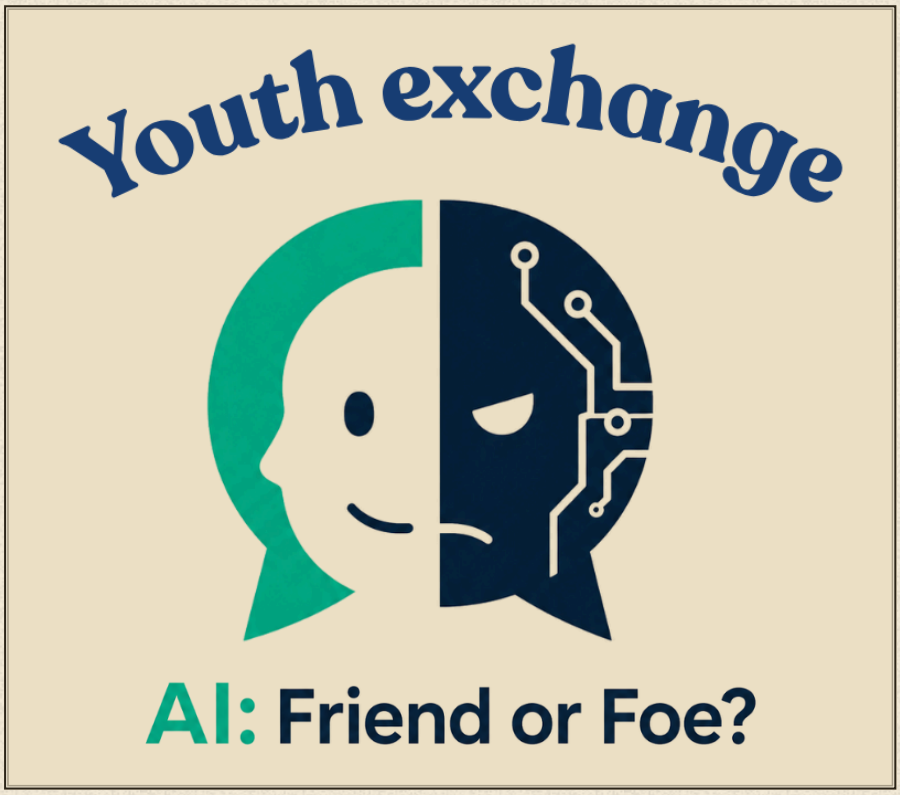
Co-funded by the European Union

Erasmus+
Enriching lives, opening minds.

A FIELD GUIDE TO ARTIFICIAL INTELLIGENCE

AI Awareness *Book*

How AI works, how to use it well, and how to question it.



Made by the participants of the youth exchange
Zadar, Croatia · 17–25 October 2026

Project no. 2026-1-HR01-KA152-YOU-000394792

PREFACE

Why a field guide *to AI?*

ZADAR 2026 KA152

> *A field guide helps you recognise what you meet in the wild: which bird is which, which mushroom is safe. Artificial intelligence is now part of our everyday world, in our phones, our feeds and our homework. This guide helps you recognise it, understand it and decide when it is a friend and when it is a foe.*

§ 0.1 The youth exchange

This booklet was born during **"AI: Friend or Foe?"**, an Erasmus+ youth exchange (Key Action 1, KA152) coordinated by **Srednja škola Glina** from Croatia. From 17 to 25 October 2026, 48 young people aged 15 to 18 from Croatia, Portugal, Romania and Serbia lived and learned together in Zadar, with their youth leaders.

The project started from a simple observation: young people use AI-driven technology every day, often for hours, but rarely get the chance to stop and think about it together. Over seven programme days, each national team designed and led its own workshops for the whole group, using non-formal methods: discussions, poster-making, theatre and role-play, hands-on testing of AI tools, and even writing a joint youth manifesto.

§ 0.2 The topics we explored

During the week we looked at artificial intelligence from many sides:

WHAT IS AI? · HISTORY OF AI · ETHICS & RESPONSIBILITY · THE FUTURE OF AI AND OUR LIVES · AI FOR STUDYING · AI FOR CREATIVE CONTENT · ACCESSIBILITY & INCLUSION · CREATIVITY & PROBLEM-SOLVING · AI & MENTAL HEALTH · ATTENTION & INSTANT GRATIFICATION · SOCIAL LIFE IN THE AGE OF AI · YOUNG PEOPLE SHAPING AI RESPONSIBLY

These topics opened many questions, some exciting and some worrying. They are also what inspired the participants to create this book. They wanted to make something **useful and educational to take back home**, for their classmates, friends, teachers and families who were not in Zadar with us.

§ 0.3 Who this guide is for

It is written for secondary school students aged about 14 to 19, but teachers, youth workers and parents can use it too. You do not need any technical knowledge. Read it from start to finish, or jump straight to the chapter you need.

PROJECT	AI: Friend or Foe?
PROGRAMME	Erasmus+ · KA152 Youth Exchange
NUMBER	2026-1-HR01-KA152-YOU-000394792
WHERE	Zadar, Croatia
WHEN	17–25 October 2026
WHO	48 young people (15–18) and their youth leaders
COUNTRIES	Croatia · Portugal · Romania · Serbia

SPECIMEN CARD *The youth exchange at a glance.*

PARTNER SCHOOLS

Srednja škola Glina
Glina, Croatia · coordinator

EB Sec. Bispo D. Manuel Ferreira Cabral
Santana, Madeira, Portugal

Colegiul Național „Kölcsey Ferenc"
Satu Mare, Romania

Sedma beogradska gimnazija
Belgrade, Serbia

FIELD NOTE

In this guide, "AI" means computer systems that do tasks we used to think needed human intelligence: understanding language, recognising images, making predictions or creating new content.

OUR PROMISE

No hype and no panic. AI is neither magic nor a monster. It is a tool made by people, and people decide how it is used. That includes you.

INDEX OF SPECIES

Contents

PART I — UNDERSTANDING AI

1 What is artificial intelligence? 5

2 How AI learns: data, algorithms and models 6

3 Generative AI: text, images, video, audio and code 7

4 What AI can do, and what it cannot do 8

PART II — AI TOOLS AND PRACTICAL USES

5 The AI toolbox 10

6 Choosing the right tool for the task 12

7 Prompting: how to ask better questions 13

8 AI for studying, creativity, accessibility and everyday life 14

9 AI models explained 15

PART III — RISKS, ETHICS AND RESPONSIBILITY

10 How to spot AI-generated content 17

11 Myths and facts about AI 18

12 Privacy, personal data and digital footprints 19

13 Bias, discrimination, misinformation and deepfakes 20

14 Mental wellbeing, social media and algorithms 21

15 Responsible AI use at school 22

PART IV — AI, INDEPENDENCE AND THE FUTURE

16 Local AI models: AI on your own device 24

17 Cloud versus local AI 25

18 AI and accessibility: technology for everyone 26

19 AI and the future of work 27

20 What kind of AI future do we want? 28

PART V — PRACTICAL GUIDE

21 Ten rules for using AI responsibly 30

22 Practical activities and experiments 31

23 Glossary of essential AI terms 32

24 Useful tools, resources and further learning 34

✦ Image credits and acknowledgements 35

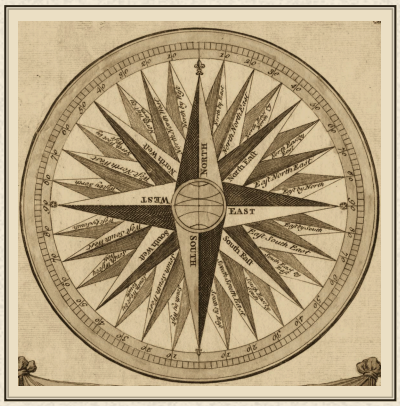


FIG. 0.1 *Compass rose from a 1748 English map. Every good guide starts with a way to find your bearings.*

How to read this guide

Each chapter is one "field note": a short, complete explanation of one topic. Along the way you will find these boxes:

FIELD NOTE
A surprising fact, number or piece of history that helps you understand the topic.

✎ TRY THIS
A quick experiment of five to fifteen minutes. Most need only a phone or a laptop.

⚠ CAUTION
A risk to watch out for, and what you can do about it.

FROM THE EXCHANGE
Ideas and conclusions from the workshops in Zadar.

MYTH	FACT
What many people believe...	...and what is actually true.

PART

I

Understanding *AI*

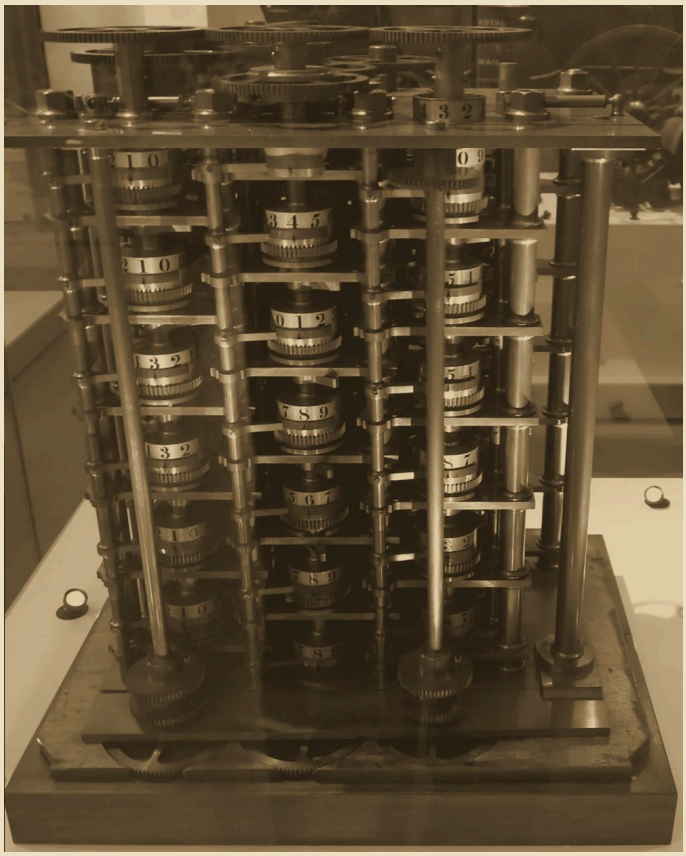


PLATE I *A section of Charles Babbage's Difference Engine No. 1, a mechanical calculator assembled in 1832. Babbage's friend Ada Lovelace imagined that such machines might one day work with music and symbols, not only numbers.*

01 What is artificial intelligence? 5

02 How AI learns 6

03 Generative AI 7

04 What AI can and cannot do 8

"I propose to consider the question, 'Can machines think?'"

ALAN TURING, 1950

CHAPTER I

What is *artificial intelligence?*

FIELD 01 NOTE

> *Artificial intelligence (AI) is the name for computer systems that do tasks we normally connect with human intelligence: understanding language, recognising images, making predictions, solving problems or creating new content.*

§ 1.1 A slippery definition

There is no single definition of AI that everyone agrees on. The European Union's AI Act describes an AI system as a machine-based system that works with some level of independence and **works out from the input it receives how to produce outputs** such as predictions, content, recommendations or decisions. The key word is works out: a normal program follows exact steps written by a programmer, while an AI system learns patterns from examples and uses them on new situations.

There is also a funny thing called the **"AI effect"**. As soon as something works reliably, people stop calling it AI. Spam filters, autocorrect, map apps that find the fastest route and the camera that recognises your face were all once "AI research". Today they are just "features".

§ 1.2 Narrow AI and general AI

All AI that exists today is **narrow** (or "weak") AI. Even the most impressive chatbot is a system built for a certain kind of task, such as working with text and images. It cannot decide on its own to learn to ride a bike or to care about you.

Artificial general intelligence (AGI) means a hypothetical AI that could learn and do any intellectual task a human can. Some tech companies say they are trying to build it. Experts strongly disagree about whether, when and how it could happen. For now it remains an open question, not a fact.

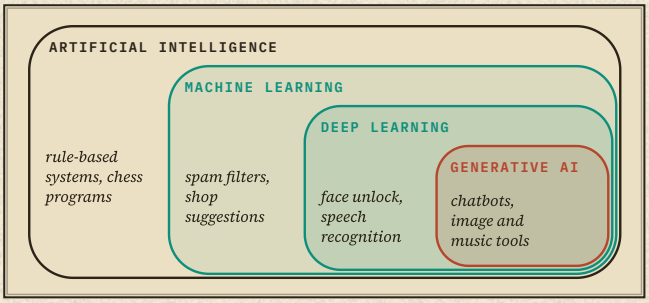


FIG. 1.1 *The AI family tree. Each circle sits inside a bigger one: generative AI is a kind of deep learning, which is a kind of machine learning, which is one way of building AI.*

§ 1.3 Where you meet AI every day

- **Your feed:** TikTok, Instagram, YouTube and Spotify use AI to guess what will keep you watching or listening.
- **Your phone:** face unlock, autocorrect, voice assistants, and the camera that makes night photos brighter.
- **Your inbox and maps:** spam filters, route planning, live traffic.
- **Your homework:** translators, chatbots, grammar checkers.
- **Behind the scenes:** bank fraud detection, online shop prices, job application screening, medical scans.

