansion

The participant applies one or more Fantasiological functions to create deliberate directions: combining services, inverting a conventional process, removing an assumed requirement, transforming a format, or multiplying possible audiences. AI may add variations, but the imaginative movement and the reason for using it should be consciously selected.

Phase 3 - AI-supported exploration

The selected directions are used to explore alternatives, visualise possibilities, compare structures, generate examples, expose conventions, or encounter unexpected associations. Detailed prompt construction is covered in Chapter 4. Here, outputs are treated as provisional material rather than completed solutions.

Phase 4 - Human interpretation and selection

The participant evaluates the generated material for relevance, accuracy, meaning, originality, feasibility, accessibility, usefulness, ethics, and professional quality. Material may be selected, combined, rewritten, transformed, or rejected; the criteria and the decision remain human responsibilities.

Phase 5 - Refinement and application

The selected idea is developed into a useful outcome and adapted to real users, resources, responsibilities, and consequences. The process may be repeated as needed: human framing → imaginative function → AI-supported generation → human interpretation → refinement → testing → revision.

Human and computational contributions are not interchangeable. AI is useful for speed, variation, comparison, and pattern-based association, while people remain responsible for intention, context, meaning, judgement, and final decisions. A short contribution record can identify the initial intention, the function used, the AI-generated material, the selection criteria, the main human transformations, and the final checks.

The workflow should also allow for stopping. If generation narrows the idea, introduces unsuitable assumptions, consumes attention without adding value, or makes the participant less able to explain the result, returning to human-only exploration may be the more creative choice.

Practice, Reflection, and Discussion

EXERCISE

1: Human-first ideation → AI variation

Choose a real professional or educational challenge. Before using AI, spend five minutes developing an initial response and record the experiences, values, images, and assumptions that shape it. Identify one element that feels essential and one element that may remain open to exploration.

Use AI to generate three variations without trying to transfer every detail of the original idea. Compare the results with the initial intention. What survived the translation? What became generic? What did the alternatives reveal that the participant had not noticed? Select, reject, or transform the material, and explain the decision.

EXERCISE

2: Generated material → human reinterpretation

Begin with several AI-generated outputs from an open request connected to the participant's field. Select one output that appears useful, one that appears irrelevant, and one that is difficult to classify. The system supplies material, but it does not determine what the material is for.

For each output, identify what a person would need to add through context, metaphor, lived experience, cultural meaning, practical application, critique, or refusal. Discuss whether reinterpretation reveals a productive direction or does too much work to rescue weak material.

EXERCISE**3: One challenge, three routes**

Choose a challenge relevant to the group, such as improving an adult-learning activity, designing a community service, communicating a complex policy, or reimagining an onboarding process. Develop three routes before comparing them:

- Route A - human-only: develop the concept without consulting AI.
- Route B - human-first hybrid: create an initial concept, then use AI to expand or challenge it.
- Route C - AI-first material: generate open alternatives, then interpret and transform one of them.

Compare the routes by examining originality, personal meaning, speed, surprise, relevance, effort, and the participant's ability to explain the result. Do not rank one route as universally best. Ask which route served this challenge, this participant, and this stage of the process-and what each route made visible.

EXERCISE**4: Fantasiological function laboratory**

Start with one ordinary professional or educational object: a classroom, meeting, application form, library, training schedule, public information leaflet, or team role. Small groups each apply a different function-combination, separation, inversion, multiplication, reduction, or transformation-first without AI.

After groups have articulated the imaginative movement in their own words, use AI to generate several expressions of it. Compare the human function with the model's interpretation. Did

the output preserve the intended movement? Did it fall back on a familiar visual or linguistic pattern? Which new association is worth keeping, and which should be rejected?

Groups then exchange outputs. A second group must infer which function guided the result. Differences between intention and interpretation become material for discussing how clearly an imaginative operation survives its translation into AI-supported generation.

EXERCISE

5: Map the hybrid workflow

Select an idea developed in an earlier exercise and take it through the five workflow phases: human framing, imaginative expansion, latent-space exploration, human interpretation and selection, and refinement and application. Use a separate colour or annotation for each type of contribution.

- Where did the purpose originate?
- Where did AI genuinely expand the possibilities?
- Where did a conventional pattern enter the process?
- Which selection criteria came from the participant or group?
- Which important decision cannot be delegated?
- At what point would stopping AI use improve the process?

End by writing a short contribution statement: “The human contribution was... The AI contribution was... The final meaning and responsibility remain with...”. Discuss whether this account feels accurate or whether the boundaries are more difficult to draw.

CASE STUDY

A vocational training centre wants to redesign its orientation programme for adult learners from different professional, linguistic, and educational backgrounds. Staff begin with their own experience: learners often receive too much information on the first day, while the most important social and practical knowledge is learned informally later.

The team applies inversion (“learners orient the institution”), reduction (“what are the three things a learner must know today?”), and transformation (“turn the handbook into a participatory journey”). They use AI to generate possible activities, visual metaphors, multilingual scenarios, and alternative programme structures. Several outputs are polished but generic. One unexpected suggestion—a map built from questions contributed by previous learners—creates a useful direction.

Discuss where the concept came from, what the AI added, and what the team had to reinterpret. Would the learner-question map still be meaningful if the questions were generated rather than collected from real learners? Which parts require participation, local knowledge, and consent? At what point could the team mistake rapid variation for genuine inclusion?

