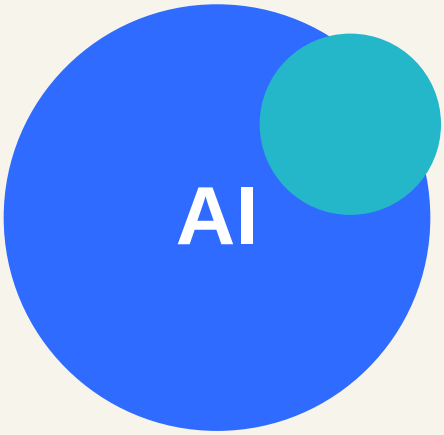
EXPLAINED VISUALLY

AI Made Simple

A visual beginner's guide



CHORAL SOFTWARE



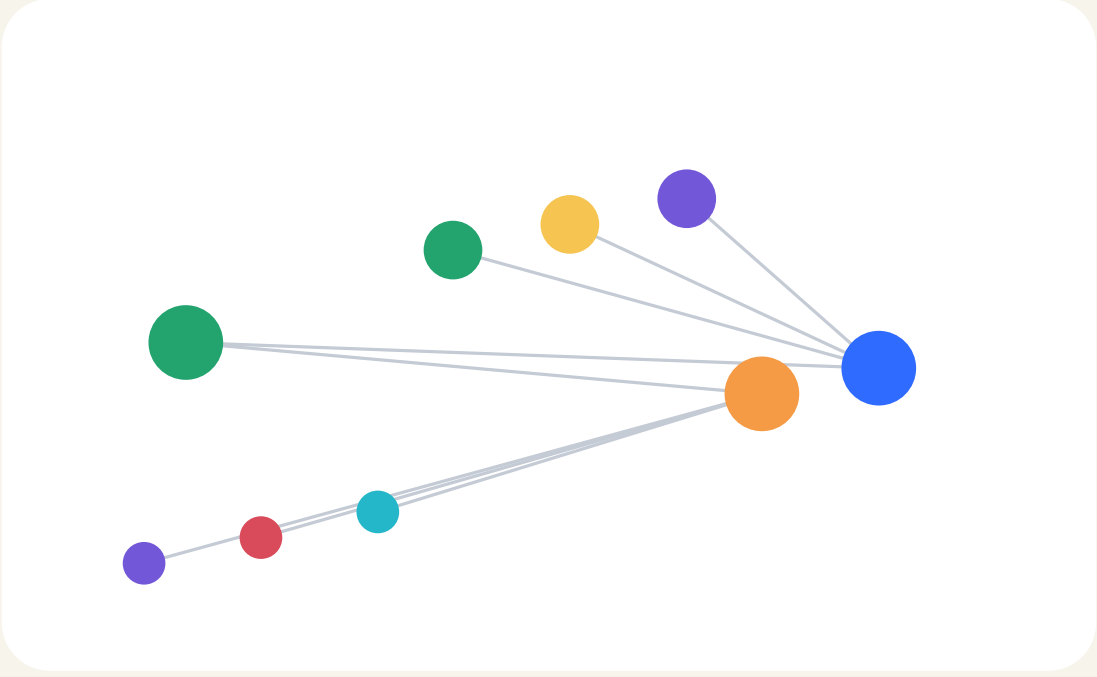
AI MADE SIMPLE

A visual beginner’s guide

A THREE-HOUR VISUAL TOUR

AI Made Simple

Generative AI, language models, coding assistants, and agents - without the jargon



Publication notes

Copyright © 2026 Choral Software. All rights reserved.

This book is an independent educational work. It is not affiliated with or endorsed by any technology company.

All diagrams, interface sketches, examples, names, organizations, datasets, and scenarios in this edition are original and fictional. No company logos, proprietary screenshots, stock illustrations, or copyrighted characters are used.

The book provides general education, not medical, legal, financial, security, or professional advice. AI systems and policies change quickly; verify important information with current authoritative sources.

Print specification: 7 × 10 inch paperback, premium color, white paper. Interior artwork is vector-based for sharp reproduction.

ASSET PROVENANCE

Every visual in this PDF is drawn from basic vector geometry by the book’s production script.

