TRAINING PROGRAM · 25 MODULES

ARTIFICIAL INTELLIGENCE

Understanding the fundamentals, mastering the use cases, anticipating tomorrow.

> *We haven't built a machine that thinks like us.*
> *We've built a mirror that computes.*

FUNDAMENTALS · TECHNOLOGY · USE CASES · ETHICS · THE FUTURE

TABLE OF CONTENTS

Training outline

What is artificial intelligence?

Artificial intelligence (AI) refers to the set of techniques that allow a machine to perform tasks associated with human intelligence: recognizing an image, understanding a sentence, generating text, making a decision. It does not **think** in the way we do: it **computes** probabilities, at massive scale, based on what it observed during training.

> *Artificial intelligence invents nothing: it recombines, at a speed our brain cannot match, everything humanity has already thought, written, drawn, or coded.*

Three circles, three levels

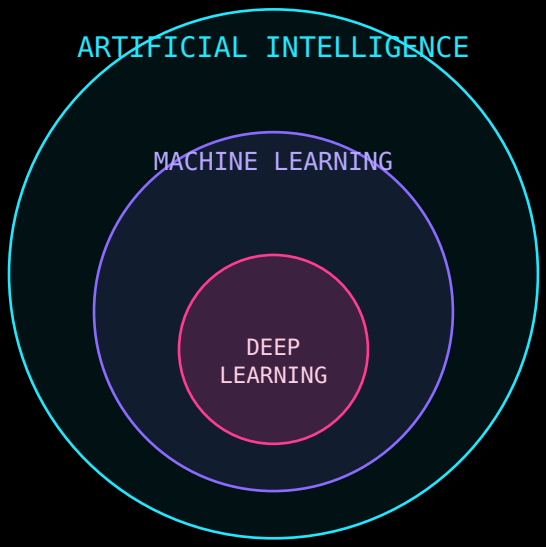

AI — the overall scientific field, born in the 1950s, aiming to simulate intelligent behavior.

Machine Learning — a family of methods where the machine learns rules from data, instead of being told the rules directly.

Deep Learning — a sub-family of ML based on deep neural networks, behind the spectacular progress made since 2012.

A short history of AI

A field made of huge hopes, two long "winters" when funding collapsed, then an abrupt awakening driven by computing power and massive data.

> *Every spring of AI was preceded by a winter when the dream seemed buried. The history of this technology isn't a straight line: it's a breathing pattern.*

The main families of AI

TODAY

Narrow AI (weak AI)

Excels at one specific task: recognizing a face, translating text, recommending a show. This is **100% of the AI that exists today**, including the most impressive models.

HYPOTHESIS

General AI (AGI)

An intelligence able to adapt to any domain, the way a human can. It **does not exist yet** and is the subject of intense debate over its feasibility and timeline.

Three ways to classify an AI

BY CAPABILITY	Narrow AI (specialized) → general AI (AGI, hypothetical) → superintelligence (speculative)
BY MECHANISM	Rule-based systems (symbolic) → machine learning (statistical) → deep learning (neural networks)
BY USE	Predictive AI (scoring, ranking) → generative AI (creates content) → agentic AI (acts, uses tools)

> *Mistaking a system that finishes your sentences for a mind that thinks is the single biggest misunderstanding of our time.*

Machine Learning: the basis of it all

Instead of hand-coding every rule, we show the machine thousands of examples and let it find the pattern itself. There are three main ways to do this.

01

Supervised learning

The machine is given **already labeled** examples (photo → "cat"). It learns to reproduce the label on new cases.

02

Unsupervised learning

No labels at all: the machine searches on its own for **groupings** and hidden structures in the data.

03

Reinforcement learning

The AI acts within an environment and receives a **reward or penalty**, much like training a behavior through trial and error.

One shared principle

> *In all three cases, the machine doesn't understand what it's doing: it optimizes a number. We are the ones who give it meaning upfront — and lend it meaning afterward.*

The artificial neural network

Loosely inspired by the brain: layers of mathematical "neurons" turn an input signal into an output decision, by adjusting thousands, even billions, of parameters.