ASSESSMENT TOOL

Assessment should focus on the participant’s process and judgement, not only on the attractiveness of an output. Participants can submit a short process portfolio containing an initial human idea, the Fantasiological function used, a brief record of AI-supported exploration, selected and rejected outputs, the final transformation, and a contribution statement.

- A clear distinction between human imagination and computational generation.
- A meaningful use of at least one Fantasiological function.
- Evidence of interpretation, selection, rejection, or transformation rather than automatic acceptance.
- A final concept connected to a real audience, context, or professional purpose.
- An ability to explain limitations, ethical questions, and the participant's responsibility for the outcome.

A strong result is not the one containing the most AI-generated alternatives. It is the one in which the participant can explain how the process developed, why particular material was retained or rejected, and how human meaning changed the generated material.

MODULE 6

BUILDING APPLICATIONS WITH AI

| Introducing the Topic

Application means using AI to build small, working tools and apps that solve concrete problems in professional and educational work. Thanks to tools such as Artifacts, Canvas, and Google AI Studio, this no longer requires knowing how to program. A user can describe in plain language what a tool should do, and the AI generates a working version that can be tested, refined, and shared.

Building a tool is not only a technical task: it is also an act of imagination. Deciding what to build, for whom, and how it should feel is where creativity and fantasy do most of the work.

Creating tools with AI is a collaborative and iterative process: human idea → description (prompt) → AI-generated tool → testing → refinement → a working, personalised outcome. The AI produces drafts of the tool, but the human participant remains responsible for defining the problem, judging whether the tool works, verifying that it is safe and accurate, and giving it a purpose worth building.

In this chapter, an application means a small, testable digital tool or prototype. It may be a quiz, checklist, calculator, visualiser, decision-support aid, or learning activity; it is not necessarily production-ready software. The aim is not to imitate professional software development or promise that every participant can build a production-ready system. It is to understand how an idea can become a testable prototype and how much human judgement is still required after the prototype appears to work.

Examples of accessible environments include Claude Artifacts, ChatGPT Canvas, and Google AI Studio Build mode. They do not offer identical capabilities or sharing conditions, so facilitators should select

one approved environment and teach the common workflow: define, describe, test, refine, and decide whether to share. Their features, availability, pricing, and data practices may change. Trainers should therefore teach the underlying process-define, describe, test, refine, and take responsibility-so that participants can transfer it to other tools.

Learning objectives

This chapter develops participants' ability to use generative AI to build small, working tools and apps without programming. Participants learn how to describe a tool in plain language, generate it with AI, test and refine it, and take responsibility for what they build.

Participants will:

Turn an idea into a tool. Participants learn to identify a concrete problem in their work or teaching and describe a small tool or app that would help solve it.

Build with AI and no code. Participants use tools such as Artifacts and Canvas to generate working tools from plain-language descriptions ("vibe coding"), without writing code themselves.

Iterate and refine. Participants test what the AI produces, identify what does not work, and improve the tool step by step through further instructions.

Bring creativity and imagination. Participants explore playful and imaginative tools, not only utilitarian ones, using fantasy to shape original ideas worth building.

Maintain responsibility for the tool. Participants recognise that AI generates the tool, but the user remains responsible for whether it works, whether it is safe and accurate, and how it is used and shared.

Test beyond the expected path. Participants learn to examine incorrect, incomplete, extreme, and unexpected inputs instead of testing only the route they designed.

Decide whether and how to share. Participants consider privacy, accessibility, transparency, maintenance, and human oversight before another person relies on the tool.

| Exploring Ideas and Examples

| From idea to tool: building apps with AI

Until recently, creating a piece of software meant writing code. This is no longer the case. Modern AI tools can turn a description written in ordinary language into a working app. The user explains what the tool should do, and the AI writes and runs the code behind the scenes. This way of working is often called “vibe coding”: you describe the result you want and let the AI handle the technical implementation.

Two tools make this especially accessible. Artifacts (in Claude) generate small, self-contained apps, visualisations, or interactive pages directly inside the conversation, where they can be used and tested immediately. Canvas (in ChatGPT and Gemini) opens a side workspace where a document or app is built and edited step by step alongside the chat. None of these require prior programming knowledge: what they require is the ability to describe clearly what you want, and to judge whether what comes back actually works.

Describing the tool well is the real skill. A vague request (“make me an app”) produces a vague result. A useful description answers a few practical questions: What problem does the tool solve? Who will use it? What should the user be able to do with it, step by step? What should it look and feel like? Are there rules or limits it must respect, such as not asking for personal data?

The clearer the description, the closer the first version will be to what you imagined. But the first version is only a starting point, not the finished tool. The most important part comes next: testing what the AI produced, noticing what is missing or wrong, and asking for specific improvements one at a time.

From a problem to a design brief

Before prompting, participants should distinguish a need from a solution. “We need an app” is already a proposed solution; “new staff repeatedly miss three steps in the onboarding process” describes a problem that can be examined. A checklist, conversation, revised document, or existing platform may solve it more simply. Building should begin only after participants can explain why an interactive tool adds value.