FROM THE EXCHANGE

The first workshop in Zadar asked mixed international groups to make posters on four questions: *What does AI mean to you? Where do you meet it every day? What are its benefits? What are its risks?* Try answering all four in one sentence each before reading on.

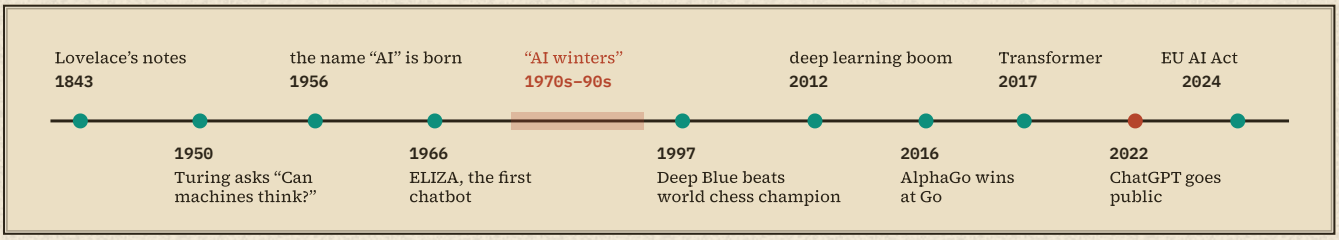


FIG. 1.3 *A very short history of AI. Progress came in waves, with "AI winters" when money and interest dried up after promises were not kept.*

CHAPTER II

How AI learns: *data, algorithms and models*

FIELD 02 NOTE

> *A normal computer program follows rules that a human wrote. Machine learning flips this around: humans give the computer lots of examples, and the computer finds the rules by itself.*

§ 2.1 Rules versus examples

Imagine a spam filter. The old way: a programmer writes rules such as *"if the email contains 'FREE MONEY', it is spam"*. Spammers write "FR33 M0NEY" and the rule breaks. The machine learning way: show the computer a million emails already marked "spam" or "not spam". It finds thousands of small clues by itself, many of which no human would think of.

§ 2.2 The three ingredients

DATA	The examples the system learns from: texts, photos, sound recordings, clicks, numbers. More and better data usually means a better model.
ALGORITHM	The learning recipe: the step-by-step method that adjusts the system so it makes fewer mistakes.
MODEL	The result of learning: a large file of numbers (called *parameters* or *weights*) that captures the patterns found in the data. This is the part you actually use.

Think of baking: data are the ingredients, the algorithm is the recipe, the model is the cake. Bad ingredients make a bad cake, however good the recipe. Computer scientists call this **"garbage in, garbage out"**: if the data is old, one-sided or full of stereotypes, the model copies those problems (see chapter 13).

§ 2.3 Inside a neural network

Most modern AI uses **artificial neural networks**, loosely inspired by brain cells. Layers of simple "neurons" take numbers in, multiply them by their weights, add them up and pass the result on. One neuron does almost nothing; billions together can recognise a cat or write a poem.

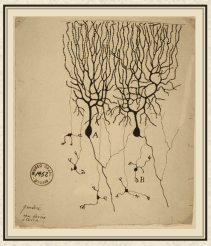


FIG. 2.1 *Brain cells drawn by Santiago Ramón y Cajal, 1899.*

The first artificial neuron that could learn, the **perceptron**, was built by Frank Rosenblatt in 1958. Its "eye" was a grid of just 20 × 20 light sensors. Today's large models have billions of parameters.

§ 2.4 Learning by mistakes

1. **Guess.** The network looks at an example and makes a prediction ("this photo is a dog").
2. **Check.** The answer is compared with the correct label ("it was a cat") and the error is measured.
3. **Adjust.** Every weight is nudged a tiny bit in the direction that would have made the error smaller.
4. **Repeat**, millions or billions of times, until the mistakes become rare.

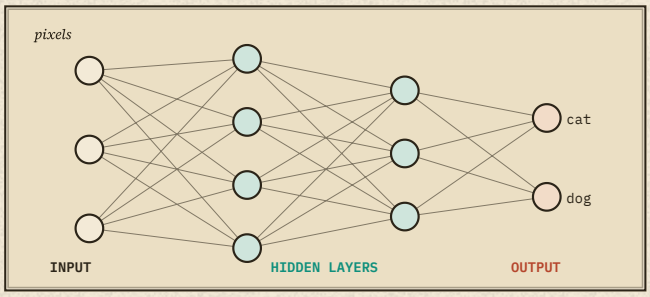


FIG. 2.2 *A tiny neural network. Every line is a weight. Training changes the weights until the right output "lights up".*

§ 2.5 Four ways to learn

Supervised learning uses labelled examples (photo + "cat"). **Unsupervised** learning finds groups in unlabelled data, such as customers with similar habits. **Reinforcement** learning works by trial and error with rewards, as in games and robots. **Self-supervised** learning hides part of the data and guesses it: chatbots learn by predicting the next word in billions of sentences.

§ 2.6 Training versus using

Training happens once, takes weeks or months and costs huge amounts of money and electricity. **Using** the model (*inference*) happens every time you ask something. A chatbot does not learn from your chat in real time, but the company may keep your chats to train future versions unless you switch this off.

FIELD NOTE

Meta said its Llama 3 models (2024) were trained on about 15 trillion tokens of text. A person reading 8 hours a day would need roughly **250,000 years** to read that much.

CHAPTER III

Generative AI: *text, images, video, audio and code*

FIELD 03 NOTE

> *Most older AI sorts, recognises or predicts: spam or not spam, cat or dog. Generative AI creates something new: an essay, a picture, a song, a video or a computer program, usually from a short instruction called a prompt.*

§ 3.1 Text: the next-word machine

Chatbots such as ChatGPT, Gemini, Claude, Copilot or Le Chat run on **large language models (LLMs)**. An LLM does one thing extremely well: it predicts which piece of text (a **token**, roughly ¾ of a word) is most likely to come next. Then it adds that token and predicts again, and again, hundreds of times per answer.

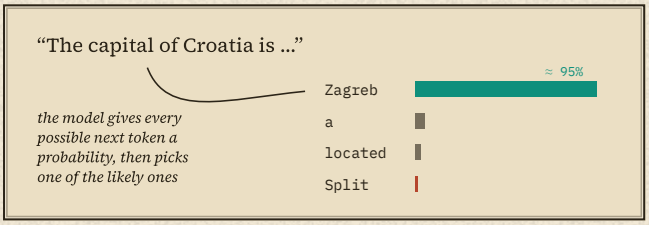


FIG. 3.1 *Next-token prediction (probabilities are illustrative). Because the model picks among likely words, the same question can get different answers.*

This is why chatbot answers sound fluent and confident. It is also why they sometimes invent facts: the model produces what *sounds* likely, not what it has checked. After the basic training, companies fine-tune models with human feedback so they follow instructions, are more helpful and refuse harmful requests.

§ 3.2 Images: from noise to picture

Image generators such as Midjourney, DALL·E, Adobe Firefly or Stable Diffusion mostly use **diffusion**. During training, the model watches millions of pictures slowly being covered in random noise and learns to reverse the process. To create a new image, it starts from pure noise and removes it step by step, guided by your text, until a picture appears.

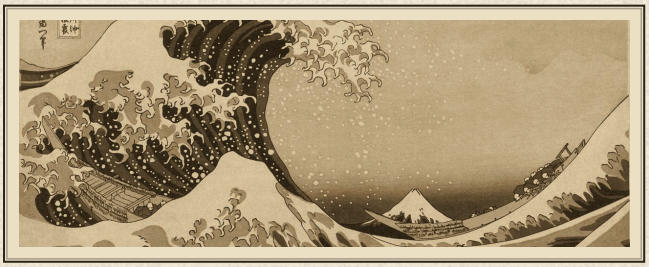


FIG. 3.2 *The Great Wave off Kanagawa, Katsushika Hokusai, about 1831. Ask an image generator for "a wave in the style of Hokusai" and it will copy his style in seconds. Is that inspiration or copying? Artists and courts are still arguing.*

§ 3.3 Video

Video generators (for example Sora, Veo, Runway or Kling) apply the same idea to many frames at once and must keep people and objects consistent from frame to frame. Clips are getting longer and more realistic, which makes fake videos much easier to produce (see chapter 13).

§ 3.4 Audio and music

AI can turn speech into text (live subtitles), text into natural speech, and generate whole songs with vocals from a one-line description (for example Suno or Udio). **Voice cloning** can copy a person's voice from a short recording. This helps people who are losing their voice, but it is also used in scam calls that imitate a family member.

§ 3.5 Code

Coding assistants such as GitHub Copilot or the big chatbots can write, explain and fix code. They are a great way to learn programming, but their code can contain bugs or security holes, so it has to be read and tested.

The newest models are **multimodal**: one model can read a photo of your maths homework, listen to your question and answer out loud.

FROM THE EXCHANGE

In the creative content workshop in Zadar, groups tested tools for images, video, animation, music and design. The workshop ended with a discussion on authorship, originality, copyright and sharing AI-made content responsibly.

✎ TRY THIS

Ask a chatbot the same question in three new chats. Are the answers identical? That is next-token prediction at work.

CHAPTER IV

What AI can do, *and what it cannot do*

FIELD 04 NOTE

AI can beat world champions at chess and still fail to count the letters in a word. Knowing its strengths and weaknesses is the most important AI skill you can have.

§ 4.1 Where AI is strong

- **Finding patterns in huge amounts of data** that no person could read in a lifetime.
- **Medicine:** a large Swedish study (2023) found that breast cancer screening supported by AI detected 20% more cancers while cutting the doctors' reading workload by 44%.
- **Science:** AlphaFold predicted the 3D shapes of hundreds of millions of proteins. Its creators shared the 2024 Nobel Prize in Chemistry.
- **Language:** translating, summarising, rewriting and explaining texts in seconds.
- **Games with clear rules:** chess (1997) and Go (2016).
- **Never getting tired or bored** of repetitive tasks.

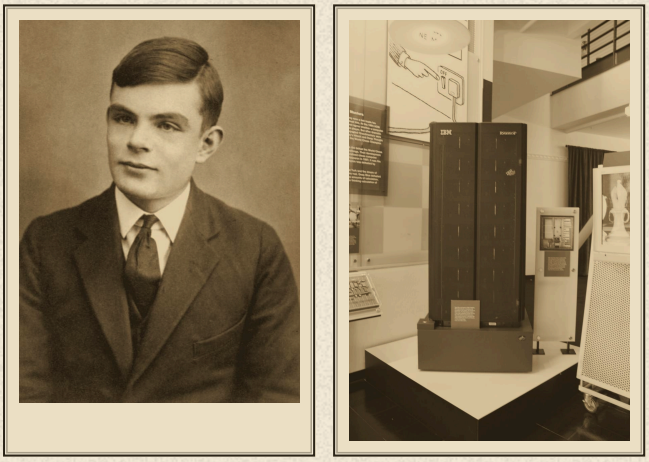


FIG. 4.1 *Left: Alan Turing at 16. In 1950 he proposed the "imitation game", now called the Turing test. Right: IBM's Deep Blue, which beat world chess champion Garry Kasparov in 1997.*

§ 4.2 Where AI is weak

- **Making things up.** Chatbots can invent facts, quotes, books and sources and present them with total confidence. This is called **hallucination**. In 2023, a New York lawyer was fined after submitting a court document with six court cases that ChatGPT had invented.
- **Knowing what it doesn't know.** A model rarely says "I'm not sure" unless it has been trained to.
- **Fresh information.** A model only knows what was in its training data, up to a cut-off date, unless it can search the web.
- **Common sense and the physical world.** It has never touched, tasted or fallen over anything.
- **Small, exact details.** Counting letters, precise maths or reading a clock face can still go wrong, because models see tokens, not single letters.
- **Values and responsibility.** AI has no feelings, no conscience and cannot be held responsible. People are.

§ 4.3 Moravec's paradox

In the 1980s, robotics researcher Hans Moravec noticed something strange: what is hard for humans (advanced maths, chess) is often easy for computers, and what is easy for any 4-year-old (walking on stairs, folding a towel, understanding a joke from someone's face) is very hard for machines. This is still largely true today.

§ 4.4 Passing the Turing test

Turing suggested that if a judge chatting by text cannot tell a machine from a human, we should take the machine's "thinking" seriously. In a 2025 experiment at the University of California San Diego, participants chose the chatbot GPT-4.5 as "the human" 73% of the time. Does this mean it understands? Most scientists say no: sounding human and understanding the world are two different things.

MYTH	FACT
"If it sounds smart and confident, it must be right."	Fluency is not accuracy. Treat every AI answer as a first draft, not a final truth.

✎ TRY THIS

Ask a chatbot about something you know very well: your hometown, your school, your favourite game. Count the mistakes. This is the best way to see hallucination for yourself.

FIELD NOTE

AI is a calculator for words, not an oracle. It is excellent at producing likely-sounding answers. You are responsible for checking whether they are true.

PART

II

AI tools and *practical uses*



PLATE II *A Jacquard loom with its chain of punched cards, National Museum of Scotland. Invented in 1804, it wove patterns automatically from instructions punched into cards: one of the first "programmable" machines and an inspiration for the first computers.*

05 The AI toolbox 10

06 Choosing the right tool for the task 12

07 Prompting: how to ask better questions 13

08 AI for studying, creativity and everyday life 14

09 AI models explained 15

"The Analytical Engine weaves algebraical patterns just as the Jacquard-loom weaves flowers and leaves."