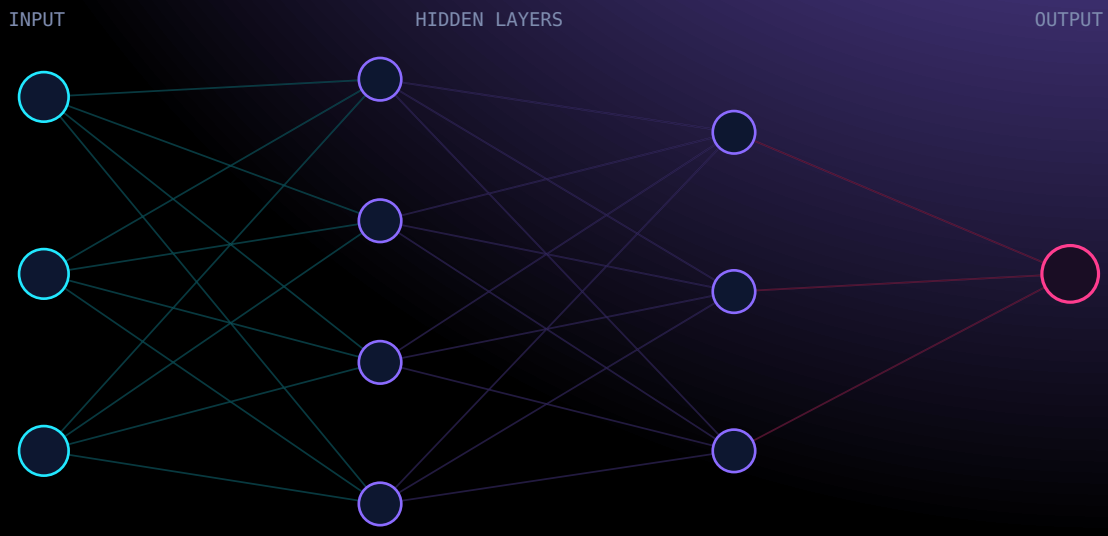

> *A single artificial neuron "knows" nothing. It adds, multiplies, compares against a threshold. The magic isn't in one neuron: it emerges from the number of them and how they're organized.*

How an AI actually learns

Training is a loop repeated millions of times: the machine predicts, its error is measured, and it adjusts its parameters to be slightly less wrong next time.

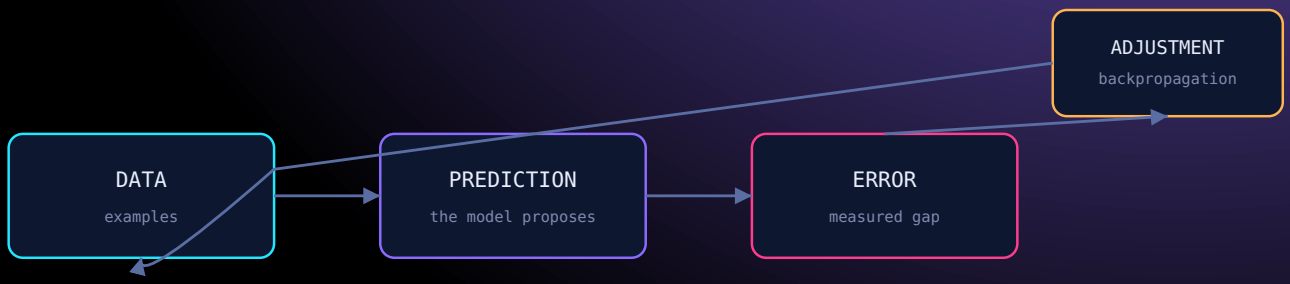

The loop repeats millions of times, until the error becomes acceptable.

> *Learning, for an AI, is not understanding. It's being wrong millions of times, faster than we could ever afford to be wrong.*

Data: the invisible fuel

An AI is never better than the data it was trained on. Too little data, biased data, or poorly cleaned data: the result will inherit it directly.