A concise design brief can contain six elements:

- Purpose: the specific problem or opportunity.
- Users: who will interact with the tool and what they already know.
- Actions: what the user should be able to do, in order.
- Outputs: what the tool should show, calculate, save, or recommend.
- Boundaries: what the tool must not ask, claim, decide, or store.
- Success: how the builder and intended users will know that it is useful.

This brief is not a technical specification. It is a way of making an imagined interaction concrete enough to test.

| Building a tool step by step

This lesson turns the ideas from 6.1 into a concrete method. We follow one example from start to finish: a small app that guides a team through the decision of whether, and how, to write a project application (a funding proposal). The same steps apply to any tool you might build.

Step 1 - Define the problem. Before opening any AI tool, be clear about what the tool is for. Here: teams often waste time debating whether an opportunity is worth applying for. The tool should walk them through a

few key questions and give a clear recommendation. Write this down in one or two sentences.

Step 2 - Describe the tool to the AI. Open the approved environment chosen by the facilitator and describe what you want in plain language. For example: “Build a simple interactive app that helps a team decide whether to apply for a funding call. Ask the user a series of yes/no and short questions-strategic fit, available time, budget match, chances of success-and at the end give a recommendation (apply / do not apply / discuss further) with a short explanation. Keep it clean and easy to use.” The AI generates a working first version you can try immediately.

Step 3 - Test it. Use the tool as a real user would. Click through it, enter realistic answers, and see what happens. Note anything confusing, missing, or wrong: a question that doesn’t make sense, a recommendation that feels off, a layout that is hard to read.

Step 4 - Improve it, one change at a time. Ask the AI for specific improvements rather than vague ones. For example: “Add a question about whether we have a partner already,” or “Show the score as a percentage,” or “Make the final recommendation more detailed.” Changing one thing at a time makes it easy to see the effect and to undo a change that doesn’t help.

Step 5 - Add imagination. A tool does not have to be purely functional. You could give the decision app a friendly tone, a visual “traffic light” result, or a playful summary. Imaginative touches often make a tool more useful because people actually want to use it. This is where the creative work of the previous modules pays off.

Other tools you could build with the same method include: an interactive checklist for onboarding new staff, a quiz that helps learners self-assess, a small calculator for a specific calculation your team repeats often, a timeline or planning visualiser, or a simple game that teaches a concept.

Step 6 - Test with other people. Ask at least one intended user to try the tool without guidance. Observe where they hesitate, misunderstand

a label, expect a different result, or use the tool in a way the builder did not predict. Do not explain the interface while they test it; their confusion is part of the evidence.

Step 7 - Sharing decision. Name the outcome honestly: experiment, prototype, internal tool, teaching demonstration, or tool ready for wider use. A working screen is not proof that the tool is reliable, accessible, secure, or suitable for real decisions.

| Testing, refining, and taking responsibility

An AI-generated app may look finished before it has been adequately examined. Buttons work, colours are consistent, and a result appears at the end. This surface completeness can hide incorrect logic, missing cases, inaccessible design, invented explanations, or unsafe handling of information. Testing is therefore not a final technical check. It is the stage at which the builder asks whether the tool deserves to be used by someone else.

Test the intended path

Begin with the journey the tool was designed to support. Complete every step using realistic information and compare the outcome with the original design brief. Check whether instructions are understandable, required actions are possible, calculations behave as expected, and the output answers the stated purpose. If the tool gives a recommendation, trace which answers influenced it. A recommendation that cannot be explained should not be treated as a decision.

Test the paths you did not imagine

People do not always use a tool as its creator expects. They may leave a field empty, paste a paragraph into a short-answer box, use an unexpected language, enter a negative number, choose contradictory answers, refresh midway, or misunderstand an instruction. Test these edge cases deliberately. The objective is not to eliminate every possible failure but to learn how the tool behaves when the ideal path breaks.

Ask: Does the tool prevent the action, explain the problem, recover safely, or produce a confident but meaningless result? What happens if a learner has limited digital confidence, uses assistive technology, or accesses the tool on a small screen? Who is excluded by the assumptions built into the interface?

Check information and logic

AI may generate code and explanations that appear plausible without implementing the intended rule correctly. Test calculations with examples whose answers are already known. Compare summaries and recommendations with the source material. If official guidance, organisational policy, or assessment criteria matter, attach or cite the authoritative material and verify that the tool follows it. Reference documents can reduce unsupported assumptions, but they do not remove the need for checking.

Protect people and data

A prototype should normally use fictional or non-sensitive data. Avoid collecting names, contact details, learner records, health information, confidential organisational material, or other personal data unless there is a justified purpose, appropriate permission, and an approved environment. Data minimisation is a design question: can the tool work without asking for this information? If data are entered, where do they go, who can access them, how long are they kept, and can the user remove them?

Make the tool understandable and accessible

A tool should not depend only on colour, small text, complex instructions, or one language variety. Use clear labels, readable contrast, meaningful error messages, keyboard-friendly controls, and plain language appropriate to the audience. Where multilingual use is expected, do not assume equal quality across languages; test with speakers and consider whether translation errors change the result or the user's access to information.

Decide what the tool may do

The higher the consequences, the stronger the safeguards must be. A playful idea generator and a tool that influences access to employment, education, funding, health support, or public services should not be treated alike. Some decisions should remain outside the prototype. A tool may structure a conversation or display criteria without ranking people or making the final choice.