ADA LOVELACE, 1843

CHAPTER V

The AI toolbox: *assistants and specialist tools*

FIELD 05 NOTE

> *There are thousands of AI tools, and new ones appear every week. You do not need to know them all. It is enough to know the main "species", what each is good for and what to watch out for.*

§ 5.1 General AI assistants (chatbots)

These are all-rounders: you can ask them to explain, write, translate, summarise, brainstorm, analyse a photo or help with code. Most have a free version with limits and a paid version (usually around 20 € per month) with stronger models and more use.

ASSISTANT	MADE BY	KNOWN FOR	KEEP IN MIND
ChatGPT	OpenAI (USA)	The best-known chatbot; text, images, voice, web search, file analysis.	Free tier has limits; check privacy settings for training on your chats.
Gemini	Google (USA)	Connected to Google Search, Docs, Gmail and Android; very long documents.	Often available through school Google accounts.
Microsoft Copilot	Microsoft (USA)	Built into Windows, Edge, Word and PowerPoint.	Often available through school Microsoft 365 accounts.
Claude	Anthropic (USA)	Careful writing, long texts and documents, coding.	Requires users to be 18+.
Le Chat	Mistral AI (France)	A European assistant; fast, with web search.	European company under EU rules.
Meta AI	Meta (USA)	Built into WhatsApp, Instagram and Messenger.	Easy to use without noticing; think before sharing.
Perplexity	Perplexity (USA)	An "answer engine" that searches the web and shows sources.	Sources make checking easier, but still check them.
DeepSeek	DeepSeek (China)	Strong free models, also downloadable (see ch. 16).	The app stores data in China; avoid sharing personal data.

§ 5.2 Free, paid and school accounts

Free versions are usually enough for school work. Before you pay, check whether your school already provides access: many schools use Microsoft 365 or Google Workspace for Education, which can include Copilot or Gemini with stronger data protection.

Remember the old internet rule: **if a service is free, you may be paying with your data**. Read what the tool does with your conversations, and look for a setting to turn off "use my data to improve the model".

⚠ AGE LIMITS

Most AI services require users to be at least 13, and many require parental consent until 18. Some are only for adults. In the EU, the age at which you can agree to data processing by yourself is between 13 and 16, depending on the country. Check the terms before you sign up.

§ 5.3 Assistants are not search engines

A search engine shows you *existing* pages written by people. A chatbot *generates* new text from learned patterns, so it can explain things in your words, and also invent things that exist on no web page at all. Many assistants can now search the web, but open the sources and check them.

FIELD NOTE

Tools come and go. In 1890, Herman Hollerith's punched-card machines counted the US census far faster than clerks by hand; his company later became part of IBM. Skills like asking clear questions and checking answers will outlive any app.



FIG. 5.1 *Hollerith's tabulating machine, 1902.*

§ 5.4 Specialist tools: a field catalogue

Specialist tools do one job very well. Many are built on top of the same big models as the chatbots, with extra features for one task. Here are well-known examples in eight "habitats". Names and prices change fast, so treat this list as a starting point, not a recommendation.

① STUDY & RESEARCH

NotebookLM (Google): upload your own notes or PDFs and ask questions only about them; can make quizzes and audio summaries.
Perplexity, **Consensus**, **Elicit**: search with sources, including scientific papers.
Khanmigo (Khan Academy): a tutor that guides you instead of giving answers.

② LANGUAGE & WRITING

DeepL (Germany): high-quality translation, also for Croatian, Portuguese and Romanian.
LanguageTool (Germany), **Grammarly**: grammar and style checks.
Duolingo: language practice with AI conversation features.

③ IMAGES & DESIGN

Canva (Magic Studio): posters and presentations.
Adobe Firefly: trained mainly on licensed images, marks its output with Content Credentials.
Midjourney, **Microsoft Designer**, the chatbots' own image tools.

④ VIDEO & AUDIO

CapCut: automatic captions and editing.
Runway, **Sora**, **Veo**: video generation.
Suno, **Udio**: songs from text.
ElevenLabs: realistic voices and dubbing (use only voices you have permission to use).

⑤ CODING

GitHub Copilot: free for verified students through the GitHub Student Developer Pack.
Replit, **Cursor**: code editors with AI built in.
Scratch + chatbot: a good start for beginners.

⑥ ACCESSIBILITY

Be My Eyes and **Seeing AI**: describe the world through the phone camera.
Live Transcribe, **Live Captions**: real-time subtitles.
Immersive Reader: reads text aloud and simplifies layouts (see ch. 18).

⑦ ORGANISING & EVERYDAY

Goblin.tools: breaks big tasks into small steps, made with neurodivergent people in mind.
Google Lens: point the camera to translate, identify plants or solve problems.
Copilot in Office, **Gemini in Docs**, **Notion AI**.

⑧ RUN IT YOURSELF

Ollama, **LM Studio**, **Jan**, **GPT4All**: free apps that run open AI models on your own computer, offline and private (see ch. 16 and 17).

§ 5.5 How to judge a new tool

Before you try a new AI app, check it like a biologist checks a new species:

- **Who made it,** and where is it based?
- **What happens to my data?** Can I delete it? Is it used for training?
- **What does it cost,** and what happens after the free trial?
- **Is it allowed** for my age and at my school?
- **Does it work in my language?** Quality in Croatian, Portuguese, Romanian or Serbian can be lower than in English.
- **Can I check its answers?** Does it show sources?
- **Does it promise too much?** "Undetectable homework", "see who viewed your profile" or apps asking for many permissions are often scams.

FROM THE EXCHANGE

In Zadar, groups tested tools hands-on in four workshops: AI for studying, for creative content, for accessibility and inclusion, and for creativity and problem-solving. Every group presented not only what the tool can do, but also its **limitations**. Do the same with every tool you use.

✎ TRY THIS

Give one task ("explain photosynthesis to a 12-year-old") to three assistants. Compare accuracy, clarity and tone. Which would you trust?

CHAPTER VI

Choosing the right tool *for the task*

FIELD 06 NOTE

> *A hammer is a great tool, but not for every job. The same is true for AI. Good users start with the task, not with the tool, and sometimes the best tool is not AI at all.*

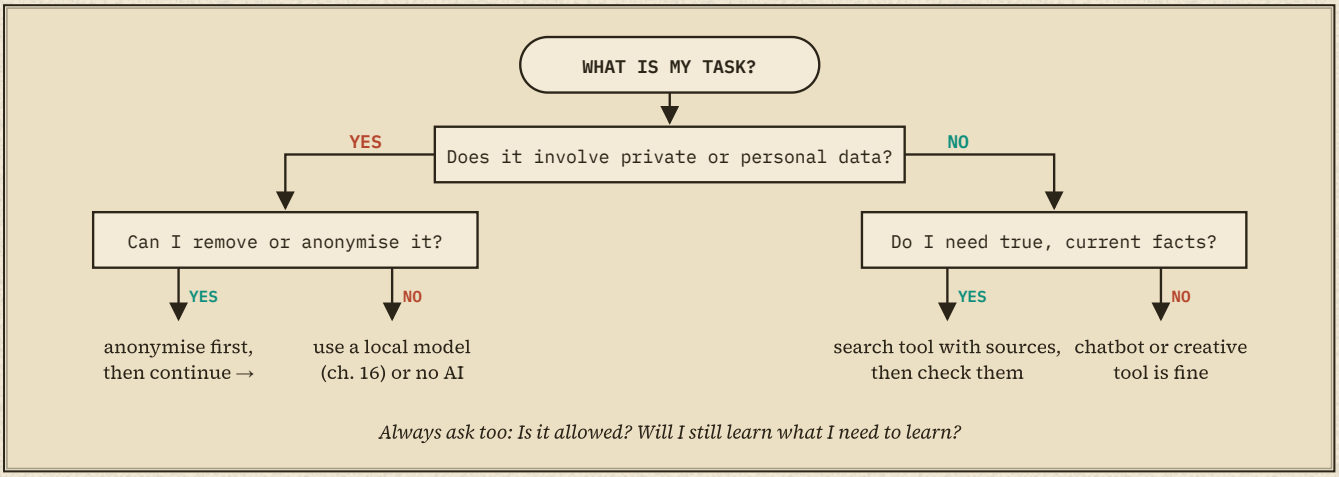


FIG. 6.1 *A simple decision path for choosing an AI tool.*

TASK	GOOD CHOICE	WHY / WATCH OUT
Understand a hard topic	Chatbot in "tutor" mode; NotebookLM with your class notes	Ask it to explain step by step and to quiz you. Check key facts in your textbook.
Find facts, dates, news	Search engine, Perplexity, chatbot with web search	You need sources you can open. A chatbot without search may be out of date.
Calculate exactly	Calculator, spreadsheet, Photomath / Google Lens	Chatbots can make arithmetic mistakes unless they run real code.
Improve your own text	LanguageTool, Grammarly, chatbot	Ask for *feedback*, not a rewrite, so the text stays yours.
Posters, slides	Canva, Firefly, image tools	Check spelling inside images; label AI-made visuals.
Anything with health, legal or money decisions	A qualified person (doctor, counsellor, teacher)	AI can help you prepare questions, but should not make the decision.

§ 6.1 Five questions before you start

1. **What exactly do I want?** An idea, an explanation, a fact?
2. **What is at stake?** A joke for friends is low-risk; a school grade or a health question is high-risk.
3. **What data am I sharing?** Never paste in passwords, ID numbers or other people's private information.
4. **Is it allowed?** Check your teacher's and school's rules (see ch. 15).
5. **Is AI even needed?** A dictionary, a library or a friend is often faster.



FIG. 6.2 *A library subject catalogue. Before search engines, finding the right drawer was the skill. Today it is choosing the right tool.*

FIELD NOTE

AI uses energy. Google reported in 2025 that a typical text question to Gemini uses about 0.24 watt-hours, like watching TV for nine seconds. Images and especially videos use much more, so generating 50 videos "just for fun" is not free for the planet.

CHAPTER VII

Prompting: *how to ask better questions*

FIELD 07 NOTE

A prompt is the instruction you give an AI. Vague prompts get vague, generic answers. Clear prompts get useful ones. The good news: prompting is just clear communication, the same skill that helps you with teachers, friends and future employers.

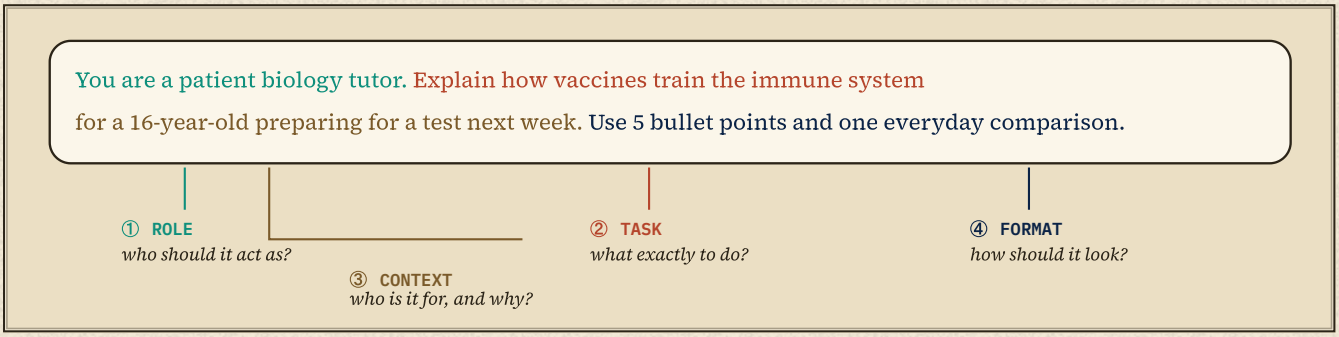


FIG. 7.1 *Anatomy of a good prompt: Role, Task, Context, Format. You don't always need all four, but each one makes the answer more useful.*

§ 7.1 Weak versus strong

✗ WEAK

"Tell me about World War I."

✓ STRONG

"I'm a 17-year-old in Croatia studying for a history test. Explain the four main causes of World War I in a short table, with one sentence on why each one mattered. Then ask me three questions to check if I understood."

§ 7.2 Seven techniques that work

1. **Be specific.** Say the topic, the level, the length and the goal.
2. **Give context.** "For a school presentation", "for my grandmother", "I already know X but not Y".
3. **Show an example.** Paste a sample of the style or format you want.
4. **Ask for steps.** "Explain step by step" makes reasoning clearer and easier to check.
5. **Let it ask you.** "Before you answer, ask me any questions you need."
6. **Iterate.** The first answer is a draft. Reply: "Shorter", "simpler", "give a counter-argument", "what did you leave out?"
7. **Ask for doubts.** "Which parts of your answer are you least sure about?"

§ 7.3 Three templates to copy

TUTOR MODE

"Act as my tutor for [topic]. Don't give me the answer. Ask me one question at a time, give hints when I'm stuck, and tell me what I got right."

FEEDBACK, NOT REWRITE

"Here is my essay. Give me three strengths and three concrete improvements. Do not rewrite it for me."

DEVIL'S ADVOCATE

"Here is my opinion: [...]. Give me the three strongest arguments against it, and sources I could check."

