

# Artificial Intelligence Basics — 2026 Complete Course

---

## Artificial Intelligence Basics — The Complete Beginner Course

TechSparkling Free Course — 10 chapters · 26 lessons · 3 hands-on labs

---

### How to use this course

AI is everywhere in 2026 — in search, in your phone, in the tools you use at work — but most people can't explain what it actually is. This course fixes that. It's written for complete beginners: no math prerequisites, no code required for the first eight chapters, and three optional hands-on labs to see real AI working. Expect 12–25 hours total. Pair it with our articles on AI search engines, open-source AI models, and AI agents for the deeper dives.

---

## Chapter 1 — What AI Actually Is

### Lesson 1.1 — AI, ML and the hype

Three terms, three layers:

- **Artificial Intelligence** — the goal: machines doing tasks that require intelligence
- **Machine Learning (ML)** — the method: programs that learn patterns from data instead of following fixed rules
- **Deep Learning** — a style of ML using neural networks with many layers

The "AI" you use daily is almost all machine learning — usually deep learning. There is no magic; there is pattern-matching at enormous scale.

### Lesson 1.2 — What AI can and cannot do in 2026

**Can do well:** recognize objects in images, transcribe and translate speech, draft and summarize text, write code, recommend content, predict trends from data.

**Cannot yet do:** reliably reason about the world, stay truthful consistently, understand cause and effect, or be held accountable. These limits are why every serious AI use comes with human review.

## Lesson 1.3 — A simple mental model

AI models are **predictors**: given what came before, they predict what comes next — the next pixel in an image, the next word in a sentence. The "intelligence" emerges from the size of the training data and the model, not from a programmed understanding of the world.

## Exercises — Chapter 1

1. Find three "AI" features you used this week and name the pattern each one predicts.
  2. Write down one task you'd trust AI to do and one you wouldn't — and why.
- 

# Chapter 2 — A Short History You Can Actually Follow

## Lesson 2.1 — The pattern of cycles

AI has had waves of hype and disappointment since the 1950s. The current wave differs because of three things: **more data** (the internet), **more compute** (GPUs), and **better algorithms** (deep learning).

## Lesson 2.2 — The last 10 years in five jumps

1. 2012 — deep learning wins image recognition, shocking the field
2. 2016 — AlphaGo beats the world's best Go player with learned strategy
3. 2020 — GPT-3 shows that scaling language models produces capable writing
4. 2022-2023 — ChatGPT goes mainstream; generative AI becomes a consumer product
5. 2024-2026 — multimodal models (text+images+audio), agents that act, and AI built into every platform

## Lesson 2.3 — Why 2026 feels different

AI moved from a feature you opt into to a layer underneath everything: search answers, phone features, office tools, developer tools. The skills that matter now are not technical — they are

**judgment:** knowing when to trust AI, when to verify, and how to direct it.

---

## Chapter 3 — How Machine Learning Works

### Lesson 3.1 — Learning from examples

Instead of rules ("if photo has wheels, it's a car"), ML models see thousands of labeled examples and adjust internal parameters until they can classify new ones. This is supervised learning: input → output pairs.

### Lesson 3.2 — Three kinds of learning

- **Supervised** — learning from labeled examples (spam/not spam)
- **Unsupervised** — finding structure in unlabeled data (customer segments)
- **Reinforcement** — learning from rewards and penalties (game-playing AI)

### Lesson 3.3 — Training vs inference

- **Training** — the expensive part: the model learns patterns from a dataset (needs GPUs, weeks, money)
- **Inference** — the cheap part: a trained model answers new inputs (milliseconds)

This distinction explains why some AI is free (inference is cheap at scale per query) and why training is only for companies.

### Exercises — Chapter 3

1. Pick three apps and classify their AI as supervised, unsupervised or reinforcement.
  2. Explain training vs inference to someone in your own words in under 60 seconds.
- 

## Chapter 4 — Data: The Fuel

### Lesson 4.1 — Garbage in, garbage out

An AI is only as good as its data. Biased, incomplete or poisoned training data produces a biased, unreliable model. This is not an edge case; it's the normal state of real-world AI.