Before sharing, state what the tool is for, what it is not for, what information it uses, what has been tested, and where human review is required. Also decide who will maintain it. A tool can become inaccurate when policies, links, platform behaviour, or users' needs change. Responsibility continues after the first successful demonstration.

Reflection questions

- What did the tool make easier, and what new work did it create?
- Which assumptions about the user became visible only during testing?
- Could the tool's output be mistaken for a decision even if it was intended only as support?
- Who can challenge or correct the tool when it is wrong?
- What evidence would be sufficient before you asked another person to rely on it?
- Would a simpler, non-AI solution now seem preferable?

| Practice, Reflection, and Discussion

EXERCISE

1: Build your first tool

Work individually or in pairs. Choose a small, concrete need from professional or educational practice. Suitable first projects include a self-assessment quiz, onboarding checklist, planning

visualiser, concept-learning game, simple calculator, or conversation guide. Avoid projects that require sensitive personal data or make high-stakes decisions.

Create a six-part design brief: purpose, users, actions, outputs, boundaries, and success. Then ask an available AI building tool to produce a first version. Test the main path and improve one element at a time. Keep a short record of at least three changes: what you noticed, what you asked the AI to change, and whether the change improved the tool.

When it matters, you can also attach reference documents to guide the tool-for example, an official EU project guide, internal rules, or a style guide-so the app follows the right sources rather than the AI's assumptions. Two example tools for inspiration: **example1**¹ (Find your ideal AI) and **example2**² (European Project Wizard).

Conclude by presenting the tool as a prototype. Explain the problem it addresses, one imaginative choice, one limitation, and one question that remains unresolved.

EXERCISE

2: Try to break the prototype

Exchange tools with another participant or pair. The tester should use the tool without instructions from its builder and try at least five cases: an ordinary case, an empty or incomplete case, an extreme case, a contradictory case, and a case

¹ <https://claude.ai/public/artifacts/244a2374-4750-4b84-87e4-fc47ec12fbb0>

² <https://claude.ai/public/artifacts/93a18a79-5323-4ef8-97dc-07c167ff0a38>

involving a different user need or language. Record what happens without immediately proposing a fix.

The builder then groups the observations under four headings: usability, logic and accuracy, inclusion and accessibility, and privacy or safety. Choose one problem to repair and one problem that should instead become a clearly stated limitation. Discuss why not every weakness can or should be solved by adding more features.

EXERCISE

3: Should this tool be shared?

Imagine that the prototype will move beyond the workshop. Decide whether it should remain a personal experiment, become an internal organisational tool, be tested with a small user group, or be shared publicly. Prepare a short statement explaining the decision.

- What evidence supports the tool's usefulness?
- What has not been tested?
- What data could users enter, intentionally or accidentally?
- Who remains responsible for checking outputs and updating the tool?
- How will users understand that it is a prototype or decision-support tool?
- Under what condition would you withdraw it?

Groups compare decisions. The objective is not to reward public release; deciding not to share can be a strong design outcome.

CASE STUDY

A small organisation had to decide, several times a year, whether to respond to funding calls. Discussions were long and often inconclusive. A team member described the decision process to an AI and, without writing any code, built a short interactive app that asks a few key questions (strategic fit, capacity, budget, chances) and returns a clear recommendation. After testing it on past decisions and refining the questions, the team began using it to structure their conversations. The app did not make the decision for them, but it made the discussion faster, fairer, and easier to document. Discuss: where in your own work could a similar tool help, and where would human judgement still be essential?

Extend the case: after six months, the organisation notices that staff treat the recommendation as a final decision and rarely discuss applications scored “do not apply”. A funding programme also changes its criteria, but the tool is not updated. Who should notice these problems? What should change in the interface, the organisational process, or the decision to use the tool at all?

ASSESSMENT TOOL

Participants are not assessed with a written test but on the tool they build. For each participant or pair, look at:

6. Problem is the tool aimed at a clear, real need?
7. Description-was the tool described to the AI clearly enough to produce a useful first version?
8. Iteration-did they test the tool and improve it with specific, one-at-a-time changes?

9. Imagination-is there a creative or engaging element beyond the purely functional?
10. Responsibility-can they explain what the tool does not do, its limits, and what they are accountable for (accuracy, privacy, appropriate use)? A strong result is not a perfect app but a working tool the participant understands, has improved, and can take responsibility for.

Suggested addition to the assessment:

11. Testing-did the participant examine unexpected inputs and learn from another user?
12. Sharing decision-can they justify whether the prototype should remain private, be tested further, or be shared, and identify who would maintain it?

MODULE 7**ETHICS AND ECOLOGY OF AI**

| Introducing the Topic

Artificial intelligence often enters our lives through an ordinary task: drafting an email, translating a sentence, generating an image, organising notes, or finding a quick explanation. The immediate question is practical: does the tool help? Yet every apparently simple interaction is connected to less visible questions. What data made the system possible? Whose language and knowledge shaped its answers? What physical resources keep it running? Who benefits from its use, who carries its risks, and who can refuse it? In this chapter, “AI” refers mainly to the generative tools explored elsewhere in the FABE curriculum, while recognising that some examples—such as automated ranking or prediction—belong to a wider family of AI-supported systems.