§ 7.4 What not to put in a prompt

Passwords, ID or health card numbers, home addresses, photos of other people without permission, private messages from friends, or anything you would not want a stranger to read. Assume that what you type may be stored.

FIELD NOTE

Most assistants understand Croatian, Portuguese, Romanian and Serbian, but they usually perform best in English because most of their training data is English. For important tasks, compare answers in both languages.

CHAPTER VIII

AI for studying, creativity, *accessibility and everyday life*

FIELD 08 NOTE

> *Used well, AI can be a patient tutor, a creative partner and a helpful assistant, available at 2 a.m. before a test. Used badly, it does your thinking for you, and then you learn nothing. The difference is how you use it.*

§ 8.1 Studying: use AI to think more, not less

Learning science shows that we remember best when we make our brain work: recalling, explaining, practising. So the most useful study prompts make *you* do the work:

- **Explain it differently:** "Explain osmosis with a football comparison."
- **Quiz me:** "Make 10 test questions from my notes, from easy to hard. Give answers only after I reply."
- **Flashcards:** turn your notes into question–answer pairs.
- **Find my gaps:** "Here is my summary. What important ideas are missing?"
- **Study plan:** "I have 4 tests in 2 weeks and 1 hour a day. Make a realistic plan."
- **Language practice:** have a conversation in the language you are learning, and ask for corrections.

FIELD NOTE

In a 2024 study with nearly 1,000 high school students in Türkiye, students who practised maths with an ordinary chatbot did much better during practice, but **17% worse** on the later test without AI than students who had no chatbot. A version designed as a tutor, giving hints instead of answers, avoided most of this harm.

§ 8.2 Creativity: a partner, not a replacement

AI is good at generating lots of ideas quickly: names for a project, plot twists, slogans, colour palettes, melodies. Use it to get unstuck, then choose, combine and change the ideas yourself. Your own experience, humour and voice are what make the work interesting. When you share AI-made content, say so.

§ 8.3 Accessibility: removing barriers

For many people, AI is not a shortcut but a door. Text-to-speech reads aloud for people with dyslexia or low vision; live captions help people who are deaf or hard of hearing; tools that simplify text or break tasks into steps help with concentration. Chapter 18 explores this in detail.

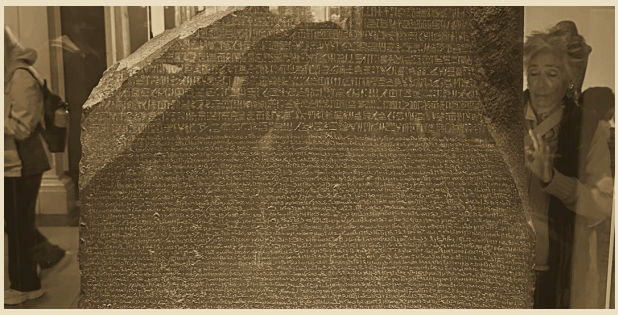


FIG. 8.1 *The Rosetta Stone (196 BC) carries the same text in three scripts. It helped scholars decode Egyptian hieroglyphs. Modern translation AI also learns from millions of "parallel texts" in two languages.*

§ 8.4 Everyday life

- Translate a menu, sign or form by pointing your camera at it.
- Plan a trip on a budget, or a week of cheap meals from what is in the fridge.
- Practise a job interview or write a first CV draft (then make it truly yours).
- Identify a plant, bird or insect from a photo, like a digital field guide.
- Get help writing a polite but firm email.

FROM THE EXCHANGE

In the workshop on AI for studying, groups tested chat tutors, language tools, summarising tools, quiz makers and planners on real school subjects. The workshop ended with a reflection on using AI as a **support** for learning rather than a replacement, and on the risk of relying on it too much.

✎ TRY THIS

Paste your notes from one lesson and ask: "Quiz me with five questions, one at a time." Then do the same lesson without AI. Which way did you remember more?

CHAPTER IX

AI models explained: *small, large, open-weight and specialised*

FIELD 09 NOTE

> *An app like ChatGPT or Le Chat is only the shop window. The engine inside is the* ***model****. Different models are like different species: some are huge and powerful, some are small and fast, some are open for anyone to download, and some are trained for one special job.*

§ 9.1 Size: counting parameters

A model's size is measured in **parameters**, the adjustable numbers inside the network (chapter 2). "7B" means 7 billion parameters. Bigger models usually know more and reason better, but need more memory, more electricity and more expensive computers. Small models are faster and cheaper, and some can run on a phone or laptop.

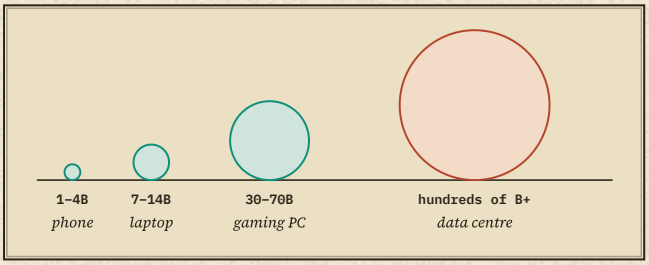


FIG. 9.1 *Model sizes and where they can "live". Exact numbers for the biggest commercial models are usually kept secret.*

§ 9.2 Open-weight versus closed

Closed models (such as those behind ChatGPT, Gemini and Claude) can only be used through the company's app or online service. **Open-weight models** (such as Llama by Meta, Gemma by Google, Mistral, Qwen and DeepSeek) can be downloaded and run by anyone, on their own computer, within the licence rules. A few projects, such as OLMo, are **fully open source**: they also publish the training data and code, so researchers can see exactly how the model was made.

Open models give privacy, independence and room to experiment; closed models are often the most powerful and easiest to use. Both can be misused or biased.

§ 9.3 A field key to model species

LARGE / FRONTIER	The most capable general models. *Habitat:* huge data centres. *Strength:* broad knowledge, complex tasks. *Weakness:* cost, energy, your data leaves your device.
SMALL	*Habitat:* phones, laptops, cars. *Strength:* fast, private, works offline. *Weakness:* knows less, makes more mistakes.
REASONING	"Think" step by step before answering. *Strength:* maths, logic, planning. *Weakness:* slower and more expensive.
MULTIMODAL	Understand and create text, images and sound together.
SPECIALISED	Trained for one job: translation, speech recognition, medical images, code, weather forecasts, protein structures.

§ 9.4 Three words you will often see

Context window: how much text a model can "keep in mind" at once, from a few pages to several books. **Knowledge cut-off:** the date its training data ends; it knows nothing newer unless it can search. **Quantisation:** compressing a model's numbers so it needs less memory, a bit like saving a photo as a smaller JPEG. This is how big models squeeze onto laptops.

FIELD NOTE

Bigger is not always better. A small model trained for one task, like grammar in one language, can beat a giant general model at it, using a fraction of the energy.

✎ TRY THIS

Ask your favourite assistant: "Which model are you, what is your knowledge cut-off, and can you search the web?" Then test it with a question about this week's news.

PART

III

Risks, ethics and *responsibility*

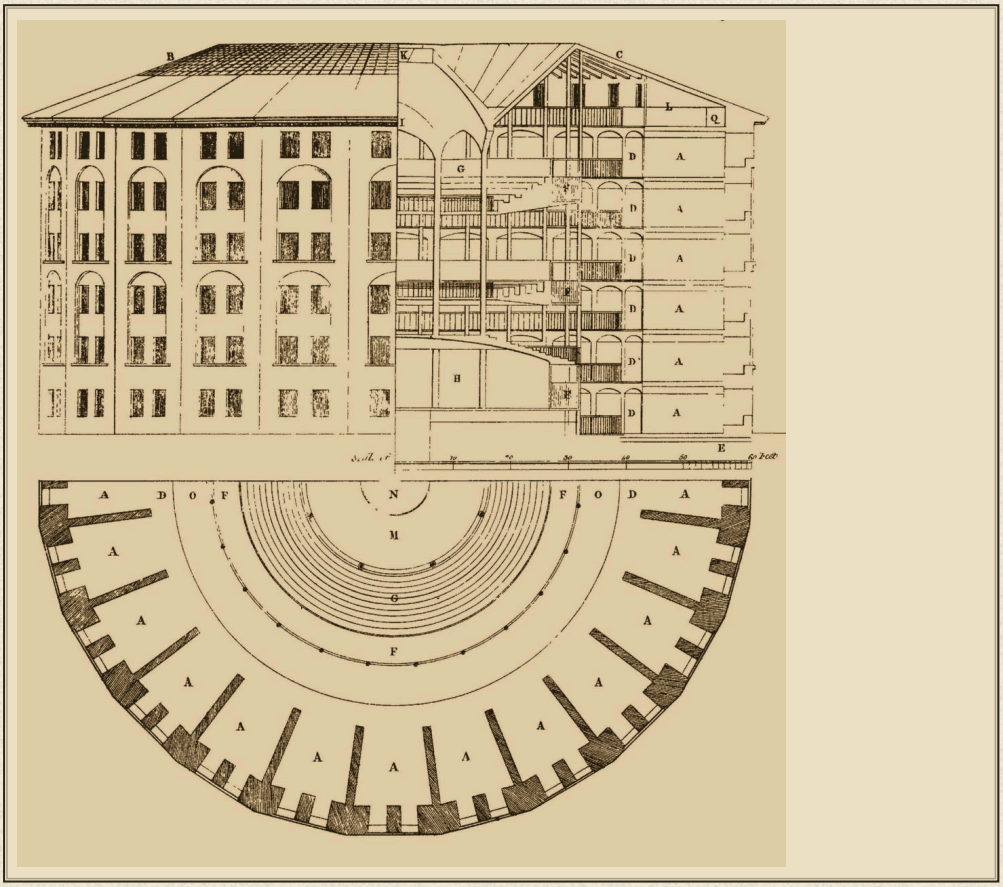


PLATE III *Jeremy Bentham's "Panopticon" (1791): a round prison where one guard in the centre can watch every cell, while prisoners never know if they are being watched. It has become a symbol of surveillance in the digital age.*

10 How to spot AI-generated content *17*

11 Myths and facts about AI *18*

12 Privacy, personal data and digital footprints *19*

13 Bias, misinformation and deepfakes *20*

14 Mental wellbeing, social media and algorithms *21*

15 Responsible AI use at school *22*

"Technology is neither good nor bad; nor is it neutral."

MELVIN KRANZBERG, HISTORIAN OF TECHNOLOGY, 1986

CHAPTER X

How to spot *AI-generated content*

FIELD 10 NOTE

AI-made images, videos, voices and texts are getting better every month, and the old tricks ("count the fingers!") no longer work reliably. The best detective skill is not staring at pixels, but checking where something comes from.

§ 10.1 Field marks: what to look for

TYPE	POSSIBLE SIGNS (NONE IS PROOF ON ITS OWN)
Images	Garbled text on signs and labels; skin that looks too smooth or "plastic"; light and shadows from different directions; earrings, glasses or buttons that don't match; objects melting into the background; strange hands, teeth or ears.
Video	Very short clips; lips slightly out of sync; flickering edges around hair and faces; objects that appear, disappear or move against physics; nobody else filmed the same event.
Voice	Flat or oddly placed emotion; no breathing or background noise; strange pronunciation of names; pressure to act *now* (send money, share a code).
Text	Smooth but generic; every answer perfectly "balanced"; repeated structures and lists; confident facts with no sources, or sources that do not exist.



FIG. 10.1 *A Cottingley Fairies photo, England, 1917–20.*

Fakes are older than AI. Between 1917 and 1920, two cousins aged 16 and 9 photographed "fairies" in their garden. The fairies were paper cut-outs held up with hatpins, yet Arthur Conan Doyle, the author of Sherlock Holmes, publicly believed they were real. The cousins only admitted the trick in 1983. AI has not invented fakes; it has made them faster, cheaper and more convincing.

§ 10.2 Why AI text detectors fail

Many websites promise to tell whether a text was written by AI. They are **not reliable**. OpenAI shut down its own AI-text detector in 2023 because it was not accurate enough. A Stanford study (2023) found that popular detectors wrongly flagged more than half of essays written by non-native English speakers as AI-generated. A detector's score should never be the only evidence against a student.

§ 10.3 The SIFT method

Developed by digital literacy expert Mike Caulfield, SIFT works for any suspicious post, AI-made or not:

S · STOP	Notice your emotion. Shock, anger or "too good to be true" are warning signs. Don't share yet.
I · INVESTIGATE	Who posted this? A real news outlet, an expert, an anonymous account created last week?
F · FIND	Find better coverage. Do trusted sources report the same thing?
T · TRACE	Trace it back to the original: the full video, the first photo, the real quote in context.

§ 10.4 Tools that help

- **Reverse image search:** Google Lens, TinEye or Bing show where else a picture appeared, and when.
- **Content Credentials** (a small "cr" icon): some cameras and AI tools attach a record of how an image was made. Check it at contentcredentials.org/verify.
- **Invisible watermarks:** Google marks images from its tools with SynthID; other companies use similar systems. But watermarks can be lost or removed.
- **InVID-WeVerify:** a free browser plug-in used by journalists to check videos.
- **Fact-checkers:** Faktograf (Croatia), Polígrafo (Portugal), Factual.ro (Romania), Raskrikavanje (Serbia), plus the European network EDMO.

⚠ VOICE SCAMS

If "your brother" calls in panic asking for money, hang up and call him back on the number you know. Agree on a family code word that only you know.