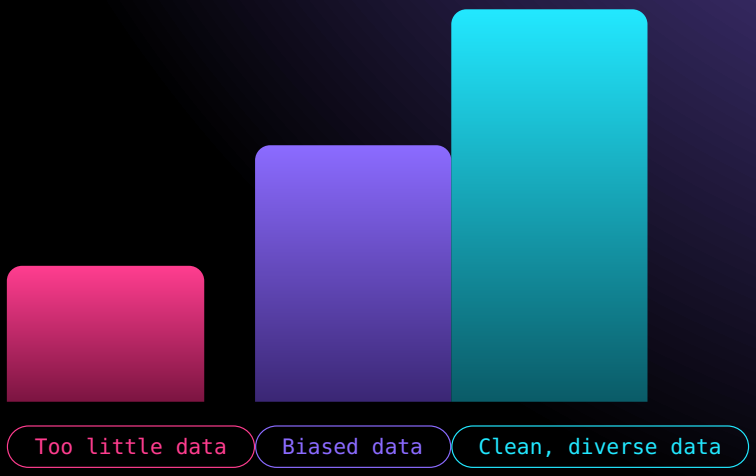

Perceived quality of the final model, depending on training data quality (illustration).

> *"Garbage in, garbage out": an AI trained on a society's blind spots will confidently reproduce those very same blind spots.*

Large Language Models (LLMs)

An LLM (*Large Language Model*) is a neural network trained on huge amounts of text to predict, word by word (or "token" by "token"), the most likely continuation of a sentence.

The token

A word is often split into several fragments called **tokens**. The model never "sees" whole words, only a sequence of these numerical fragments.

The parameter

A **parameter** is an internal setting adjusted during training. Recent models have anywhere from several billion to several hundred billion of them.

Inference

This is when the already-trained model **generates a response** to your question, token by token, in a matter of seconds.

Context

The **context window** is the amount of text the model can "look at" at once when formulating its response.

> *A language model isn't looking for the truth: it's looking for the most likely continuation. Most of the time, the two coincide. Not always.*

The Transformer architecture

Published in 2017, this architecture made today's explosion of generative AI possible. Its core mechanism: **attention**.

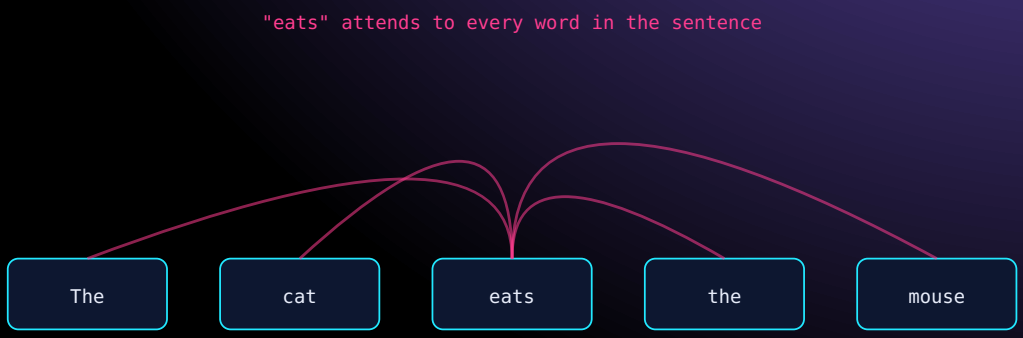

Instead of reading a sentence in strict linear order, the Transformer evaluates, for each word, **how relevant every other word** in the sentence is, all at once. This is what lets it grasp long-range context, a distant negation, or an indirect reference.

> *Attention isn't reading one word after another: it's taking in the whole sentence at once, and deciding, for each word, what actually matters.*

The prompt: talking to an AI

The "prompt" is the instruction you give the model. Its quality directly determines the quality of the response.

1. **Give clear context**
   Who are you, for what purpose, for whom? An AI doesn't guess your situation.

2. **Specify the expected format**
   List, table, email, desired length: say so explicitly.

3. **Provide examples**
   A concrete example is often worth more than a long abstract description.

4. **Iterate**
   A first answer is a draft: refine it through successive exchanges.

> *An AI does not read your mind. It reads exactly what you wrote: no more, no less.*

AI in business

A customer service team receives thousands of messages every day. A generative AI drafts a first response in seconds: a human agent reviews it, corrects it, sends it. Processing time drops sharply, but a new role emerges: supervising the machine rather than writing from scratch.

WHAT THE COMPANY GAINS

Faster processing, round-the-clock availability, consistent answers, and the ability to absorb activity spikes.

WHAT STAYS HUMAN

Judgment on sensitive cases, empathy with an upset customer, and the final decision that carries real accountability.

