

What Is Artificial Intelligence?

A clear map of what AI actually means -- narrow vs. general AI, how it differs from traditional software, and where Machine Learning, Deep Learning, and Generative AI fit.

Beginner 15 min

Generated on 5 October 2026

Read online: <https://www.learnforgex.com/en/topics/what-is-ai>

What Is Artificial Intelligence?

This is the first lesson of the new "AI" course, and the first lesson of its "AI Fundamentals" category. If you already know Java, Spring, or React, you're coming into this course with a strong technical background -- but the topic itself is brand new here, so we're starting from zero, exactly the way "React Fundamentals" did for React. There's no code in this lesson. The goal isn't to write anything yet -- it's to build an accurate mental map of what "AI" actually means, before the next three lessons (Machine Learning, Deep Learning, Generative AI) each zoom into one piece of that map, and before later categories (Large Language Models, Tools & MCP, AI Agents) build real, runnable systems on top of it.

What Is Artificial Intelligence?

Artificial intelligence (AI) is the field of computer science concerned with building systems that perform tasks which normally require human intelligence -- recognizing an image, translating a sentence, playing a game, recommending a product, answering a question. That definition is intentionally broad, because AI isn't one single technique. It's an umbrella term that covers everything from a simple `if`-based chess-move picker written in 1970 to a modern system that writes working code from a plain-English description.

Under the AI umbrella there are also older approaches based on hand-written rules (like the 1970 chess-move picker above). But today, when people say "AI," they almost always mean -- and the rest of this course focuses on -- the following approach: **software whose behavior is learned from data or examples, rather than explicitly written line by line by a programmer.** A traditional program's behavior is exactly what its `if / for / switch` statements say it is. This learning-based approach has its own name -- **Machine Learning** -- and we'll zoom into it in the very next lesson. Keep that one sentence in mind; it's the thread that connects everything in this course.

Why Does It Exist?

Some problems are genuinely hard to solve with explicit rules, no matter how many `if` statements you're willing to write. Consider "is this photo a picture of a cat?" You could try to hand-write rules -- "if it has pointy ears and whiskers and a certain fur texture" -- but real photos are messy: different angles, lighting, breeds, partial occlusion, blurry backgrounds. Nobody has ever successfully hand-coded a rule-based cat detector that works reliably on real-world photos. Or consider "translate this sentence from Turkish to English" -- human language is full of ambiguity, idiom, and context that no fixed rule table captures.

AI exists because, for problems like these, it turns out to be far more effective to show a system **many examples** and let it find the pattern itself, than to try to describe the pattern by hand. This single idea -- learning from examples instead of being told the rules -- is why AI has succeeded at problems (vision, language, game-playing) that decades of traditional, rule-based software engineering struggled with.

History

The idea of a "thinking machine" is old, but the field of AI, as a named discipline, is usually dated to a 1956 summer workshop at Dartmouth College, where the term "artificial intelligence" was coined. A few years earlier, in 1950, Alan Turing had proposed his famous "Turing Test" as a practical way to ask "can a machine think?" without getting stuck in philosophy.

The decades that followed were not a straight line of progress. AI went through two well-documented "AI winters" (roughly the mid-1970s and the late 1980s/early 1990s) -- periods when early techniques hit their limits, funding dried up, and enthusiasm collapsed. What changed things was **data and computing power catching up with the ideas**. Around 2012, a technique called "deep learning" (which we'll meet properly in the "Deep Learning" lesson) started dramatically outperforming older approaches on image recognition, once large labeled datasets and powerful GPUs became available. That moment kicked off the modern AI era, which accelerated further in the 2020s with large language models -- systems trained on enormous amounts of text, capable of writing, reasoning about, and generating human-like language and code (the subject of this course's next category, "Large Language Models").