This chapter does not offer a formula for perfectly ethical or ecological AI use. Such a formula would be misleading. AI systems differ, their infrastructures change quickly, evidence is incomplete, and the consequences of a decision depend on its context. A harmless experiment and a decision that affects employment, education, health, welfare, or public information cannot be judged in the same way. Even the wish to act responsibly may produce competing demands: a tool may improve accessibility while increasing dependence on a commercial platform; it may save time for one worker while raising productivity expectations for everyone; it may support environmental research while consuming energy, water, minerals, and hardware.

The purpose of reflection is therefore not to arrive quickly at approval or rejection. It is to slow down the moment in which a tool becomes normal. What changes when AI is present? What becomes easier, and what may become harder to see? Which choices remain ours, and which have already been made by developers, employers, institutions, markets, and public policy?

Throughout FABE, AI has been approached as a space for imagination, experimentation, and professional practice. This final chapter places those possibilities beside responsibility. It asks participants to consider not only what AI can produce, but what forms of attention, judgement, creativity, labour, and material infrastructure make that production possible. It therefore returns to three ideas developed in earlier modules: latent space is a computational representation rather than human imagination; the human-AI feedback loop depends on human purpose and interpretation; and the human editing layer is where generated material is critiqued, refined, personalised, and verified.

Learning objectives

- Recognise ethical, social, and ecological tensions in AI-supported work.
- Examine whose perspectives and interests shape an AI system and its use.
- Distinguish lower-stakes assistance from uses that may seriously affect people.
- Discuss privacy, authorship, transparency, labour, and environmental responsibility in context.
- Connect ethical and ecological questions to the creative and professional workflows developed in earlier chapters.
- Articulate a reasoned position without assuming that every dilemma has one universal answer.

Opening pause

When did AI stop feeling unusual in your work or daily life?

Which of your uses feel chosen, and which feel expected?

What would you want to know about an AI system before trusting it-and why is that knowledge often difficult to obtain?

| Exploring Ideas and Examples

| Assistance, authority, and human accountability

An AI system can propose, classify, summarise, predict, or recommend, but it cannot carry moral or professional responsibility. That responsibility remains with the people and institutions that select the system, provide the data, interpret the output, and decide how it will be used. This is easy to say and harder to practise. Fluent language can look authoritative even when it is inaccurate, incomplete, or poorly suited to the situation. A human reviewer may technically approve an output while lacking the time, knowledge, or freedom to challenge it.

Earlier in the curriculum, the Museum of Hallucination treats unexpected semantic deviation as creative material. In a factual, educational, or professional context, however, “hallucination” usually names unsupported or fabricated content. The same surprising output may therefore invite exploration in one setting and require rejection in another. The difference lies not only in the output, but in its purpose, consequences, and the claims made about it.

Consider two uses. A facilitator asks AI for possible workshop titles. A recruitment team asks it to rank applicants. Both involve generated suggestions, but the consequences are not comparable. In the first case, an unsuitable result can be discarded with little harm. In the second, patterns in the data, the criteria, or the interpretation may alter somebody’s livelihood. The important distinction is not simply whether a human clicks “approve”. It is whether meaningful oversight is possible: can the person understand the basis of the output, notice an error, disagree without penalty, and explain the final decision to those affected?

Accountability also becomes blurred across organisations. A user may be told to work faster with an approved tool, while responsibility for privacy, accuracy, and harm remains individual. A vendor may describe

the system as merely supportive, while its recommendation becomes decisive in practice. When many actors contribute to an outcome, does responsibility become shared-or diluted until nobody can be held answerable?

Questions of accountability

At what point does assistance become authority?

Who can challenge an AI-supported decision, and what happens when they do?

Is human oversight meaningful if the reviewer lacks expertise, time, information, or institutional power?

Who should answer when an AI-supported outcome harms someone: the user, the organisation, the developer, the provider, or the regulator?

What would meaningful human oversight require in your own workplace, classroom, or training setting?

| Bias, knowledge, and the concentration of power

AI systems are trained on material produced within unequal societies. Their outputs can therefore reproduce stereotypes and exclusions, but bias is not limited to an obviously discriminatory sentence or image. It also appears in what is treated as normal, professional, universal, or worth knowing; in which languages work well; in whose histories are visible; and in which experiences are reduced to exceptions.

For example, an image prompt for a “successful leader” may repeatedly return a narrow demographic pattern. A request for “professional language” may privilege one cultural register and make other forms of expression appear less competent. A historical explanation may reproduce a dominant Western narrative while presenting it as neutral. Asking for decolonial, peripheral, or underrepresented perspectives can

interrupt this pattern, but it cannot guarantee that those perspectives are represented accurately. The system may imitate their vocabulary without carrying their lived, political, or historical context.

This makes bias a question of power as much as representation. A small number of companies decide how systems are trained, filtered, priced, and made available. Their tools increasingly mediate knowledge, creative production, workplace practices, and public communication. Wider access can be empowering, especially across linguistic, educational, or disability-related barriers. At the same time, dependence on privately governed systems may narrow public choices. Who is included in the promised accessibility, and on what terms? For trainers and VET professionals, unequal performance across languages also raises a practical question: does a tool that increases access for some learners create a second, less reliable level of support for others?

Questions of knowledge and power

Whose knowledge appears without being named, and whose knowledge must be specifically requested?

What is lost when a system turns conflicting histories and experiences into one smooth answer?