> *AI doesn't replace a job: it redistributes its tasks. What used to be invisible — judgment, accountability, arbitration — suddenly becomes the real value of human work.*

AI and healthcare

An AI analyzes an X-ray and flags an anomaly that the human eye, tired after the fiftieth image of the day, might have missed. It doesn't make a diagnosis: it flags, alerts, and prioritizes. The doctor keeps the final decision.

Here, AI's role isn't to replace medical expertise, but to serve as a **tireless, systematic second pair of eyes**. The risk: excessive trust that leads people to forget to check what the machine didn't see.

> *When a machine touches human life, the question is no longer just "does it work?" but "who remains responsible when it doesn't?"*

AI and education

A student stuck on a math problem asks an AI for help at 10pm on a Sunday. It doesn't hand over the answer: used well, it guides, asks questions in return, and helps the student find the reasoning themselves. A tutor available around the clock, at almost no cost.

THE PROMISE

Personalized support, at each student's own pace, available at any hour and without judgment.

THE RISK

The temptation to have the AI do the homework instead of understanding it: ease can undermine the effort of learning.

> *A good tutor, human or artificial, doesn't give the answer: it makes the student capable of finding it alone one day. That's where the real stakes of AI in education lie.*

Limitations: hallucinations and errors

A "hallucination" is a false answer generated with the same confidence as a true one. The model doesn't know that it doesn't know.

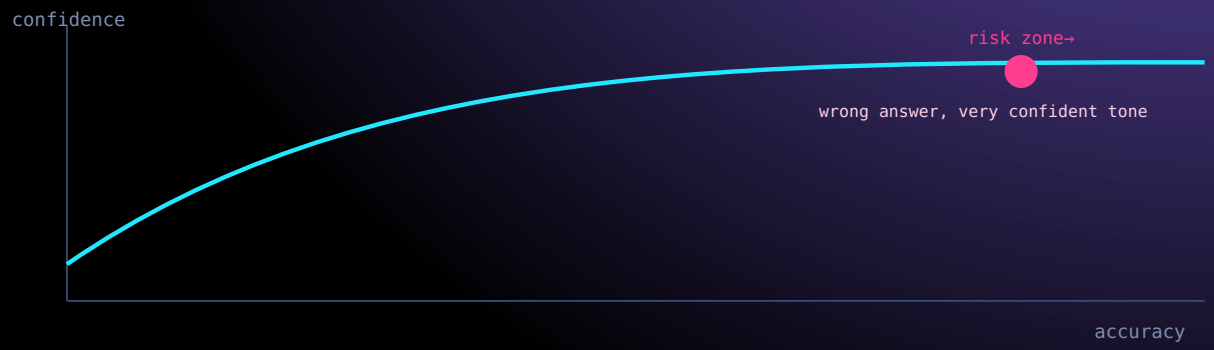

Illustration: an AI's apparent confidence is not correlated with the actual accuracy of its answer.

Why it happens

The model generates the statistically most likely word, not the verified one. On a topic poorly represented in its training data, or a very precise fact (a date, a number, a quote), it may "invent" a plausible answer rather than admit uncertainty.

> *An AI never lies — lying requires knowing the truth. It gets things wrong with exactly the same confidence it gets things right. It's up to us to keep the critical thinking it lacks.*

Algorithmic bias

An AI trained on data produced by humans inevitably inherits the imbalances, stereotypes, and blind spots present in that data.

DATA BIAS

Certain groups or situations are underrepresented in the training examples.

HISTORICAL BIAS

The data reflects past inequalities: the AI learns them as if they were the norm.

USAGE BIAS

The way an AI is used can amplify certain outcomes at the expense of others.

> *A biased AI isn't a broken AI: it's an AI working exactly as intended, on data that already carried our flaws.*

Ethics and AI governance

The more powerful a technology becomes, the more central questions of accountability, transparency, and control become.

TRANSPARENCY	Knowing you're interacting with an AI, and understanding, at least broadly, the basis on which it produces a response.
ACCOUNTABILITY	Clearly identifying who answers for the consequences of a decision made with AI assistance: the vendor, the operator, or the user.
HUMAN OVERSIGHT	Keeping a human able to understand, correct, or stop the system, especially for high-impact decisions.
FAIRNESS	Checking that systems don't systematically disadvantage certain groups of people.
REGULATORY FRAMEWORK	Regulations are emerging at national and international level to govern the most sensitive uses of AI.