AI vs. Traditional Software

It's worth being explicit about a distinction that will come up constantly in this course, because it's the single biggest difference between AI systems and every other kind of software you've built in the Java, Spring, and React courses on this platform:

- **Traditional software is deterministic and rule-driven.** You write the exact logic. Given the same input, a correctly written `TopicController` (see this platform's own real controller, covered in the Spring MVC category) always produces the same output, and you can read the source code to know exactly why.
- **AI systems are trained, not programmed, for the specific task.** A developer doesn't write "if the image has these pixel patterns, it's a cat." Instead, a training process (see "Machine Learning") adjusts the system's internal parameters based on examples, until it produces good outputs on its own. The result is a system whose exact reasoning for any single output is much harder to read directly from "source code," because there isn't a human-written rule to point to.

This doesn't mean AI systems are "unpredictable magic" -- they follow precise mathematical rules, and their behavior can be tested, measured, and improved (a theme this course returns to in its "AI Evaluation" lesson later on). It means the way you build and reason about them is fundamentally different from writing a `for` loop.

Narrow AI and General AI

You'll often hear AI split into two categories:

- **Narrow AI** (also called "weak AI") is a system built to do one specific task well -- recognize handwritten digits, recommend a movie, translate a sentence, generate an image from a text prompt. Every AI system that exists and works today, including the most capable large language models, is narrow AI in this technical sense: extremely capable at a wide range of tasks, but still a trained system operating within the bounds of its training, not a general-purpose reasoning agent that can autonomously master any new domain the way a human can.
- **General AI** (also called "AGI," artificial general intelligence) refers to a hypothetical system with human-level (or beyond) reasoning ability across *any* intellectual task, not just the ones it was specifically trained or designed for. As of this course's writing, general AI does not exist -- it remains a research goal and a subject of active debate, not a product you can use today.

When you read a headline claiming an AI system "thinks" or "understands" something, it's almost always describing a narrow AI system performing impressively within its trained domain -- not evidence of general intelligence. Keeping this distinction in mind will make you a much more accurate reader of AI news.

Where Is AI Used Today?

AI already sits inside a lot of software you use daily, often invisibly:

- **Recommendations:** what a streaming service suggests you watch next, what an online store suggests you buy.

- **Search and ranking:** how a search engine decides which results are most relevant to your query.
- **Language:** spell-checkers, machine translation, and -- increasingly -- AI assistants that can write, explain, and refactor code.
- **Fraud and anomaly detection:** flagging an unusual credit card charge or fraud in a bank transaction.
- **Computer vision:** unlocking a phone with face recognition, detecting obstacles in a self-driving car's camera feed.
- **Speech:** voice assistants converting spoken words to text and back.

Notice that these examples span very different kinds of tasks -- text, images, numbers, sound. That variety is possible because "AI" isn't one technique; it's an umbrella over several distinct sub-fields, which is exactly what the next section maps out.

A Map of the Field: AI, Machine Learning, Deep Learning, and Generative AI

The next three lessons in this category each cover one term you'll see used (and sometimes confused) constantly: **Machine Learning**, **Deep Learning**, and **Generative AI**. Before diving into each one separately, it helps to see how they relate to each other -- they are not four different, competing technologies. They're **nested inside one another**, each one a more specific subset of the one before it:

- **Artificial Intelligence** is the broadest term -- any system that performs tasks requiring apparent intelligence, by any technique.
- **Machine Learning** is a subset of AI: specifically, systems that learn their behavior from data, rather than through hand-written rules (covered in depth in the next lesson).
- **Deep Learning** is a subset of Machine Learning: specifically, machine learning using multi-layered "neural networks" as the learning mechanism (covered in depth two lessons from now).
- **Generative AI** is a particular *application* of (usually deep learning based) models: models trained not just to classify or predict, but to *create* new content -- text, images, audio, code (covered in depth in this category's fourth and final lesson).