Your route through AI

01 What AI actually is 5

02 How machines learn 14

03 Inside a language model 25

04 Generative AI beyond words 38

05 Working well with AI 49

06 AI that writes code 62

07 From chatbots to agents 73

08 Build a simple workflow 84

09 What can go wrong 93

10 Responsible use 104

11 Where AI is going 113

CHAPTER 01

What AI actually is

A map of the field, without the mythology

CHAPTER COMPASS

Read the explanation, study the visual, then answer the two-minute check.

CORE IDEA

A useful definition

Artificial intelligence is the broad name for computer systems designed to perform tasks that normally require human judgment, perception, language, or planning. The label covers many different techniques. A spam filter, a navigation system, a speech recognizer, and a text generator can all be called AI even though they work in different ways.

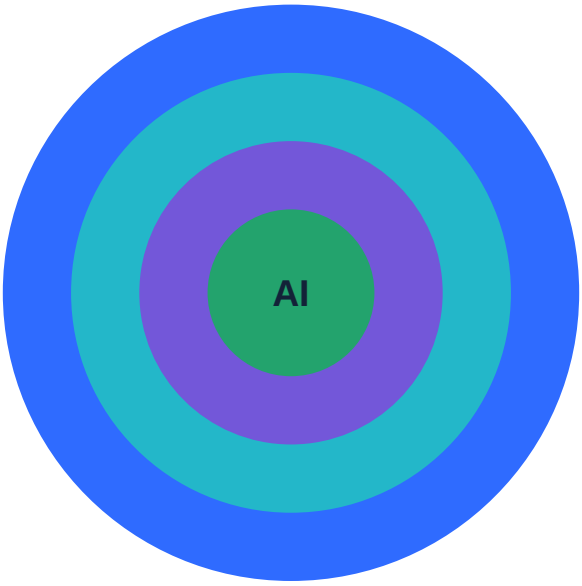
IN PLAIN LANGUAGE

Think of AI as a large toolbox, not a single machine. Some tools sort, some predict, some recommend, and some generate. The right first question is not “Is this AI?” but “What task is the system doing, what information does it use, and how is its result checked?”

KEEP IN MIND

AI is a field and a design goal; it is not a promise that a machine has a mind.

VISUAL MODEL



Two-minute check

How would you explain this idea to a friend using an everyday example?

MEMORY HOOK

AI is a field and a design goal; it is not a promise that a machine has a mind.

CORE IDEA

Rules and learned patterns

Traditional software follows rules written explicitly by people. Machine-learning software derives useful patterns from examples. A tax calculator may apply fixed formulas; a fraud detector may learn statistical clues from past transactions. Modern products often combine both approaches: learned components handle ambiguity, while ordinary code enforces limits and business rules.