✎ TRY THIS

Play "Which face is real?" (whichfaceisreal.com). Score 20 rounds, read the tips, then play again.

CHAPTER XI

Myths and facts *about AI*

FIELD 11 NOTE

Films, adverts and social media have filled our heads with ideas about AI: killer robots, all-knowing machines, digital friends who truly understand us. Some ideas are exaggerated, some are wrong, and some hide the real problems. Here are the most common ones, checked.

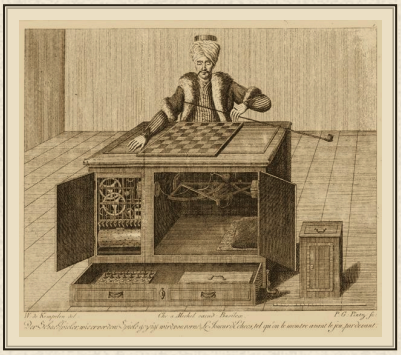


FIG. 11.1 *The "Mechanical Turk", a chess-playing machine built in 1770, engraving from 1783.*

A fake "AI" from 1770

The Turk amazed audiences in Europe and America for 84 years and reportedly played against Napoleon. Its secret: a human chess master hidden inside the cabinet. Amazon later named its online micro-work platform "Mechanical Turk" as a joke about the people behind the "machine".

Today, "AI" still often depends on hidden human work. Thousands of people around the world label images, rate chatbot answers and filter violent content, often for low pay. In 2023, *TIME* reported that workers in Kenya making a chatbot safer earned less than 2 US dollars an hour.

Nº	MYTH	FACT
01	*AI is conscious and has feelings.*	There is no scientific evidence that today's AI feels anything. It produces text that *sounds* emotional because it learned from human writing.
02	*Computers are neutral, so AI is objective.*	AI learns from human data and human choices, including our prejudices (chapter 13).
03	*AI knows everything and always searches the internet.*	Most models work from training data with a cut-off date and can invent facts. Some can search, but may still misread sources.
04	*AI learns from me while we chat.*	Not in real time. But companies may store your chats and use them to train future models.
05	*AI detectors can always catch AI-written homework.*	Detectors make many mistakes and often accuse innocent students (chapter 10).
06	*AI will take all jobs in a few years.*	AI changes tasks inside jobs much more often than it removes whole jobs, and creates new ones too (chapter 19).
07	*You have to be good at maths or coding to understand AI.*	Anyone can understand the basics and use AI critically. That is the point of this book.
08	*AI is just hype and will disappear.*	Some promises are hype, but AI already saves time in medicine, science, translation and accessibility.
09	*The biggest danger is robots taking over the world.*	Today's real problems are less cinematic: misinformation, bias, privacy, scams, energy use and power in the hands of a few companies.
10	*If a product says "AI-powered", it must be smart.*	"AI-washing" is a marketing trick: some products call simple rules or even human work "AI".

✎ TRY THIS

Ask a chatbot: "Do you have feelings?" Then: "Pretend you have feelings and describe them." What does the difference tell you about myth 1?

CHAPTER XII

Privacy, personal data *and digital footprints*

FIELD 12 NOTE

Every like, search, location and voice message leaves a footprint. AI makes these footprints more valuable, because it can combine millions of small traces and guess things about you that you never said.

§ 12.1 What counts as personal data?

Under the EU's General Data Protection Regulation (**GDPR**), personal data is any information about a person who can be identified: your name, photo, voice, email, phone number, location, school, IP address or online username. Some data is especially sensitive and has extra protection: health, religion, political opinions, sexual orientation, biometric data such as your face or fingerprint.

§ 12.2 Your two footprints

Your **active footprint** is what you share on purpose: posts, comments, photos. Your **passive footprint** is collected without you noticing: where your phone was, how long you looked at a video, what time you go to sleep, which ads you scrolled past. Apps and data brokers can combine these traces into a detailed profile.

FIELD NOTE

As early as 2013, Cambridge researchers showed that Facebook "likes" alone could predict a person's political views, religion and sexual orientation with surprising accuracy. Modern AI is far better at this kind of guessing.

§ 12.3 What happens to your chats?

When you type into an online AI tool, your words travel to the company's servers. Depending on the service and your settings, they may be **stored**, **read by human reviewers** for quality and safety checks, and **used to train** future models. Many apps have a "temporary chat" mode and a setting to opt out of training. Find them and use them.



FIG. 12.1 *From a 1905 guide to fingerprint identification. Today you can be identified not only by your fingerprints but by your face, your voice and even the rhythm of your typing.*

§ 12.4 Faces, voices and other people

Uploading a photo of a friend to an AI app, or cloning a classmate's voice "as a joke", means sharing *their* personal data without permission. It can hurt them and in many cases it is against the law. Ask before you share anyone's photo, voice or messages.

§ 12.5 Your rights

In the EU (Croatia, Portugal, Romania) the GDPR gives you the right to **know** what data a company has about you, to get a **copy**, to **correct** it, to have it **deleted** ("right to be forgotten") and to **object** to some uses. Serbia's Law on Personal Data Protection (2018) gives very similar rights. If a company ignores you, you can complain to your national authority: AZOP (Croatia), CNPD (Portugal), ANSPDCP (Romania) or the Commissioner for Information of Public Importance and Personal Data Protection (Serbia).

§ 12.6 A privacy check-up

- Turn off "use my chats to improve the model" where you can.
- Never type passwords, ID numbers, health details or addresses into a chatbot.
- Check which apps can see your location, microphone, camera and contacts.
- Use strong, different passwords and two-factor authentication.
- Think twice before posting photos of your face, school uniform or home.
- Google yourself once a year and see what others can find.

FROM THE EXCHANGE

Privacy came up again and again in Zadar: in the ethics workshop, where one group discussed how AI systems collect and use personal data, and in the youth manifesto, which had its own section on how personal data should be handled.

CHAPTER XIII

Bias, discrimination, misinformation *and deepfakes*

FIELD 13 NOTE

> *AI can be unfair without anyone meaning it to be, and it can make lies look real. These are among the most serious risks of AI today, and they affect real people's lives, jobs and elections.*

§ 13.1 How bias gets in

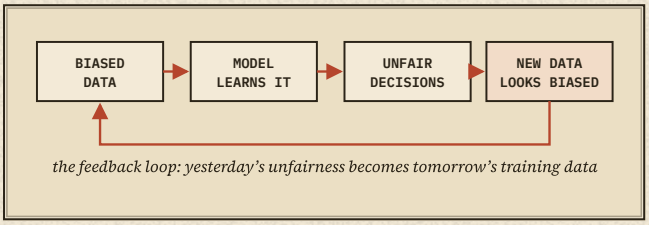


FIG. 13.1 *How bias can grow in a loop.*

Bias can enter through **data** (some groups are missing or shown in stereotypes), through **labels** (people's judgements when marking data), through **design** (what the system is told to optimise) and through **use** (using a tool in a situation it was not built for).

§ 13.2 Real cases

- **Faces:** the 2018 *Gender Shades* study found commercial face-analysis systems misclassified darker-skinned women up to 34.7% of the time, but lighter-skinned men only 0.8%.
- **Jobs:** Amazon stopped using an experimental CV-screening tool after finding it downgraded CVs that included the word "women's", because it had learned from ten years of mostly male hires.
- **Welfare:** in the Netherlands, a risk-scoring system that treated dual nationality as a risk factor contributed to thousands of families being wrongly accused of benefits fraud. The government resigned over the scandal in 2021.

Old mistake, new tools. Phrenology claimed to read character and intelligence from the shape of the skull. It was pseudoscience, and it was used to justify racism. AI systems that claim to read personality, honesty or "criminality" from a face repeat the same mistake.

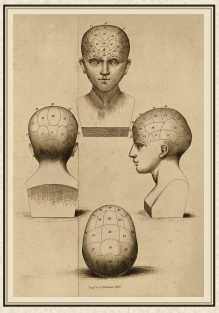


FIG. 13.2 *Phrenology head, 19th century.*

§ 13.3 Misinformation at machine speed

Misinformation is false information shared by mistake; **disinformation** is spread on purpose. AI makes both cheaper: one person can now produce thousands of fake posts, fake news sites or fake "witnesses" in many languages. There is also the **liar's dividend**: once people know fakes exist, real evidence can be dismissed as "just AI".

§ 13.4 Deepfakes

Deepfakes are realistic fake videos, images or audio of real people. In 2022 a fake video showed Ukraine's president telling soldiers to surrender. Before Slovakia's 2023 election, fake audio of a party leader spread just before voting. Most harmful of all are fake intimate images, mostly targeting women and girls, sometimes made by classmates. This is abuse, and in many countries a crime.

⚠ IF IT HAPPENS TO YOU OR A FRIEND

Don't share it further. Save evidence (screenshots, links). Report it on the platform. Tell a trusted adult or teacher, and contact the police or a helpline. It is never the victim's fault.

FROM THE EXCHANGE

The ethics workshop in Zadar split into groups on privacy, bias and fairness, responsibility when AI causes harm, misinformation and deepfakes, and the ethical choices young people make every day.

§ 13.5 What you can do

- **Ask "who is missing?"** when you look at AI results, images or recommendations.
- **Prompt for diversity** on purpose, and notice what the AI shows by default.
- **Pause before sharing** anything that makes you angry or shocked (SIFT, ch. 10).
- **Report** unfair results and fakes to the platform; they rely on reports.
- **Support transparency:** people affected by an AI decision should be told, and should be able to ask a human to review it.

✎ TRY THIS

Ask an image generator for "a scientist", "a nurse", "a criminal" and "a family", four images each. Count genders, ages and skin colours. What stereotypes appear?

CHAPTER XIV

Mental wellbeing, social media *and algorithms*

FIELD 14 NOTE

The AI you meet most often is not a chatbot. It is the recommendation algorithm that decides what appears in your feed. It is very good at its job, which is keeping your attention. Whether that is good for you is another question.

§ 14.1 How your feed decides

Recommendation systems learn from every signal you give: what you watch to the end, what you replay, where you pause, what you like, share or skip, and what similar users enjoyed. Their main goal is usually **engagement**: more time in the app means more ads seen. Showing you what is good for you is not their main goal.

§ 14.2 A slot machine in your pocket



FIG. 14.1 *A "Liberty Bell" slot machine, invented about 1898.*

Slot machines are addictive because rewards are **unpredictable**: sometimes you win, often you don't, so you keep pulling. Pull-to-refresh, endless scrolling, autoplay, streaks and likes work in a similar way. You never know if the next video will be the great one. The designer of infinite scroll, Aza Raskin, later said he regretted how it is used to keep people scrolling.

§ 14.3 Comparison and self-image

Feeds show the highlights of other people's lives, often edited by beauty filters and AI retouching that can change face shape, skin and body in real time. Comparing your real everyday life with someone else's edited best moments can damage self-esteem. Remember: *you are comparing your behind-the-scenes with their trailer.*

§ 14.4 Rabbit holes

Because algorithms show "more of what you engaged with", a few clicks on extreme dieting, conspiracy or hateful content can quickly fill a feed with it. The EU's Digital Services Act now requires very large platforms to offer at least one feed option that is **not based on profiling** you, and bans ads targeted at minors based on their profile. Look for the "Following" or non-personalised feed option.

§ 14.5 AI "friends"

Companion chatbots are always available, always nice and never tired of you. For some lonely people they can be a comfort. But they are products, often designed to keep you chatting. They are not therapists, they can give harmful advice, and they cannot replace friends who know you in real life.

§ 14.6 Take back control

- Turn off non-essential notifications.
- Use "Not interested" and unfollow accounts that make you feel bad. Your clicks train your feed.
- Set app timers, and keep the phone out of the bedroom at night.
- Try grayscale mode: a grey screen is much less tempting.
- Before opening an app, ask: "What am I here for?"

FROM THE EXCHANGE

Three workshops in Zadar looked at this topic. Through theatre scenes on comparison, digital overload, online pressure and fear of the future, posters on attention and instant gratification, and discussions on friendship, face-to-face contact and FOMO, the discussions kept coming back to one question: how can technology **support** young people's wellbeing instead of exploiting it?

NEED TO TALK?

116 111 is the free, confidential helpline for children and young people in Croatia, Portugal, Romania and Serbia. In an emergency, call your local emergency number (112 in the EU).

CHAPTER XV

Responsible AI use *at school*

FIELD 15 NOTE

AI at school is not simply "cheating" or "allowed". It depends on the task, the teacher's rules and how honest you are about it. The golden rule: AI may help you learn, but it must not learn instead of you.



FIG. 15.1 *"At School", a French picture card from around 1900 imagining the year 2000: books are fed into a machine and the knowledge flows through wires into the pupils' heads. Over a century later, learning still doesn't work like that.*

LIGHT	USUALLY FINE (CHECK YOUR TEACHER'S RULES)	TYPICAL EXAMPLES
• **Green**	Using AI to understand, practise and get ideas	Explaining a topic in a new way; quizzing yourself; brainstorming; making a study plan.
• **Amber**	Allowed only if your teacher agrees, and you say how you used it	Feedback on your draft; help with structure; translating sources; checking grammar.
• **Red**	Not acceptable	Handing in AI-written work as your own; using AI in tests; inventing sources or data; uploading classmates' work or data.