Picture four concentric circles, from largest to smallest: AI contains Machine Learning, which contains Deep Learning, which is the technique behind most of today's Generative AI. As plain text, that same map looks like this:

```

Artificial Intelligence
├─ Machine Learning      (learns from data)
│   └─ Deep Learning    (uses neural networks)
│       └─ Generative AI (creates new content)

```

Every generative AI system is deep learning; not every deep learning system is generative; not every machine learning system is deep learning; not every AI system is machine learning at all (that hand-written chess-move picker from the first section is AI, but it's not machine learning -- nothing in it was learned from data).

Running one real problem -- **spam email detection** -- through all four layers makes the difference concrete. Traditional software solves it with fixed rules ("is the sender on a blocklist?", "does it contain certain words?"). Machine Learning derives those rules itself from thousands of labeled emails instead of having them hand-written (covered in depth in the next lesson). Deep Learning does the same job by feeding a neural network the raw text of the email -- needing less manual selection of which words matter. Generative AI solves a different task: instead of predicting whether an email is spam, it could, for example, *generate* an explanation of why an email was flagged as spam -- generation, not classification.

Best Practices

- When you hear the word "AI" used in a product pitch or news headline, ask yourself which of the four circles above it's actually describing -- vague "AI" marketing claims become much easier to evaluate once you have this map.
- Judge an AI system by what it was actually trained and evaluated to do, not by how impressively it's described. A narrow AI system doing one task extremely well is still a narrow AI system.
- When learning this field, prioritize understanding the relationship between concepts (this map, and the "Training vs. Inference" distinction in the next lesson) over memorizing individual technique names -- the names change fast, the underlying relationships don't.

Common Mistakes

- **Treating "AI" and "AGI" as the same thing.** Every AI system in production use today, including the most capable language models, is narrow AI (see "Narrow AI and General AI"). Claims of "true understanding" or general intelligence should be read skeptically.
- **Assuming AI systems are simply "if-statements in disguise."** Traditional rule-based automation and AI solve problems in fundamentally different ways (see "AI vs. Traditional Software") -- confusing the two leads to wrong expectations about how these systems behave, fail, and improve.
- **Assuming deep learning is the only kind of machine learning, or that machine learning is the only kind of AI.** As the map in this lesson shows, each is a specific subset of the one before it -- useful, well-established AI and machine learning techniques exist that involve no deep learning (and no neural network) at all.

Summary, Cheat Sheet, and Glossary

Summary

- AI is software whose behavior comes from learning patterns in data or examples, rather than from explicitly hand-written rules.
- AI exists because some real-world problems (vision, language) are far easier to solve by learning from examples than by hand-coding rules.
- The field dates to 1956 (Dartmouth), went through two "AI winters," and entered its modern era around 2012 with deep learning, accelerating further in the 2020s with large language models.
- All AI systems in production use today are **narrow AI** -- highly capable at specific trained tasks, not general-purpose reasoning agents.
- AI, Machine Learning, Deep Learning, and Generative AI are **nested** concepts, each a more specific subset of the one before it -- not four separate technologies.

Cheat Sheet

- AI = any system performing tasks that normally need human intelligence.
- ML = AI that learns its behavior from data (next lesson).
- DL = ML using multi-layered neural networks (two lessons from now).
- GenAI = models trained to create new content, usually deep-learning based (this category's fourth lesson).
- Narrow AI = capable at one task/domain. AGI = hypothetical, doesn't exist yet.

Glossary

- **Artificial Intelligence (AI):** the broad field of building systems that perform tasks requiring apparent human intelligence.
- **Narrow AI:** an AI system built and trained for a specific task or domain -- the only kind of AI that exists in practice today.
- **Artificial General Intelligence (AGI):** a hypothetical system with human-level reasoning across any intellectual task; not yet achieved.
- **Turing Test:** a 1950 proposal by Alan Turing for judging whether a machine's behavior is indistinguishable from a human's.
- **AI winter:** a historical period of reduced funding and enthusiasm for AI research, following unmet expectations.