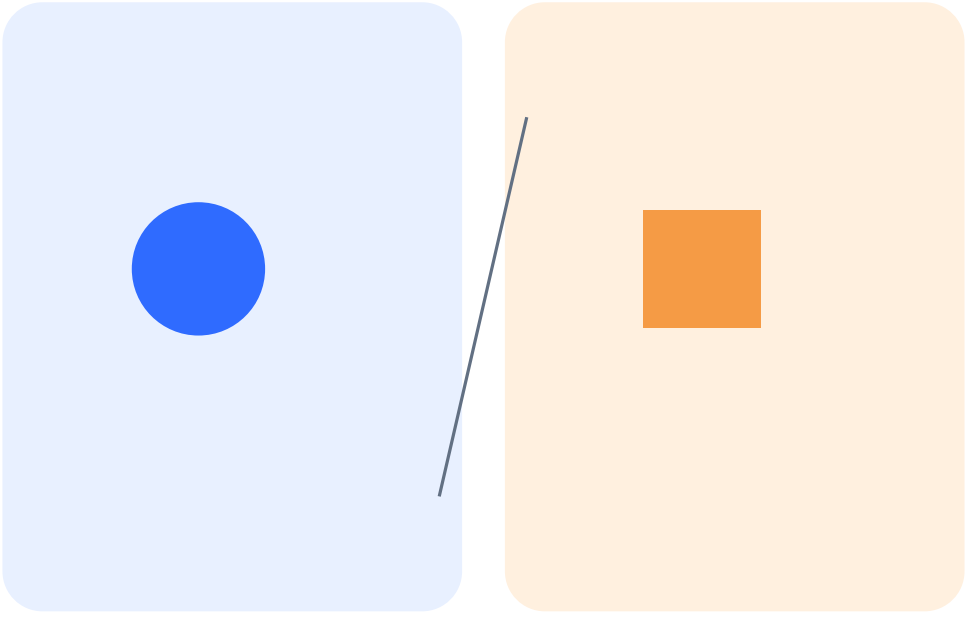
IN PLAIN LANGUAGE

Imagine sorting fruit. A rule-based program might say “if it is round and red, call it an apple.” A learned model sees many labeled examples and builds a more flexible boundary. It may work better on varied fruit, but its boundary can still be wrong.

KEEP IN MIND

Learning from examples does not remove human choices. People choose the data, objective, tests, and acceptable errors.

VISUAL MODEL



Two-minute check

What evidence would make you trust a system using this idea?

MEMORY HOOK

Learning from examples does not remove human choices. People choose the data, objective,…

CORE IDEA

The nested vocabulary

Machine learning sits inside the broader field of AI. Deep learning is a family of machine-learning methods built from many layers of adjustable numerical operations. Generative AI is defined by what it does: it creates new content. Large language models are deep-learning systems trained on sequences of tokens. Agentic systems connect models to tools, memory, and control loops.

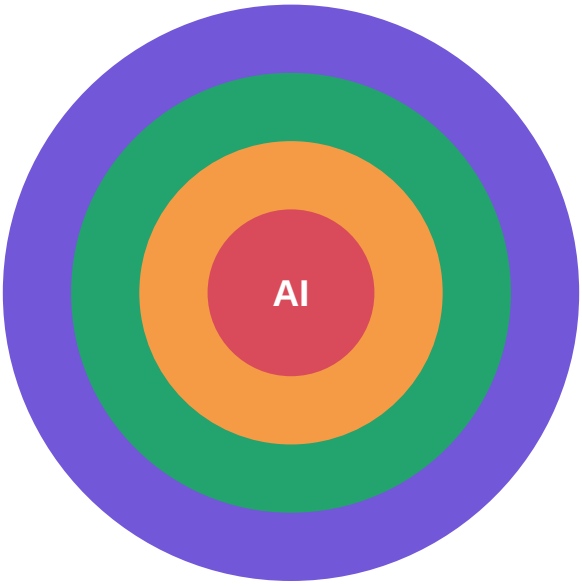
IN PLAIN LANGUAGE

The terms overlap, so a system can be several things at once. A language model can be deep learning, generative AI, and one component inside an agent. A recommendation engine may be AI and machine learning without being generative.

KEEP IN MIND

When a headline uses “AI,” look for the more precise category underneath.

VISUAL MODEL



Two-minute check

Where could a hidden assumption change the result?

MEMORY HOOK

When a headline uses “AI,” look for the more precise category underneath.

CORE IDEA

Capability is not consciousness

AI systems can produce fluent explanations, imitate styles, and solve unfamiliar-looking problems. Those behaviors are impressive, but they do not by themselves establish awareness, feelings, intentions, or human-like understanding. A useful working stance is to judge observable capabilities while remaining cautious about claims concerning inner experience.

IN PLAIN LANGUAGE

A weather model may be excellent at predicting a storm without experiencing rain. Likewise, a language model can describe grief without feeling grief. The output can still be useful; the distinction matters when we decide how much authority or trust to give it.

KEEP IN MIND

Human-sounding language invites projection. Treat warmth and confidence as interface qualities, not proof of a mind.

VISUAL MODEL



Two-minute check

Which part belongs to the model, and which part belongs to ordinary software?

MEMORY HOOK

Human-sounding language invites projection. Treat warmth and confidence as interface…

CHAPTER 02

How machines learn

From examples to a working model

CHAPTER COMPASS

Read the explanation, study the visual, then answer the two-minute check.

CORE IDEA

Training, validation, and testing

Training adjusts a model using examples. Validation helps designers choose settings and notice overfitting while development is underway. Testing estimates performance on held-out examples that were not used to tune the system. After deployment, inference is the act of applying the trained model to a new input.