§ 15.1 Why honesty matters

Handing in AI work as your own is plagiarism. But the bigger loss is invisible: the skills you were meant to practise, such as writing, arguing and problem-solving, don't grow. Remember the study in chapter 8: students who let a chatbot give them the answers did worse on the real test.

§ 15.2 How to disclose AI use

When AI helped, add a short note at the end of your work:

EXAMPLE

"I used Microsoft Copilot to brainstorm possible topics and to check my grammar. The research, arguments and final wording are my own."

Keep your chat history and use version history in Word or Google Docs: it shows how your work grew and protects you if you are wrongly suspected.

§ 15.3 Check and cite

Never use a chatbot as the source of a fact. Find the original source (book, article, official website), check it and cite *that*. If the AI gives you references, look up every single one: they may not exist.

§ 15.4 Rules for schools, too

The EU AI Act **bans** AI systems that try to recognise students' emotions in schools (except for medical or safety reasons). AI used to grade exams, decide admissions or monitor students during tests counts as **"high-risk"** and must meet strict rules on accuracy, human oversight and transparency. UNESCO recommends that generative AI in classrooms should be used only from age 13.

FROM THE EXCHANGE

The youth manifesto written in Zadar had a section on education and learning: how AI should be used responsibly in schools. Another section demanded that young people be included in decisions about AI.

✎ TRY THIS

With your class, sort ten real tasks (essay, lab report, presentation…) into green, amber and red. Compare with your teacher.

PART

IV

AI, independence *and the future*



PLATE IV *The "Maschinenmensch" (machine-human) from Fritz Lang's film Metropolis (1927), one of the first robots in cinema. For a hundred years, films have imagined AI as a robot that looks like us. Real AI turned out to be stranger, and much less visible.*

16 Local AI models: AI on your own device 24

17 Cloud versus local AI 25

18 AI and accessibility 26

19 AI and the future of work 27

20 What kind of AI future do we want? 28

"The best way to predict the future is to invent it."

ALAN KAY, COMPUTER SCIENTIST, 1971

CHAPTER XVI

Local AI models: *AI on your own device*

FIELD 16 NOTE

> *You don't always need a giant data centre to use AI. Many open-weight models can run directly on a laptop or even a phone: offline, private and free. This chapter shows how.*

§ 16.1 What "local" means

With an online assistant, your question travels over the internet to a company's servers, and the answer travels back. With a **local model**, you download the model file once and it runs on your own processor and memory. Your questions and documents **never leave your device**. Once it is downloaded, it even works in airplane mode.

§ 16.2 What you need

The most important thing is memory (RAM). A rule of thumb: a compressed (4-bit) model needs roughly **0.6 GB of memory per billion parameters**, plus some extra.

YOUR COMPUTER	COMFORTABLE MODEL SIZE
8 GB RAM laptop	1–4B: simple chat, summaries, grammar
16 GB RAM laptop	7–14B: good everyday assistant
32 GB+ or gaming graphics card	20–35B: strong, close to older online models
Modern phone	1–4B, with a suitable app

Model files are big (usually 1 to 20 GB), so download them over Wi-Fi. Computers with a recent graphics card or an Apple M-series chip run them much faster.

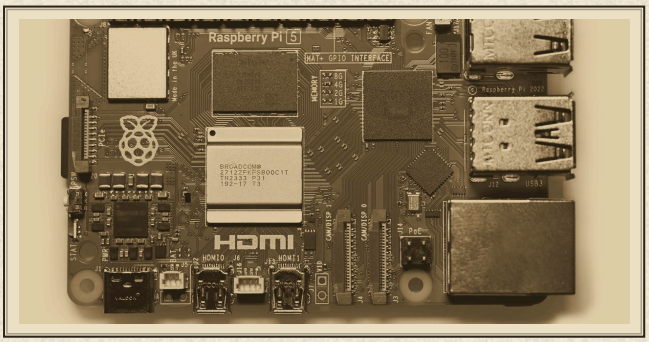


FIG. 16.1 *A Raspberry Pi 5, a credit-card-sized computer that costs less than 100 €. With 8 GB of memory it can run small language models, slowly but completely offline.*

§ 16.3 Your first local model in four steps

1. **Install a free app:** *LM Studio* (lmstudio.ai) has a simple chat window; *Ollama* (ollama.com) works from the command line and with many other apps. *Jan* and *GPT4All* are other options.
2. **Choose a small model** from a well-known family: Gemma (Google), Llama (Meta), Qwen (Alibaba), Phi (Microsoft) or Mistral. Start with 3–4B parameters.
3. **Download** it inside the app. In Ollama, one line is enough:
   `ollama run gemma3:4b`
4. **Chat.** Switch off Wi-Fi and ask again. It still works!

§ 16.4 What local AI is good for

- Private things: your diary, personal notes, sensitive school projects.
- Travelling or places with poor internet.
- Learning how AI really works, comparing models and testing their limits.
- No subscription, no account and no ads.

§ 16.5 Limits

Small models know less and make more mistakes than the big online ones. They are slower on older computers, drain laptop batteries and usually cannot search the web. You are also responsible for where you download them from: use official sources such as the model makers' pages or Hugging Face.

FIELD NOTE

Hugging Face, a platform for sharing AI models, hosts well over a million of them, from tiny models for one language to large general assistants. It is like a public library of AI.

✎ TRY THIS

Ask a local model and an online assistant the same three questions: one about general knowledge, one about today's news and one about your town. Where does the local model struggle, and why?

CHAPTER XVII

Cloud versus local AI: *privacy, costs and limitations*

FIELD 17 NOTE

> *"The cloud" sounds light and fluffy. In reality it is someone else's computer: huge buildings full of servers that use electricity, water and land. Choosing cloud or local AI means balancing power against privacy.*

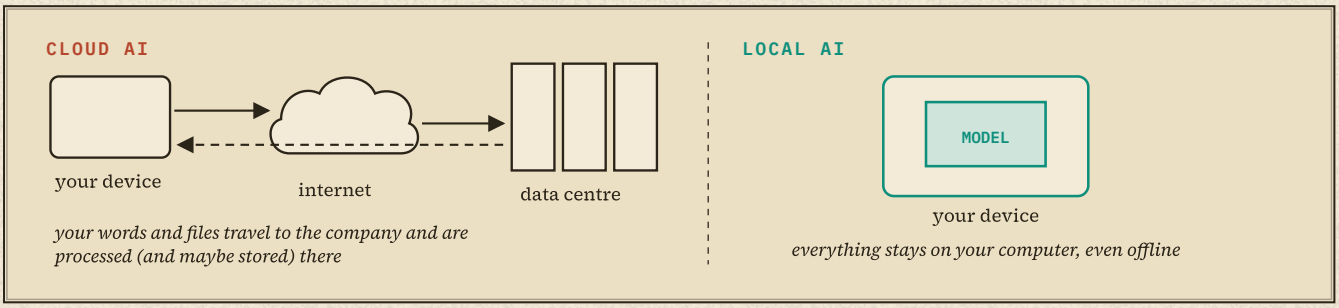


FIG. 17.1 *Where your data goes: cloud AI versus local AI.*

QUESTION	CLOUD AI (CHATGPT, GEMINI, COPILOT…)	LOCAL AI (LM STUDIO, OLLAMA…)
Privacy	Data is sent to and processed by a company; may be stored, reviewed or used for training.	Data stays on your device. Best for anything personal.
Power	The strongest models, huge knowledge, web search, images and video.	Smaller models; good for many tasks, weaker at complex ones.
Cost	Free tiers with limits; around 20 € a month for more. "Free" may mean paying with data.	Free software and models; you pay with your own hardware and electricity.
Internet	Required.	Not needed after download.
Ease of use	Very easy: open the app and type.	Needs some setup and a good enough computer.
Control	The company can change the model, the price or the rules at any time.	You choose the model and version, and it does not change unless you change it.

§ 17.1 The hidden costs of the cloud



FIG. 17.2 *Server racks in the CERN computer centre, Geneva. AI data centres are full of rows like these.*

The International Energy Agency estimates that data centres used about 1.5% of the world's electricity in 2024, and expects this to roughly double by 2030, mainly because of AI. Many data centres also use large amounts of water for cooling. A single chat question uses little energy; billions of them, plus training ever-larger models, add up.

§ 17.2 Independence and "digital sovereignty"

Most of the world's leading AI models are made by a handful of companies in the USA and China. If everyone depends on them, those companies have great power over information, prices and rules. That is why Europe invests in its own AI (such as Mistral in France and EU supercomputers), and why open and local models matter for schools, media and public services.

§ 17.3 The best of both

Many new phones and computers already mix both: simple tasks run on the device, while harder ones are sent to the cloud. A sensible personal rule: **use local or no AI for private things; use cloud AI for general, non-personal tasks**, with training on your chats switched off.

✎ TRY THIS

Write down five things you used AI for this week. Next to each, write "cloud OK" or "should be local". Were you surprised?

CHAPTER XVIII

AI and accessibility: *technology for everyone*

FIELD 18 NOTE

> *For about one in six people in the world who live with a significant disability, and for many others with temporary or invisible needs, AI can be much more than a gadget. It can be a way to read, hear, speak, move and take part more independently.*



FIG. 18.1 *Left: Helen Keller (left), who was deafblind, aged 8 with her teacher Anne Sullivan in 1888. They communicated by spelling words into each other's hands. Right: physicist Stephen Hawking, who spoke through a computer. From 2014 his system used predictive-text AI to guess his next words.*

§ 18.1 Seeing

Be My Eyes connects blind and low-vision people with volunteers, and its AI feature describes photos in detail. **Seeing AI** (Microsoft, free) reads text aloud, recognises money and describes scenes. Screen readers now use AI to describe images that have no text description.

§ 18.2 Hearing

Live captions on phones, laptops and video calls turn speech into subtitles in real time. Apps such as **Live Transcribe** show conversations as text and can alert users to sounds like a doorbell, a crying baby or a fire alarm.

§ 18.3 Speaking

People who are losing their voice through illness can create a synthetic copy of it in advance (for example Apple's *Personal Voice*). Speech recognition trained on unusual speech, such as Google's *Project Relate*, helps people whose speech others find hard to understand.

§ 18.4 Learning and thinking

Text-to-speech and tools like **Immersive Reader** help people with dyslexia; AI can rewrite complex texts in easy-to-read language; task-splitting tools and reminders support people with ADHD. These help many other learners too.

§ 18.5 Moving and language

Voice control and eye tracking let people with limited movement use phones and smart homes. Live translation helps migrants and refugees deal with doctors, schools and offices in a new language.

§ 18.6 Not a magic fix

- **Mistakes matter more:** a wrong caption in a lesson or a misread medicine label can cause real harm.
- **Bias:** speech recognition often works worse for strong accents, older people or speech disabilities.
- **Privacy:** describing your surroundings means sending camera images to a company.
- **Cost:** accessibility should not depend on an expensive subscription.
- **Design with, not for:** the disability movement's motto "Nothing about us without us" applies to AI too.

FROM THE EXCHANGE

In the accessibility and inclusion workshop, groups tested text-to-speech and speech-to-text, automatic translation and subtitles, tools for learning difficulties, image description and communication aids, and asked who benefits from each. The reflection focused on why technology must be designed in an inclusive, human-centred way.

✎ TRY THIS

Turn on live captions on your phone or laptop and watch a video with the sound off. Then use a screen reader for five minutes. What works well, and what is frustrating?

CHAPTER XIX

AI and the future of work: *jobs, skills and opportunities*

FIELD 19 NOTE

> *Will AI take my future job? The honest answer: AI will change almost every job a little, some jobs a lot, and create jobs that do not exist yet. History helps us understand what might happen, and what we can do about it.*

§ 19.1 Lessons from history

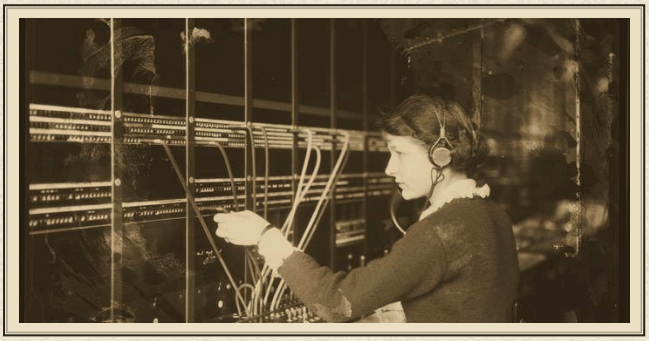


FIG. 19.1 *Telephone switchboard operators in Seattle, about 1922. Automatic exchanges later replaced this huge job category, yet new kinds of work grew elsewhere.*

Between 1811 and 1816, English textile workers called **Luddites** smashed new weaving machines. They were not simply "against technology": they protested because factory owners used the machines to cut wages and replace skilled workers with cheaper labour. Today "Luddite" is used as an insult, but their question is still important: *who benefits from new technology?*

Later, machines replaced telephone operators, typists and many factory jobs. Each time, some people lost their work while new jobs appeared in other places. Overall employment grew, but the change was hard for the people caught in the middle.