## Lesson 4.2 — The data behind the giant models

Large language models train on vast slices of the public web: books, articles, forums, code. That training data includes errors, opinions and copyrighted material — which is why models can be confidently wrong and why copyright lawsuits continue.

## Lesson 4.3 — Your data is training data

Anything you paste into a consumer AI tool may be used for training (or at least stored). Rule for sensitive material: use enterprise/paid tiers with data controls, or local models. Our article on open-source AI models explains how to run models locally so your data never leaves your machine.

---

# Chapter 5 — Neural Networks Without the Math

## Lesson 5.1 — The neuron analogy

A neural network is layers of simple units ("neurons") connected by weights. During training, the network adjusts the weights to reduce errors — millions of tiny nudges that produce pattern recognition.

## Lesson 5.2 — Layers and depth

- **Input layer** — the raw data (pixels, words)
- **Hidden layers** — intermediate patterns (edges, then shapes, then objects)
- **Output layer** — the prediction (label, next word, answer)

"Deep" learning = many hidden layers = more abstract patterns.

## Lesson 5.3 — Why scale matters

More data + more parameters + more compute tends to produce more capable models — a pattern that surprised researchers repeatedly. This is why the biggest models come from the companies with the biggest compute budgets, and why open-weight models (like Llama or DeepSeek, covered

in our open-source AI guide) are such a big deal: they give everyone else access to large-scale capability.

---

## Chapter 6 — Large Language Models Explained

### Lesson 6.1 — What an LLM is

An LLM (large language model) predicts the next token (roughly, the next piece of a word) given the previous tokens. That sounds trivial. The surprise: at massive scale, next-token prediction produces models that can translate, summarize, write code and hold a conversation.

### Lesson 6.2 — How ChatGPT actually works

1. Your prompt is converted to tokens
2. The model computes the probability distribution of the next token
3. It samples a next token, appends it, repeats — building the answer step by step

There is no internet search built in by default, and there is no "reading" happening at reply time — the model recalls patterns from training.

### Lesson 6.3 — Hallucinations explained

Because the model predicts plausible text rather than checking facts, it can generate confident nonsense — "hallucinations." Mitigations: retrieval (give it your documents), grounding (search/evidence), and always verifying important claims. This is why our AI search engines article emphasizes clicking the citations.

### Lab 1 — See it think

Use a free AI chat tool. Ask it to explain a concept, then ask it to explain again with a different example. Notice: same words? different words? invented details? Write down one hallucination risk you spot.

---

## Chapter 7 — Practical AI Skills: Prompting

## Lesson 7.1 — Prompting is a skill

The same model gives wildly different results depending on how you ask. Learn the craft:

- **Be specific** — "write a 200-word summary for a busy manager" beats "summarize this"
- **Give context** — who is the reader? what's the goal? what's the tone?
- **Provide structure** — "output as a table with three columns"
- **Show an example** — few-shot: give it a sample before asking

## Lesson 7.2 — Prompt patterns that work

| Pattern    | Example                                          |
|------------|--------------------------------------------------|
| Role       | "Act as a cybersecurity analyst..."              |
| Format     | "Answer as a numbered list of max 5 items"       |
| Constraint | "Do not mention prices"                          |
| Iteration  | "Make it shorter / more formal / for a beginner" |
| Verify     | "Which sources are you using? List them"         |

## Lesson 7.3 — AI in your workflow, honestly

Great uses: first drafts, brainstorming, summarization, translation, code scaffolding, learning explanations. Bad uses: writing final medical/legal/financial answers, generating content you can't verify, replacing your thinking.

## Lab 2 — Before and after

Take one task you do regularly (an email, a report, a lesson plan). Write it first without AI. Then prompt AI to improve the draft. Compare: what did it actually improve? What did it keep wrong?