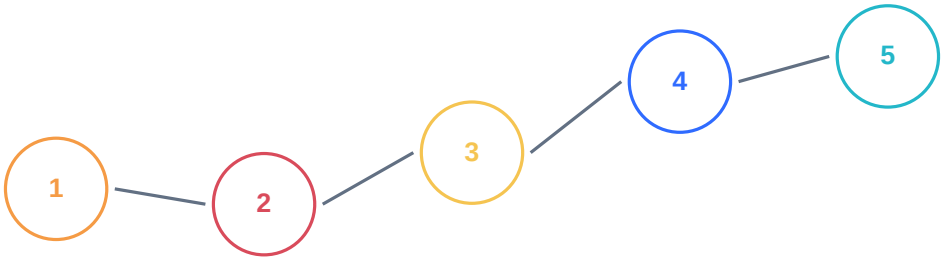
IN PLAIN LANGUAGE

A student who memorizes the practice sheet may fail a new exam. In the same way, a model that performs beautifully on familiar data may fail in the world. A separate test set makes the exam more honest.

KEEP IN MIND

A test score is evidence about a particular dataset and task, not a lifetime guarantee.

VISUAL MODEL



Two-minute check

What would a responsible human need to review?

MEMORY HOOK

A test score is evidence about a particular dataset and task, not a lifetime guarantee.

CORE IDEA

Features, labels, and objectives

A feature is information supplied to a model. A label is the target answer in supervised learning. The objective is the numerical goal used during training. For an email filter, features may represent words and sender patterns, the label may be spam or not spam, and the objective rewards correct classification.

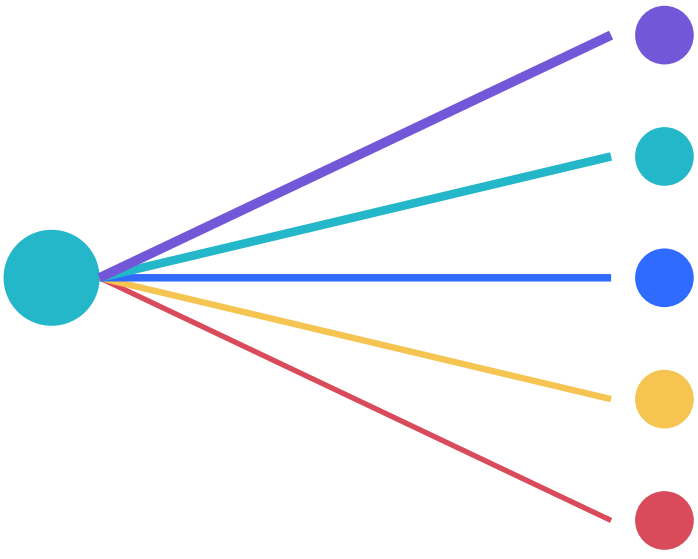
IN PLAIN LANGUAGE

Objectives are compressed definitions of success. If a delivery system is rewarded only for speed, it may ignore cost or safety. Designers often need several measures and explicit constraints.

KEEP IN MIND

A model becomes good at the target it is given, which may not perfectly match the outcome people truly want.

VISUAL MODEL



Two-minute check

How would you explain this idea to a friend using an everyday example?

MEMORY HOOK

A model becomes good at the target it is given, which may not perfectly match the outcome…

CORE IDEA

Parameters and optimization

A model contains adjustable numbers called parameters. During training, an optimization process repeatedly compares predictions with desired outcomes and makes tiny parameter changes that reduce error. Large neural networks may contain millions, billions, or even trillions of these values. The learned knowledge is distributed across them rather than stored as a tidy list of facts.