§ 19.2 What the forecasts say

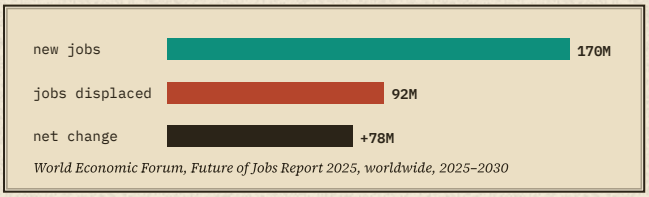


FIG. 19.2 *Forecast job changes by 2030 caused by technology, the green transition, demographics and the economy together.*

The World Economic Forum expects that by 2030 about 39% of the key skills workers need will change. The International Monetary Fund estimated in 2024 that about 40% of jobs worldwide, and around 60% in rich countries, are "exposed" to AI. For some, AI will be a helpful tool; for others it may replace important tasks.

§ 19.3 Think in tasks, not jobs

AI is best at tasks that are digital, repetitive and text- or pattern-based: drafting standard texts, translating, sorting documents, basic coding, answering routine customer questions. It is weak at tasks that need hands-on skill in messy real places, deep human contact, care, trust and responsibility: nurses, electricians, teachers, therapists, chefs. New roles are growing too: AI trainers, data specialists, AI auditors, cybersecurity experts, robotics technicians and green-energy jobs.

§ 19.4 Skills that grow in value

TECH SKILLS	AI and data literacy, cybersecurity, basic coding, using AI tools well in your field.
THINKING SKILLS	Critical and creative thinking, analysing problems, judging information.
HUMAN SKILLS	Communication, teamwork, empathy, leadership, working across cultures and languages.
SELF SKILLS	Curiosity, resilience, flexibility and lifelong learning.

FROM THE EXCHANGE

Many of these "human skills" are exactly what an international youth exchange trains: working in mixed teams, presenting in a foreign language, solving problems with people from other cultures. Write them in your Youthpass and your CV!

✎ TRY THIS

Pick a job you might want. List its ten main tasks. Mark which AI could do, which it could help with, and which need a human. What does that job look like in 2035?

CHAPTER XX

What kind of AI future *do we want?*

FIELD 20 NOTE

> *The future of AI is not a train on fixed rails. It is shaped by decisions made by companies, governments, scientists and ordinary people every day. So the real question is not only "What will happen?" but "What do we want to happen?"*

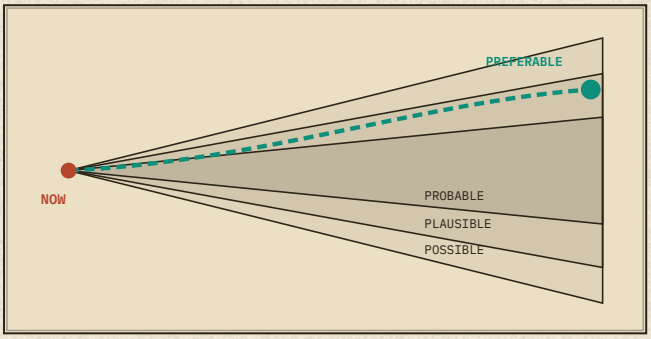


FIG. 20.1 *The "futures cone" used by futures researchers. Many futures are possible; fewer are plausible or probable. The preferable future, the one we want, has to be chosen and worked for.*

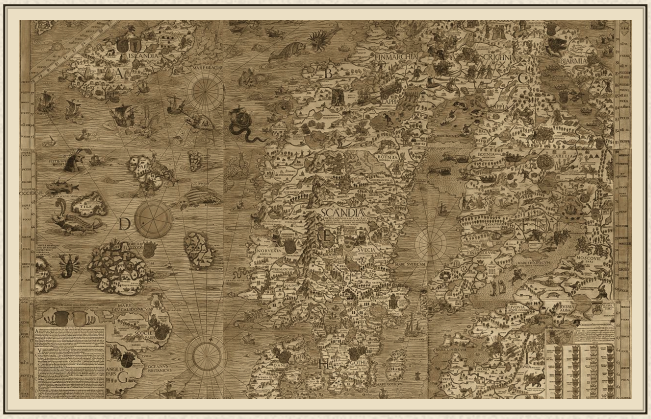


FIG. 20.2 *Detail of the Carta Marina (1539) by Olaus Magnus, a map of northern Europe full of sea monsters where knowledge ran out. Our map of the AI future is still being drawn.*

§ 20.1 Two futures, side by side

AREA	AI AS A FRIEND	AI AS A FOE
Education	Every student has a patient tutor that adapts to them; teachers have more time for people.	Students stop thinking for themselves; schools watch and score every move.
Work	Boring, dangerous tasks are automated; people work less and better; benefits are shared.	A few companies collect the profits; many workers lose income and security.
Social life	Translation and tools connect people across borders and help shy or isolated people.	People retreat into feeds and AI companions; loneliness and comparison grow.
Democracy	Citizens understand complex issues and take part more easily.	Deepfakes and targeted propaganda make it hard to know what is true.
Health	Earlier diagnosis, new medicines, better support for mental health.	Sensitive health data is sold; apps replace real care.

§ 20.2 Who decides?

Laws matter. The EU AI Act (2024) bans some uses, such as social scoring and most real-time face recognition by police in public spaces, and sets strict rules for high-risk uses. In 2024 the Council of Europe also opened the first international treaty on AI and human rights. But laws follow what people demand. Every time you choose a tool, share (or don't share) a fake, or speak up at school, you are voting for one future or another.

§ 20.3 Questions to discuss

- Which decisions should never be left to AI alone?
- Should AI-made content always be labelled?
- Who should profit when AI replaces human work?
- What would make you trust an AI system?

FROM THE EXCHANGE

In Zadar, groups acted out possible futures of AI in education, work, social life, democracy and health. Later, everyone wrote a joint youth manifesto on responsible AI, including a demand that **young people must be included** in decisions about AI.

PART

V

Practical *guide*

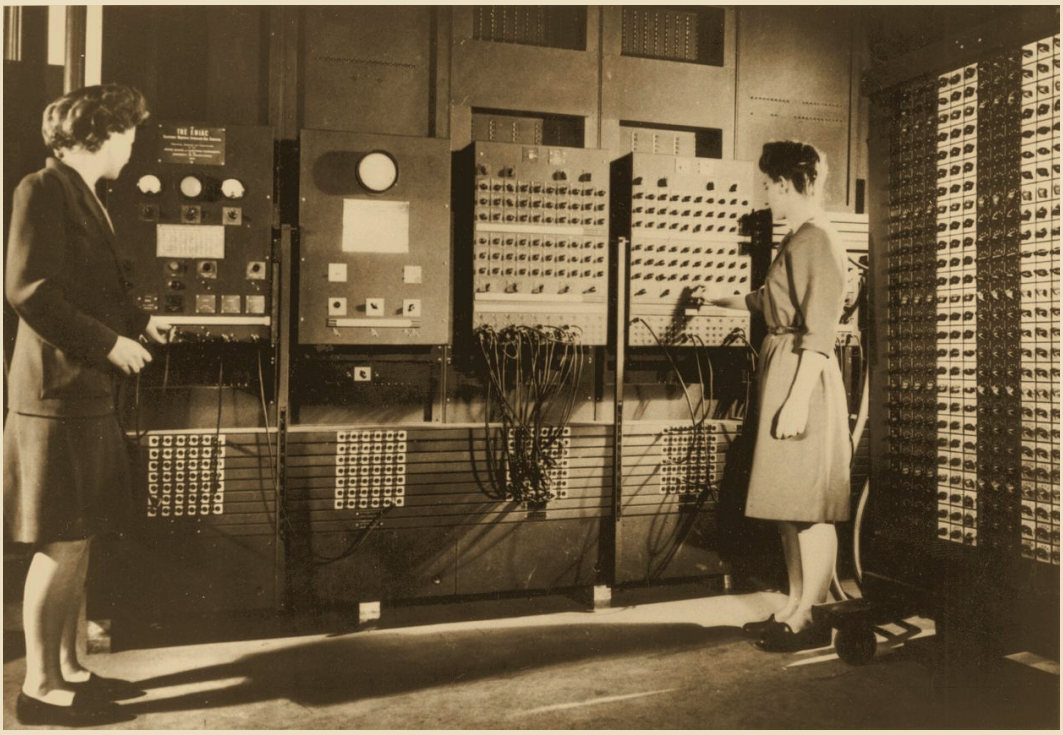


PLATE V *Programmers at work on ENIAC, one of the first electronic general-purpose computers, in the 1940s. Its first programmers, six young women, had no manuals or programming languages: they learned the machine hands-on, from its wiring diagrams. The best way to understand technology is still to try it yourself.*

21 Ten rules for using AI responsibly *30*

22 Practical activities and experiments *31*

23 Glossary of essential AI terms *32*

24 Useful tools, resources and further learning *34*

"Extraordinary claims require extraordinary evidence."

CARL SAGAN, ASTRONOMER, 1980

CHAPTER XXI

Ten rules for using *AI responsibly*

FIELD 21 NOTE

Everything in this guide, in ten rules. Cut them out, put them above your desk, or share them with your class.

I — You are the pilot, AI is the co-pilot.
Use AI to help you, but stay in charge. You decide what to use, and you are responsible for what you hand in, post or send.

II — Check before you trust.
AI can sound sure and still be wrong. Verify important facts with reliable sources, and look up every reference it gives you.

III — Protect your data, and other people's.
No passwords, ID numbers, health details or addresses in prompts. Never upload someone's photo, voice or messages without permission.

IV — Be honest about AI use.
Follow your school's rules and say when and how AI helped you. Hiding it damages trust; disclosing it protects you.

V — Learn with AI, not instead of learning.
Ask it to explain, quiz and give feedback, not to do the thinking for you. If it feels too easy, you are probably not learning.

VI — Never create or share fakes that hurt people.
Deepfakes, fake screenshots and fake nudes can destroy lives and are often crimes. If you receive one, don't forward it: report it.

VII — Question the output.
Ask yourself: Whose view is missing? Is this a stereotype? Who benefits if I believe this? AI copies the biases in its data.

VIII — Protect your attention and wellbeing.
Feeds and apps are designed to keep you hooked. Set limits, turn off notifications and spend time with real people. Ask for help if you need it.

IX — Choose the right tool, or no tool.
Use local or no AI for private things. Remember that AI uses energy, and that sometimes a book, a teacher or a friend is the better choice.

X — Stay curious and speak up.
Keep learning how AI works, and take part in decisions about it at school and in your community. The future of AI is also yours.



WHY RULE I COMES FIRST

In 1843 Ada Lovelace, often called the first computer programmer, wrote about Charles Babbage's Analytical Engine: *"The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform."*

Today's AI can surprise us far more than Babbage's machine, but her point still matters: humans design these systems, humans choose how to use them, and humans are responsible for the results.

FIG. 21.1 *Ada Lovelace, portrait by Alfred Edward Chalon, about 1840.*

MY AI PLEDGE

I have read the ten rules and I will try to use AI as a friend, not a foe.

NAME _____ DATE _____

CHAPTER XXII

Practical activities *and experiments*

FIELD 22 NOTE

Eight activities for a class, a youth club or a group of friends. Several are based on the non-formal methods we used in Zadar. Most need only phones or laptops and an internet connection.

A Teach a machine (and break it)

30 MIN · LAPTOP WITH CAMERA

Open **Teachable Machine** (teachablemachine.withgoogle.com). Train an image model with two classes, e.g. "thumbs up" and "thumbs down", using only one person's hand in one room. Then test it with other people, other light, other backgrounds.

Discuss: When did it fail? What does this show about training data and bias?

B Hallucination hunt

20 MIN · ANY CHATBOT

Each person asks a chatbot five questions about a topic they know very well (their town, school, sport, a local tradition) and asks for sources. Check every fact and every source.

Discuss: How many mistakes? Did the AI sound less confident when it was wrong?

C Real or fake? Detective lab

30 MIN · PHONES

The facilitator collects ten images or short clips, some real, some AI-made. In pairs, participants use the SIFT method and reverse image search (ch. 10) and must explain their verdict, not just guess.

Discuss: Which clues were reliable? Which tricked you?

D Prompt duel

25 MIN · ONE CHATBOT PER TEAM

Teams get the same task (e.g. "a one-page guide to our town for exchange visitors"). Ten minutes to write the best prompt, using Role, Task, Context and Format (ch. 7). The group votes on the results.

Discuss: Which details in the prompt made the biggest difference?

E Feed audit

10 MIN A DAY · ONE WEEK

Each day, note the first 20 posts in your main feed: topic, who posted it, and how it made you feel (😊 😐 😕). At the end of the week, look for patterns, then use "Not interested" and unfollow to retrain your feed.

Discuss: Who is really choosing what you see?

F Data diary

ONE DAY · NOTEBOOK

For 24 hours, write down every moment you leave a data trace: unlocking your phone, paying by card, a search, a like, a voice message, a location check-in. Mark the ones AI might use.

Discuss: What could a company guess about you from one day?

G Friend or foe? Human line

30 MIN · OPEN SPACE

Mark a line on the floor: one end "Friend", the other "Foe". Read statements such as "AI should grade our tests" or "AI companions are good for lonely people". Everyone stands where they agree, then volunteers explain their position. People may move after hearing others.

Discuss: Did anyone change your mind? Why?

H Write your own AI manifesto

60 MIN · PAPER, MARKERS

As in the final workshop in Zadar, split into groups: fairness, mental health, privacy, education, youth participation. Each group writes three principles or demands. Combine them into one class manifesto and share it with your school.