> *A technology without governance isn't neutral: it simply moves forward without anyone having chosen the direction.*

AI and work

Every major technological wave has destroyed jobs — and created others, often unpredictable at the moment the wave began. AI is no exception, but its distinctive feature is that, for the first time, it touches **intellectual and creative tasks**, long thought to be beyond automation.

MOST EXPOSED TASKS

Repetitive writing, sorting and summarizing information, first drafts of code, first-line customer support.

WHAT GAINS VALUE

Judgment, human connection, creativity that breaks the mold, and the ability to ask the AI the right questions.

> *The question is no longer "will AI take my job?" but "which part of my work is still worth doing myself?"*

Generative AI: text, image, sound, code

The same core principle — predicting the most likely next element — now applies to almost every type of content.

TEXT

Writing, summarizing, translation, conversational answers, and computer code generation.

IMAGE

Creating illustrations from a text description, retouching, and visual style transformation.

SOUND & VOICE

Realistic voice synthesis, assisted music composition, and automatic transcription.

CODE

Code generation, completion, and correction, speeding up software development.

> *Creating has never been this fast. Discerning what deserves to be created has never been this important.*

Best practices for everyday use

Using AI effectively means staying in control of the tool rather than submitting to it.

01

Check sensitive facts

Numbers, dates, quotes, sources: cross-check them before sharing.

02

Protect sensitive data

Don't share confidential or personal information in a prompt.

03

Keep your critical thinking

A fluent answer is not proof of accuracy.

04

Iterate rather than settle

Rephrase, clarify, and challenge the answer you get.

> *The best user of AI isn't the one who trusts it the most: it's the one who knows precisely where that trust should stop.*

Security: deepfakes and disinformation

The same technology that generates useful content can fabricate a convincing voice, face, or video that is entirely fake.

Vigilance reflexes

1. **Check the source**
   Viral content not confirmed by an identifiable outlet is suspect.
2. **Look at the details**
   Inconsistent lighting, motion, or lip movement in fabricated videos.
3. **Be wary of emotional urgency**
   Disinformation seeks an immediate reaction, not reflection.

> *Seeing has never required this much verifying. The perfect image has, paradoxically, become the least trustworthy one.*

Looking ahead: a general intelligence?

Some researchers believe that an artificial general intelligence (AGI), able to match humans at any cognitive task, could emerge within the coming decades. Others see this prospect as distant, or even ill-defined. The debate remains open: no one today can predict the timeline with any certainty.

> *Perhaps the real question isn't when the machine will become as intelligent as we are, but what we will choose to remain, once it is.*

Glossary of key terms

AI	Artificial Intelligence: the set of techniques simulating intelligent behavior.	**ML**	Machine Learning: automated learning from data.
DEEP LEARNING	Deep learning, based on multi-layer neural networks.	**LLM**	Large Language Model: a large model trained on massive amounts of text.
TOKEN	A fragment of text (word or sub-word) processed by the model.	**PROMPT**	An instruction or question submitted to a generative AI.
HALLUCINATION	A false answer generated with apparent confidence.	**BIAS**	A systematic distortion inherited from training data.
AI AGENT	A system able to use tools and act to achieve a goal.	**RAG**	Retrieval-Augmented Generation: the AI consults external documents before answering.
FINE-TUNING	Lightly re-training an existing model on specific data.	**AGI**	Artificial general intelligence, hypothetical, comparable to human intelligence.
TRANSFORMER	A neural network architecture based on the attention mechanism.	**INFERENCE**	The moment an already-trained model generates a response.
PARAMETER	An internal model setting, adjusted during training.	**CONTEXT WINDOW**	The amount of text the model can consider at once.

CONCLUSION

What this training leaves you with

Artificial intelligence is neither magic nor inherently threatening: it's an extraordinarily powerful statistical tool, built on data, probabilities, and calculations. Understanding its **fundamentals** lets you get the most out of it without blind trust, grasping its **real-world uses** without ignoring the risks, and approaching its **future** with clarity rather than fascination or fear.

> *AI will not replace your critical thinking. It extends it — or numbs it — depending on how you use it.*

Understand | Question | Use with discernment