Can a tool be inclusive if the people represented in its data have little influence over its design?

What forms of power follow from controlling the systems through which others search, write, learn, and create?

When the quality of AI support differs by language, who is expected to notice and correct the inequality?

| Privacy, authorship, and transparency

Using AI often feels like speaking privately to a machine, yet prompts may contain names, personal histories, organisational knowledge, unpublished work, images, or confidential documents. Removing a name does not always make information anonymous: a role, city, unusual event, and project title may still identify a person. The deeper issue is not only whether a user intends to disclose sensitive information, but whether the system encourages disclosure by making conversation feel informal, helpful, and frictionless.

Creative work raises another uncertainty. Generative systems learn from enormous collections of writing, images, music, code, and other cultural material. Many creators question whether their work was used with meaningful consent, attribution, or compensation. At the output stage, authorship becomes difficult to describe with a simple label. There is a difference between using AI to test a structure after developing an argument and publishing a generated text with minimal review, but where should a line be drawn? Does transparency require naming the tool, explaining the stages in which it was used, or describing how the human contribution changed the result?

Disclosure can build trust, but it can also become symbolic. “Made with AI” says little about whether the system supplied a few alternatives or most of the work. Full disclosure may be impossible when providers themselves reveal little about training data, resource use, or system behaviour. What, then, can transparency reasonably mean in an opaque technological environment?

Questions of privacy and authorship

Why do we share information with a conversational system that we would hesitate to share with a stranger?

Questions of privacy and authorship

Can consent be meaningful when users do not know how their prompts or uploaded files may be processed?

What makes a work yours: the first idea, the wording, the selection, the revision, the experience behind it, or responsibility for the result?

What would an honest account of AI involvement need to tell an audience

| Work, education, creativity, and the value of difficulty

AI may remove repetitive work and create more time for learning, imagination, and human contact. It may also enable employers to demand more output in the same amount of time. Efficiency is not distributed automatically: the person who works faster may not be the person who receives the saved time or economic benefit. The transformation of work therefore cannot be understood only as a question of which tasks a system can perform. It also concerns working conditions, bargaining power, deskilling, job loss, new forms of monitoring, and the value assigned to human expertise.

A similar tension appears in education and creative practice. Immediate answers can support curiosity and make difficult material more accessible. Yet learning also depends on slower experiences: trying, getting stuck, searching, making mistakes, tolerating uncertainty, and developing an unfinished thought. If AI supplies polished language before a person has found their own position, does it support expression or quietly define what is worth saying? FABE's movement between exploration and exploitation offers one way to notice this boundary: AI may broaden the field of possibilities during exploration, while selection, meaning, verification, and final responsibility remain human-led.

The concern is not nostalgia for every difficult task. Some forms of difficulty exclude people and deserve to be reduced. Translation,

transcription, explanation, and adaptive support may expand participation. Other forms of difficulty are part of how judgement, confidence, patience, and an individual voice develop. The challenge is that the boundary between unnecessary friction and meaningful struggle is different for each person and context.

Questions of work and human capacity

Who gains the time saved by AI?

Which skills are we happy to outsource, and which would we regret losing?

When does support expand a person's agency, and when does it create dependence?

What is the value of an imperfect human voice in a world of instantly polished language?

Could producing more content make communication less meaningful rather than more?

Which parts of the FABE creative process-exploration, interpretation, selection, editing, or verification-would you want to keep deliberately human-led?

| The ecology of an apparently immaterial technology

AI is often described through the language of clouds, models, and digital services, but it depends on material systems. Minerals are extracted; chips and servers are manufactured; data centres occupy land and use electricity and water; networks move data; equipment is replaced; and obsolete hardware becomes waste. Environmental impact is therefore a lifecycle, not a number attached only to the moment of prompting.

The desire to know the footprint of a single prompt is understandable, but there is no universal figure. Resource use varies by model, hardware, data-centre location, electricity source, cooling method, length and type of output, and the number of versions generated. Text, image, audio, and video do not require the same computation. Providers often disclose too little for confident comparison. Uncertainty should not be used to claim that the impact is negligible, but neither should one dramatic number be treated as true for every system and situation. Here, “model” means the computational system that generates the output, while “provider” means the company or organisation that makes the system available. Users can rarely observe the model’s infrastructure directly, which is why organisational procurement and public disclosure matter alongside personal choices.

Local consequences matter. Water used for cooling has a different meaning in a water-stressed region than in a place with abundant supply. Electricity demand may influence prices and decisions about fossil, renewable, or nuclear generation. Mining and manufacturing costs may be borne far from the people who receive the benefits. Rapid infrastructure expansion can also reshape climate commitments and public investment. The geography of harm and benefit is part of the ethical question.

AI can also contribute to climate modelling, biodiversity monitoring, energy optimisation, scientific discovery, agricultural planning, and the analysis of large environmental datasets. These benefits do not cancel the resources consumed, just as resource use does not make every application equally unjustifiable. The relevant comparison is not “AI or no impact”, because all digital and institutional alternatives have material costs. It is: which method, at what scale, for whose benefit, compared with what alternative?

Individual restraint may matter, but focusing only on whether a person generates one more image can hide the decisions of companies, governments, investors, and procuring organisations. Users influence demand; providers determine infrastructure and disclosure; organisations determine scale; regulators determine obligations. Where

should responsibility sit when the largest choices are not available to the individual user?