---

# Chapter 8 — AI Tools You Should Know in 2026

## Lesson 8.1 — The five categories

1. **Chat assistants** — general purpose: answer, write, reason (ChatGPT, Claude, Gemini)
2. **Search with AI** — cited answers over the live web (Google AI Mode, Perplexity, ChatGPT Search)
3. **Generators** — images, video, audio, code from descriptions
4. **Agents** — tools that act: code, research, schedule, automate (covered in our agents article)
5. **Embedded AI** — features inside the apps you already use (Office, Photoshop, phone OS)

## Lesson 8.2 — Free vs paid, honestly

Free tiers are for learning and light use; paid tiers add context length, speed, priority and data controls. If you use AI professionally with sensitive data, the paid tier's data-privacy guarantees can be worth the money. For everything else, free is genuinely fine.

## Lesson 8.3 — A starter stack (free)

- A chat assistant (any of the big three)
- An AI search engine (cited answers)
- An image generator for concepts
- Your platform's built-in AI (it's already in your phone and tools)

## Exercises — Chapter 8

1. Use two different chat assistants on the same prompt and compare.
  2. Use an AI search engine with citations enabled for tomorrow's decision.
  3. List the AI already built into tools you use daily — you'll be surprised how much you had.
- 

# Chapter 9 — AI Ethics in Real Life

## Lesson 9.1 — Bias

Models learn bias from data. A hiring tool trained on historical hires may reproduce historical discrimination. Bias is a systems problem: audit outputs, add human review, and demand transparency from vendors.

## Lesson 9.2 — Privacy

You (or your organization) are responsible for what you put into an AI. Assume consumer tools store and may train on input. Use local models or enterprise tiers for sensitive information — our open-source AI models article covers running models on your own hardware.

## **Lesson 9.3 — Deepfakes and you**

AI can fake voices (seconds of audio), faces (photos and video), and text in your style. Defenses: be skeptical of urgent voice/video requests, verify through a channel you initiated, and protect your biometric data. In organizations: verify financial instructions out-of-band — the classic deepfake scam targets finance teams.

## **Lesson 9.4 — Your rules of thumb**

- Verify anything that matters (double-check the source, not the answer)
  - Disclose AI use where honesty requires it
  - Never put secrets into tools without data controls
  - Treat AI output as a draft from a confident intern — talented, not reliable
- 

# **Chapter 10 — Your AI Action Plan**

## **Lesson 10.1 — This week**

- Pick two free tools: one chat assistant, one AI search engine
- Write your personal AI-use policy (3 sentences about what you will and won't trust it with)
- Try one prompt each in role, format, constraint and iteration patterns

## **Lesson 10.2 — This month**

- Complete one real project with AI: a report, a script, a lesson, a translated doc
- Test one local open-source model (Ollama on any computer) to experience offline AI
- Set a rule for sensitive data and publish it to your family/team

## **Lesson 10.3 — This year**

- Build a review habit: every quarter, re-check what tools changed and what you should verify more
- Practice the most valuable AI skill: knowing when NOT to use it

- Follow the layer cake: chat for drafts, search-with-AI for facts, agents for tasks, review always

### Lab 3 — The verification drill

Ask an AI chat assistant for three specific factual claims (dates, statistics, quotes) about any topic. Then verify each one against a primary source using an AI search engine with citations. Write down how many were right. This single drill is the most important exercise in this course.

---

### Course glossary

| Term              | Meaning                                           |
|-------------------|---------------------------------------------------|
| AI                | Machines performing intelligent tasks             |
| ML                | Learning patterns from data                       |
| Deep learning     | Multi-layer neural networks                       |
| Neural network    | Layers of simple units with adjustable weights    |
| Training          | The expensive learning phase                      |
| Inference         | Using a trained model to answer                   |
| LLM               | Large language model — predicts next tokens       |
| Token             | A chunk of text an LLM processes                  |
| Hallucination     | Confident but wrong output                        |
| Prompt            | Your instruction to an AI model                   |
| Retrieval         | Giving the model your documents to ground answers |
| Open-weight model | Public model files you can run yourself           |

---

**Course complete!** You now understand what AI is, how machine learning and LLMs work, how to prompt effectively, which tools matter, and how to use AI ethically and safely. You've also run three real labs. The next step is practice: use AI for a real project every week, and verify everything that matters. Our articles on AI search engines, AI agents and open-source AI models will take you deeper whenever you're ready.