Discuss: Who needs to hear your manifesto?

CHAPTER XXIII

Glossary of *essential AI terms*



*A pocket dictionary of the words you will meet when people talk about AI. Words in **bold** inside a definition have their own entry.*

A

AI agent
An AI system that can plan and carry out several steps on its own, such as searching the web, filling in forms or booking something, to reach a goal you give it.

AGI (artificial general intelligence)
A hypothetical AI that could learn and do any intellectual task a human can. It does not exist today.

Algorithm
A step-by-step set of instructions for solving a problem, like a recipe for a computer.

Alignment
Research on making AI systems behave according to human values and intentions.

API
A way for one program to talk to another. Many apps use AI models through an API.

Artificial intelligence (AI)
Computer systems that do tasks normally linked to human intelligence, such as understanding language, recognising images or making predictions.

B

Benchmark
A standard test used to compare how well different AI models perform.

Bias
A systematic unfairness in an AI system's output, often copied from its **training data**.

Bot
A program that does automated tasks online, such as posting messages or answering questions.

C

Chatbot
A program you can have a conversation with in text or voice, usually powered by a **large language model**.

Cloud AI
AI that runs on a company's servers, which you reach over the internet.

Content Credentials
Information attached to a file that records how it was created or edited, including whether AI was used.

Context window
How much text (in **tokens**) a model can take into account at once.

Copyright
The legal right of creators to control how their work is copied and used. How it applies to AI training and AI output is still being decided in courts.

D

Data
Information such as text, images, numbers or sounds, used to train and run AI.

Dataset
A collection of data prepared for training or testing a model.

Deep learning
Machine learning with **neural networks** that have many layers.

Deepfake
A realistic fake image, video or audio recording of a real person, made with AI.

Diffusion model
A type of generative model that creates images or video by gradually removing noise.

Disinformation
False information spread on purpose to deceive. Compare **misinformation**.

E

Engagement
How much people interact with content: watching, liking, sharing, commenting. Many **recommendation algorithms** try to maximise it.

EU AI Act
The European Union's law on AI (2024). It bans some uses, sets strict rules for high-risk systems and requires transparency, for example labelling deepfakes.

F

Fine-tuning
Extra training of an existing model on a smaller, specialised dataset to adapt it to a task.

Foundation model
A large model trained on broad data that can be adapted to many tasks.

G

GDPR
The EU's General Data Protection Regulation, which protects personal data and gives people rights over it.

Generative AI
AI that creates new content: text, images, audio, video or code.

GPU
Graphics processing unit: a computer chip, first made for video games, that is very good at the maths AI needs.

Guardrails
Rules and filters that stop an AI system from producing harmful or unwanted output.

H

Hallucination
When an AI produces false or invented information and presents it as true.

Human in the loop
A setup where a person checks or approves an AI system's decisions.

I – J

Inference
Using a trained model to produce an answer. Every chat message is an inference.

Jailbreak
A trick to make an AI ignore its **guardrails**. Often against the terms of use.

L

Large language model (LLM)
A very large neural network trained on huge amounts of text to predict the next **token**. The engine behind chatbots.

Local model
A model that runs on your own device instead of in the cloud.

M

Machine learning
A way of building AI in which computers learn patterns from examples instead of following hand-written rules.

Misinformation
False information shared by mistake, without the intention to deceive.

Model
The result of training: a file of **parameters** that turns inputs into outputs.

Multimodal
Able to work with several kinds of data, such as text, images and sound, together.

N – O

Neural network
A computing system made of layers of connected "neurons", loosely inspired by the brain.

Open source
Software whose code, and for AI ideally also data and training process, is public and free to use and change.

Open-weight model
A model whose **weights** can be downloaded and run by anyone, even if its training data is not public.

Overfitting
When a model memorises its training examples so closely that it fails on new ones.

P

Parameter
One of the adjustable numbers inside a model. Large models have billions.

Personal data
Any information about an identifiable person: name, photo, voice, location and much more.

Prompt
The instruction or question you give an AI system.

Prompt injection
Hidden instructions in a web page or file that try to take control of an AI that reads them.

Q – R

Quantisation
Compressing a model's numbers so it needs less memory and runs on smaller devices.

RAG (retrieval-augmented generation)
A technique where an AI first looks up relevant documents and then writes its answer based on them.

Reasoning model
A model trained to work through a problem step by step before answering.

Recommendation algorithm
A system that chooses which posts, videos, songs or products to show you.

Reinforcement learning
Learning by trial and error with rewards. **RLHF** (reinforcement learning from human feedback) uses people's ratings to make chatbots more helpful.

S

Supervised learning
Learning from examples that are labelled with the correct answer.

Synthetic data
Artificial data generated by computers, often by AI, used for training or testing.

System prompt
Hidden instructions from the developer that set how a chatbot behaves.

T

Token
A piece of text a model reads and writes: a word, part of a word or a symbol. In English, roughly ¾ of a word.

Training
The process of adjusting a model's parameters using data so it gets better at a task.

Training data
The examples a model learns from.

Transformer
The neural network design (2017) behind most modern language models. It uses "attention" to weigh how words relate to each other.

Turing test
Alan Turing's 1950 idea: if a judge cannot tell a machine from a human in a text conversation, the machine passes.

U – W

Unsupervised learning
Finding patterns or groups in data without labels.

Watermark
A visible or invisible mark that shows content was made by AI.

Weights
Another word for a neural network's **parameters**, the strengths of its connections.

CHAPTER XXIV

Useful tools, resources *and further learning*



> *This guide is only a beginning. Here are trusted places to keep exploring, most of them free. Web addresses change, so if a link stops working, search for the name.*

FREE COURSES

Elements of AI · elementsofai.com
Free online course by the University of Helsinki and MinnaLearn, available in many European languages, including Croatian, Portuguese and Romanian.

Experience AI · experience-ai.org
Lessons for young people by the Raspberry Pi Foundation and Google DeepMind.

CS50's Introduction to AI with Python · cs50.harvard.edu/ai
Harvard's free course, for those who want to code.

Crash Course AI (YouTube)
Short, fun videos explaining the basics.

3Blue1Brown: Neural networks (YouTube)
Beautiful visual explanations of how neural networks and LLMs work.

PLAY AND EXPERIMENT

Teachable Machine · teachablemachine.withgoogle.com
Train your own image, sound or pose model in the browser.

Quick, Draw! · quickdraw.withgoogle.com
Draw, and a neural network guesses what it is.

Which Face Is Real? · whichfaceisreal.com
Test your skill at spotting AI-generated faces.

Moral Machine · moralmachine.net
MIT's game about the moral choices of self-driving cars.

CHECKING FACTS

Faktograf (HR) · faktograf.hr — **Polígrafo** (PT) · poligrafo.sapo.pt — **Factual** (RO) · factual.ro — **Raskrikavanje** (RS) · raskrikavanje.rs — **EDMO** · edmo.eu

TinEye · tineye.com and **Google Lens** for reverse image search; **InVID-WeVerify** plug-in for videos; **Content Credentials** · contentcredentials.org/verify.

ONLINE SAFETY AND WELLBEING

Better Internet for Kids · betterinternetforkids.eu
The EU portal with links to every national Safer Internet Centre: **Centar za sigurniji internet** (HR) · csi.hr, **Internet Segura** (PT) · internetsegura.pt, **Ora de Net** (RO) · oradenet.ro. In Serbia: **Pametno i bezbedno** · pametnoibezbedno.gov.rs.

116 111 · free helpline for children and young people in all four partner countries.

RIGHTS AND RULES

EU AI Act explorer · artificialintelligenceact.eu
AI Incident Database · incidentdatabase.ai: real cases where AI went wrong.
DigComp 2.2: the European framework of digital skills for citizens, including AI.

YOUTH AND ERASMUS+

European Youth Portal · youth.europa.eu — **SALTO-YOUTH** · salto-youth.net — **Youthpass** · youthpass.eu: record what you learned in projects like ours.

BOOKS

You Look Like a Thing and I Love You, Janelle Shane (2019): a funny introduction to how AI gets things wrong.
Unmasking AI, Joy Buolamwini (2023): the researcher behind *Gender Shades* on bias.
Weapons of Math Destruction, Cathy O'Neil (2016): how algorithms can deepen inequality.
Atlas of AI, Kate Crawford (2021): the hidden costs of AI for people and the planet.

DOCUMENTARIES

AlphaGo (2017) · *Coded Bias* (2020) · *The Social Dilemma* (2020)

KEEP THE CONVERSATION GOING

Share this guide with your class, run one of the activities from chapter 22, and tell us what you discovered. AI will keep changing. Your curiosity and judgement are what will keep you ready.

COLOPHON

Credits *and thanks*

About this book

AI Awareness Book: A Field Guide to Artificial Intelligence was created as an output of the Erasmus+ youth exchange **"AI: Friend or Foe?"** (project no. 2026-1-HR01-KA152-YOU-000394792), held in Zadar, Croatia, from 17 to 25 October 2026.

It was made by and for the young people who took part, together with their youth leaders, from **Srednja škola Glina** (Croatia, coordinator), **EB Sec. Bispo D. Manuel Ferreira Cabral** (Portugal), **Colegiul Național „Kölcsey Ferenc"** (Romania) and **Sedma beogradska gimnazija** (Serbia). Thank you to every participant, youth leader and teacher who shared ideas, questions and energy during the week.

In the spirit of chapter 15: AI tools were used to help draft and edit parts of the text; all content was reviewed by the project team.

Facts in this guide were checked in October 2026. AI changes fast: tools, prices and laws may have changed since. Always check the latest information.

Licence

The text of this booklet is shared under the Creative Commons licence **CC BY-SA 4.0**: you may copy, adapt and share it, also in translation, as long as you credit the project and share your version under the same licence. Images keep their own licences, listed on the right. All images were colour-toned for this booklet.

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 Agency for Mobility and EU Programmes (AMPEU). Neither the European Union nor the granting authority can be held responsible for them.

Image credits

All photographs and engravings come from Wikimedia Commons. Diagrams (Fig. 1.1, 1.3, 2.2, 3.1, 6.1, 7.1, 9.1, 13.1, 17.1, 19.2, 20.1) were drawn for this booklet. Licences: *CC BY / CC BY-SA* = Creative Commons Attribution (ShareAlike); *CC0* and *Public domain* = free of copyright.

Fig. 0.1 Compass rose, from Emanuel Bowen's map, 1748. Emanuel Bowen (via Geographicus). Public domain.

Plate I Difference Engine No. 1, Science Museum, London. Marcin Wichary. CC BY 2.0.

Fig. 2.1 Purkinje cells, drawing by Santiago Ramón y Cajal, 1899. Santiago Ramón y Cajal. Public domain.

Fig. 3.2 The Great Wave off Kanagawa, after Katsushika Hokusai. Katsushika Hokusai. Public domain.

Fig. 4.1 Alan Turing aged 16. Unknown (possibly A. R. Chaffin). Public domain.

Fig. 4.1 IBM Deep Blue, Computer History Museum. Anton Chiang. CC BY 2.0.

Plate II Jacquard loom with punchcards, National Museum of Scotland. Stephencdickson. CC BY-SA 4.0.

Fig. 5.1 Hollerith electric tabulating machine, 1902. U.S. Census Bureau. Public domain.

Fig. 6.2 Subject catalogue (Schlagwortkatalog). Dr. Marcus Gossler. CC BY-SA 3.0.

Fig. 8.1 Rosetta Stone, British Museum. APK. CC BY 4.0.

Plate III Panopticon, plan by Jeremy Bentham / Willey Reveley, 1791. Jeremy Bentham. Public domain.

Fig. 10.1 Cottingley Fairies photograph. Elsie Wright. Public domain.

Fig. 11.1 The Mechanical Turk, engraving, 1783. From K. G. von Windisch's book on the Turk. Public domain.

Fig. 12.1 Guide to finger-print identification, 1905. Internet Archive Book Images. No known restrictions.

Fig. 13.2 Numbered phrenological head, 19th century, Wellcome Collection. Wellcome Collection. CC BY 4.0.

Fig. 14.1 Liberty Bell slot machine. Nazox. CC BY-SA 3.0.

Fig. 15.1 France in the 21st century: At School. Jean-Marc Côté or Villemard. Public domain.

Plate IV Metropolis (1927), the Maschinenmensch. Horst von Harbou. Public domain.

Fig. 16.1 Raspberry Pi 5. Simon Waldherr. CC BY 4.0.

Fig. 17.2 CERN Computer Center. Simon Waldherr. CC BY-SA 4.0.

Fig. 18.1 Helen Keller with Anne Sullivan, July 1888. Unknown photographer (New England Historic Genealogical Society). Public domain.

Fig. 18.1 Stephen Hawking (colorized). NASA. Public domain.

Fig. 19.1 Telephone switchboard operators, Seattle, ca. 1922, MOHAI. Webster & Stevens. Public domain.

Fig. 20.2 Carta Marina, 1539. Olaus Magnus. Public domain.

Fig. 21.1 Ada Lovelace, portrait, about 1840. Alfred Edward Chalon. Public domain.

Plate V Two women operating ENIAC. U.S. Army photographer. Public domain.

Co-funded by the European Union Erasmus+ Enriching lives, opening minds.