Questions of ecological responsibility

What becomes invisible when AI is described as a cloud service?

Who receives the benefits of AI infrastructure, and which communities bear its extraction, water, energy, and waste costs?

How should we decide when uncertain environmental costs are proportionate to a hoped-for benefit?

Does efficiency reduce total resource use, or can cheaper and faster generation lead to more consumption?

Why is responsibility so often framed as an individual choice when infrastructure is controlled collectively and commercially?

What environmental information would an organisation need before it could call the procurement of an AI tool responsible?

| Practice, Reflection, and Discussion

The activities below are not designed to produce a single responsible-use protocol. Their purpose is to make assumptions, conflicts, and power relations visible. Facilitators should allow disagreement and uncertainty. Participants may revise their position during the discussion; a change of mind is evidence of reflection, not failure. Each activity can be adapted to the participants' sector and does not require technical expertise. Facilitators may select two or three activities rather than treating all five as a fixed sequence.

EXERCISE

1: The ordinary use, seen differently

Choose one recent or realistic use of AI from your work, studies, or daily life. Describe the task in one sentence without judging it. Then examine the same use through five lenses: the person using it; a person represented in or affected by the output; the organisation introducing the tool; the company providing it; and a community connected to its physical infrastructure.

For each lens, write down what may be gained, what may be lost, what remains unknown, and who has the power to change the situation. Do not try to balance the columns mathematically. Discuss which perspective was easiest to imagine and which was most absent from your first description.

- Did the task still appear simple after changing perspectives?
- Which consequences were distant in space, time, or institutional responsibility?
- Whose interests would define whether this was a successful use?

EXERCISE

2: A spectrum without fixed answers

Place the following scenarios on a line from “AI use feels readily justifiable” to “AI use feels difficult to justify”. Participants should position themselves before discussing the cases. They may stand physically along a line in the room or mark a shared scale. The objective is not consensus but the reasons behind different positions.

- A learner uses AI to translate instructions into a language they understand better.
- A communications team generates fifty image variations to choose one social-media post.
- A social worker uses AI to rewrite a public information leaflet in plain language.
- An employer uses an AI score to decide which applicants receive interviews.
- An artist uses AI trained on unlicensed creative work to explore a commercial concept.
- A research team uses a large model to identify patterns in environmental data.
- A colleague asks AI to summarise a confidential meeting transcript.

After discussion, participants may move. Ask what changed their position. Which facts would they need before deciding? Were some concerns treated as more important than accessibility, speed, livelihood, creativity, privacy, or ecological cost? Who was imagined as the decision-maker?

EXERCISE

3: Read against the smooth answer

Generate or provide an AI response to a question connected to the participants' field. The response should sound plausible and complete. In small groups, read it once for information and a second time for absence. Mark where the text sounds certain, where it generalises, and where it turns a disagreement into a neutral summary. This revisits the human editing layer from Chapter 6 but deliberately pauses before correction: participants

first examine framing, absence, and authority, not only accuracy and style.

- Which voices, regions, languages, or forms of knowledge are centred? >
- Which people are spoken about rather than allowed to speak?
- What would be difficult to verify?
- Does asking for “another perspective” genuinely change the frame, or only add a paragraph to it?
- What knowledge or lived experience does the group have that the answer cannot supply?

Rewrite one part only if rewriting helps reveal the issue. The point is not to perfect the AI response, but to examine why polished language can make its limits difficult to notice.

EXERCISE

4: The material journey of one prompt

Begin with a single prompt, such as generating an image for a workshop. On a large sheet, trace what must exist before, during, and after the interaction: training material and human labour; mineral extraction; chip manufacture; model development; data centres; energy and cooling; networks and devices; storage; hardware replacement; and electronic waste. Add people and places wherever possible. If time allows, compare two plausible ways of meeting the same purpose—for example, a short text response and several generated images—without assuming that one format is always environmentally preferable.

Next, circle the parts about which the group has reliable information. Put a question mark beside the parts that depend

on provider disclosure or changing technical conditions. Finally, draw arrows showing who can make decisions at each stage: user, organisation, vendor, infrastructure provider, investor, public authority, or affected community.

- Does uncertainty change the group's judgement, and in which direction?
- Where is individual choice meaningful, and where is it mostly symbolic?
- What information should providers or public authorities be required to disclose?
- How would the discussion change if the output supported a public benefit rather than casual consumption?

EXERCISE

5: A conversation across roles

Assign participants different roles in a fictional organisation considering an AI system: employee, manager, learner or service user, artist or knowledge creator, data-protection officer, environmental advocate, technology provider, and local resident near a data centre. Present a proposal-for example, using AI to produce educational materials and personalise learner feedback.

Each role should identify what they would need to know, what they fear losing, what benefit they hope for, and what conditions would make the proposal acceptable or unacceptable. The group then holds a meeting, but it must end without a vote. Instead, record the disagreements that remain and the evidence or authority needed to address them.

- Which roles were expected to prove harm, and which were allowed to promise benefit?
- Whose concerns were treated as practical, emotional, technical, or political?
- Would consensus necessarily produce a responsible decision?