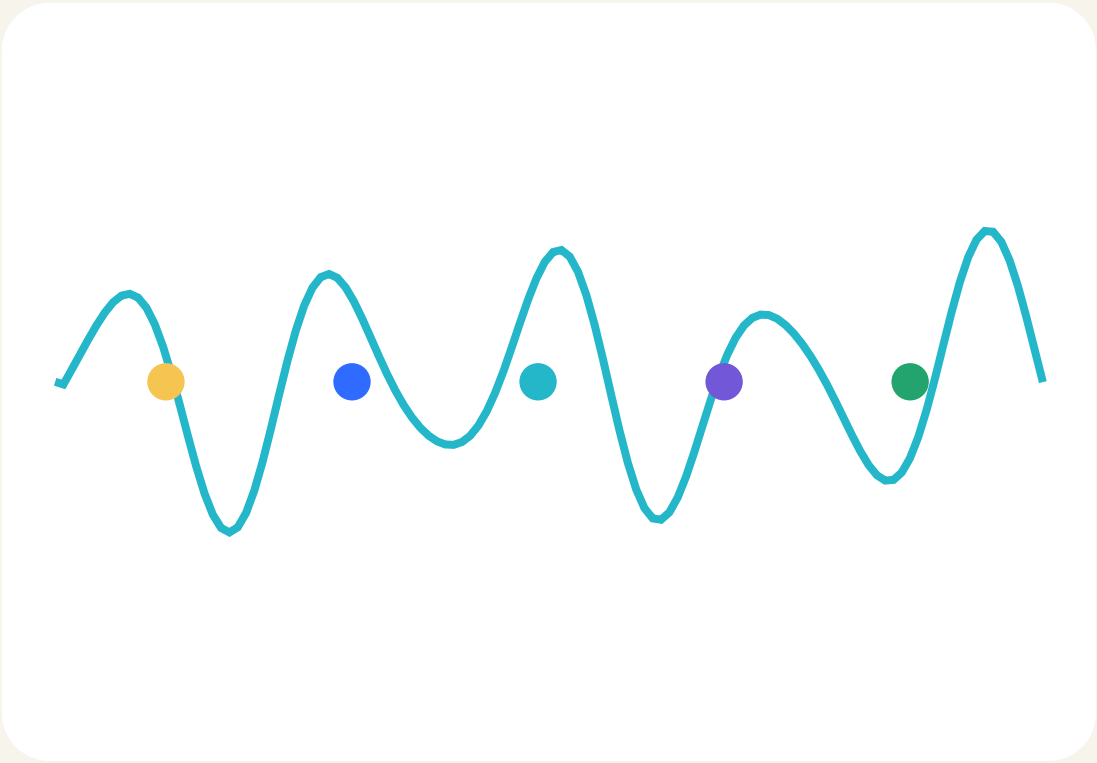
IN PLAIN LANGUAGE

Picture a huge sound-mixing board. Training turns many knobs by small amounts until the overall output better matches the examples. No single knob means “cat” or “grammar”; behavior emerges from combinations.

KEEP IN MIND

Parameters are not database rows. Inspecting one number rarely reveals a human-readable fact.

VISUAL MODEL



Two-minute check

What evidence would make you trust a system using this idea?

MEMORY HOOK

Parameters are not database rows. Inspecting one number rarely reveals a human-readable…

CORE IDEA

Generalization and overfitting

Generalization means performing well on new examples that resemble the intended real-world task. Overfitting occurs when a model becomes too tailored to its training examples or accidental quirks. More complexity is not automatically better; useful systems balance expressive power with reliable behavior on unseen data.