CASE STUDY

The multilingual community programme

A community organisation receives funding to expand a multilingual education programme. Staff are enthusiastic about using generative AI to translate materials, create illustrations, adapt exercises for different literacy levels, and draft individual feedback. The tool could reduce costs and help the programme reach more people. The funder expects visible innovation and rapid delivery.

The organisation has limited information about how the tool handles uploaded learner work. Several participants use minority languages for which output quality is uneven. Local translators and artists worry that their work will be replaced or used only to correct machine output. Staff are already overstretched and welcome anything that saves time. Management proposes producing many variants and testing them quickly. No environmental information is available from the provider beyond a general commitment to efficiency.

Discuss the case without designing a checklist. What promises are being made, and by whom? What would count as evidence that access has improved? Which forms of labour become visible or invisible? How might errors affect speakers of different

languages unequally? What information can the organisation reasonably demand? If the project proceeds, which uncertainties remain ethically significant? If it does not proceed, who loses an opportunity? Where would you place the boundary between experimentation and institutional dependence? Finally, ask whether the proposed system is functioning as augmentation, replacement, or a mixture of both-and whether different participants would answer that question differently.

| Final reflection: Keeping the question open

Return to one AI use that you currently accept, avoid, or feel uncertain about. Write a short reflection beginning with: “I used to frame this as a question of..., but I now also see...”. Continue with three sentences: one about a person or community affected; one about a form of power or responsibility; and one about a material or environmental condition.

End with a question you cannot yet answer. Do not convert it into a personal rule. Share it with the group and listen for connections between questions. The aim is to leave with a more precise uncertainty: not paralysis, but an awareness of what any future decision would need to take seriously.

GLOSSARY

The glossary provides working definitions as the terms are used in this curriculum. Some concepts have broader or contested meanings in other disciplines.

Accountability: The responsibility of people and institutions for selecting, supervising, interpreting, and acting on AI-supported processes and outputs.

AI hallucination: Unsupported, fabricated, or inaccurate content produced by an AI system. In a clearly framed creative exercise, unexpected deviation may become material for reflection; it must not be treated as reliable factual evidence.

AI literacy: The ability to understand, use, question, and evaluate AI systems in relation to their purpose, limitations, consequences, and context.

Archetype: In the latent-space metaphor used here, a dense and familiar region where recurring patterns make conventional outputs more likely.

Artificial Intelligence (AI): Computational systems that perform tasks associated with capabilities such as classification, prediction, language processing, or generation. The curriculum focuses mainly on generative AI.

Bias: A systematic tendency in data, system design, outputs, or use that privileges some patterns or groups and may disadvantage, exclude, or misrepresent others.

Creativity: The organisation, transformation, and development of ideas into outcomes that carry meaning, purpose, or use.

Embodied knowledge: Understanding shaped through bodily experience, perception, action, material interaction, and situated life.

Exploitation: The creative mode concerned with selecting, testing, refining, optimising, and implementing promising possibilities.

Exploration: The creative mode concerned with opening possibilities, experimenting, making connections, and tolerating uncertainty before premature selection.

FABE: Fantasy/Imagination and Artificial Intelligence for Business Enhancement, an educational approach that combines human imaginative and critical capacities with responsible use of AI.

Fantasiology: The interdisciplinary framework developed by Massimo Gerardo Carrese for studying imagination/fantasy, imagery, and creativity as observable and trainable human faculties.

Fantasy/imagination: The human capacity to generate possibilities, alternatives, connections, and transformations beyond what is immediately present.

Generative AI: AI systems that produce new text, images, audio, code, or other outputs from learned patterns and user-provided direction.

Human-AI feedback loop: An iterative process in which a person provides direction, examines an AI output, interprets it, and decides how to revise, redirect, accept, or reject the next step.

Human editing layer: The stage in which people critique, verify, contextualise, personalise, and take responsibility for AI-assisted material.

HUPP: An experimental application created by Massimo Gerardo Carrese to compose, decompose, transform, and recombine information in support of Fantasiology-based project development.

Imagery: The formation of mental, sensory, visual, or symbolic representations through which possibilities acquire perceptible form.

Latent space: A multidimensional computational representation in which a generative model organises learned relationships among features, patterns, or concepts.

Logic of “Also”: A Fantasiology method that expands an object or concept by asking what else it could be or do without necessarily erasing its original identity.

Overlap zone: In the curriculum’s latent-space geography metaphor, an area where apparently distinct concepts share features or associations.

Polarity management: The deliberate movement between complementary modes-such as exploration and exploitation-without assuming that one mode permanently resolves the other.

Prompt: A user-provided description, question, instruction, example, or constraint that gives an AI system an approximate direction for generation.

Semantic proximity: The degree to which concepts are represented as related through learned patterns, contexts, or characteristics.

Seed: A value that helps initialise a generative process; using the same seed may support controlled comparison when other settings remain stable.

Transparency: The extent to which people can understand that AI is being used, for what purpose, under what conditions, and with which limits or uncertainties.

Unexplored territory: In the curriculum’s metaphor, less frequently activated combinations or relationships that may produce less predictable possibilities.

Vocational education and training (VET): Education and training oriented toward knowledge, skills, and competences for professional



**almost
▶ there**

Comparative
Research
Network:

**Ω⁺ ON
THE
EDGE**



© 2026 FABE project partners
CC BY-NC-SA 4.0



Co-funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor the granting authority can be held responsible for them.