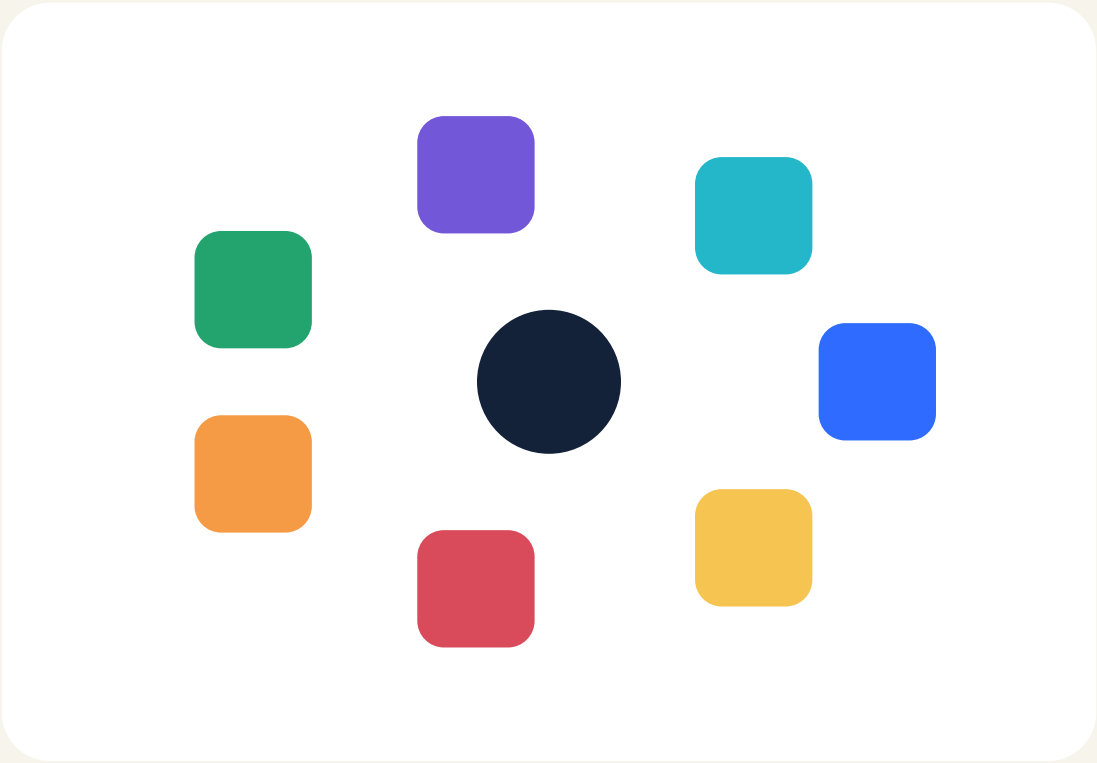
IN PLAIN LANGUAGE

A model trained to identify boats might accidentally learn that blue backgrounds mean “boat” if nearly every training photo shows water. It can then misclassify a blue wall. Diverse tests expose shortcuts like this.

KEEP IN MIND

Ask what changed between the training environment and the environment where the model is used.

VISUAL MODEL



Two-minute check

Where could a hidden assumption change the result?

MEMORY HOOK

Ask what changed between the training environment and the environment where the model is…

CORE IDEA

Three broad learning styles

Supervised learning uses examples paired with desired answers. Unsupervised or self-supervised learning finds structure or creates prediction tasks from unlabeled data. Reinforcement learning improves behavior using feedback tied to actions and outcomes. Modern systems often combine these styles at different stages.

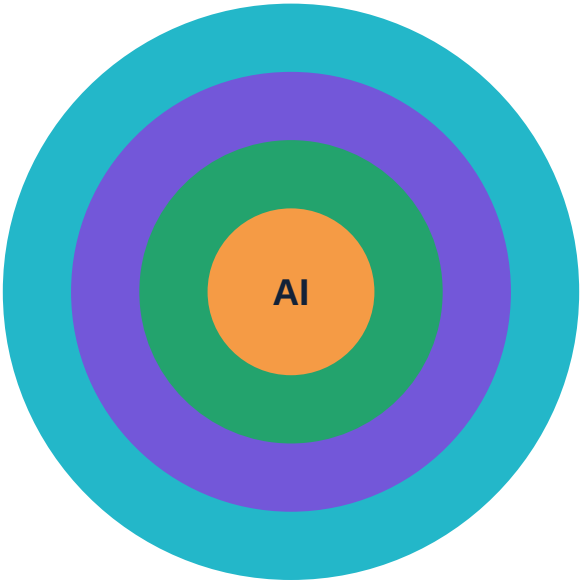
IN PLAIN LANGUAGE

A language model can learn from predicting missing or next tokens, then receive instruction examples, then be tuned using preference feedback. The final assistant is the result of a pipeline, not one magical lesson.

KEEP IN MIND

The name of a learning style tells you how feedback is provided, not whether the result is safe or intelligent.

VISUAL MODEL



Two-minute check

Which part belongs to the model, and which part belongs to ordinary software?

MEMORY HOOK

The name of a learning style tells you how feedback is provided, not whether the result is…

CHAPTER 03

Inside a language model

Tokens, context, prediction, and attention

CHAPTER COMPASS

Read the explanation, study the visual, then answer the two-minute check.