NO PREP

What Is Artificial Intelligence, Really?

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

What Is Artificial Intelligence, Really?

A Word That Gets Tossed Around

You have probably heard the phrase "artificial intelligence" hundreds of times by now. It shows up in news headlines, in conversations about your phone, in school discussions, and in debates about the future of work. Politicians talk about regulating it. Tech companies brag about building it. Your friends might joke that ChatGPT wrote their essay. But if someone asked you right now — "What is artificial intelligence, actually?"

Name: **Period:** **Date:**

What Is Artificial Intelligence, Really?

A Word That Gets Tossed Around

You have probably heard the phrase "artificial intelligence" hundreds of times by now. It shows up in news headlines, in conversations about your phone, in school discussions, and in debates about the future of work. Politicians talk about regulating it. Tech companies brag about building it. Your friends might joke that ChatGPT wrote their essay. But if someone asked you right now, "What is artificial intelligence, actually?", could you answer with real confidence? Most people cannot, and that includes most adults. The term gets stretched and squeezed to mean almost anything, which makes it nearly impossible to have an honest conversation about it. That confusion is where this course begins. Artificial intelligence is not magic. It is not a robot with feelings. It is not a digital brain that thinks the way you think. And despite what the movies suggest, it is not plotting to take over the world. What it actually is turns out to be both simpler and more interesting than most people expect, and understanding the real definition gives you a serious advantage in every conversation about AI from this point forward.

The Real Definition, and Why It Matters

At its core, artificial intelligence refers to software (that is, computer programs) that can perform tasks that would normally require human judgment or learning. Notice that phrase: "normally require human judgment." That is the key. For most of computing history, programs could only do exactly what a programmer explicitly told them to do, step by step. If a programmer forgot to include a rule, the program failed. AI changes that. AI programs can learn from examples, recognize patterns, and make decisions even in situations they have never seen before. A more technical way to put it: AI is a field of computer science focused on building systems that can perceive their environment, process information, and take actions to achieve a goal.

That definition comes from decades of academic research stretching back to the 1950s. But the practical meaning is simpler: AI is software that gets better at a task the more data it processes, rather than requiring a human to manually update every rule. Why does the definition matter? Because if you think AI is just "smart computers," you will misunderstand what it can and cannot do. AI systems are incredibly powerful within narrow tasks: recognizing a face, translating a sentence, predicting what video you want to watch next. But they have no understanding, no common sense, and no awareness. Knowing that distinction will help you use AI better and avoid being fooled by it.

Not All AI Is the Same

Here is something most people do not realize: there is no single thing called "AI." The phrase covers a huge family of different techniques, each suited to different problems. Some AI systems use rules explicitly programmed by humans, called rule-based systems or expert systems, where engineers write out thousands of "if-then" instructions. If the patient's temperature is above 103 degrees and they have a rash, flag for review. These systems can be effective, but they break down when reality does not match the programmer's assumptions. Other AI systems, the kind you interact with most today, use machine learning, which means the system is trained on massive amounts of data rather than hand-coded rules. A spam filter trained on millions of emails learns to recognize patterns in suspicious messages without anyone explicitly writing the rule "watch out for the phrase 'click here to claim your prize.'" A medical imaging AI trained on a million X-rays learns to spot signs of pneumonia that human radiologists might miss when they are tired. The learning happens through a process of trial and error at enormous scale, which we will dig into in a later unit. Then there is deep learning, a subfield of machine learning that uses structures loosely inspired by the human brain (called neural networks) to tackle incredibly complex tasks like understanding language, generating images, and translating speech in real time. The large language models behind ChatGPT, Claude, and Gemini are all deep learning systems. The differences between these types matter when you are evaluating what a particular AI tool can and cannot do.

What AI Is Not

Spend five minutes watching a Hollywood science fiction film and you will come away with a deeply misleading picture of AI. Movie AI is typically either a helpful servant with human-level understanding (think J.A.R.V.I.S. in Iron Man) or an existential threat with its own desires and plans (think Skynet in Terminator). Neither of these exists in 2026. Current AI systems — even the most impressive ones — do not have feelings, desires, consciousness, or goals of their own. They are sophisticated pattern-matching machines that have been trained on enormous datasets. It is also important not to confuse AI with automation in general. Automation has been happening for centuries — the assembly line, the ATM, the self-checkout machine. Automation replaces a specific repetitive task with a machine. AI is different because it introduces the ability to handle variability and uncertainty. An ATM can dispense cash using a fixed set of rules. An AI customer service chatbot can respond to thousands of different questions it has never seen before, in different languages, using context clues to figure out what the customer needs. The difference is meaningful. Finally, AI is not infallible. One of the most important things to understand right from the start is that AI systems make mistakes — sometimes spectacular ones. They can be confidently wrong. They can reflect the biases of the data they were trained on. They can be gamed or manipulated. And they have no way of knowing when they do not know something. Understanding these limitations is not pessimism. It is exactly the kind of critical thinking that makes you a smarter, safer user of these tools.

Where This Is All Headed

Artificial intelligence has existed as an academic field since the mid-twentieth century, but the current wave of AI tools — the ones you interact with on your phone, in search engines, and in school — are the product of a handful of breakthroughs that mostly happened between 2012 and 2023. In that roughly decade-long window, the combination of vastly more data, far more powerful computer chips, and new algorithmic techniques transformed AI from a niche research subject into a technology reshaping entire industries. Over the next few months in this course, you will build a complete picture of this technology: how it works under the hood, how to use it effectively and ethically, how to spot when content has been generated by AI, and how to protect yourself from its misuse. But

everything starts here, with one honest answer to the question you started with. Artificial intelligence is software that learns from data to perform tasks that would otherwise require human judgment. Keep that in your head, and you are already ahead of most people in the room.

Key Information: Artificial Intelligence

Detail	Information
Field of origin	Computer science; formally named in 1956 at Dartmouth College
Core definition	Software that learns from data to perform tasks requiring human judgment
Key distinction from regular software	AI learns from examples; traditional software follows fixed rules
Main types of AI	Rule-based systems, machine learning, deep learning
What AI lacks	Consciousness, feelings, desires, common sense
Common misconception	AI is not the same as automation; automation follows fixed rules
Current AI examples	ChatGPT, Gemini, Claude, Siri, spam filters, recommendation engines
Most important limitation	AI can be confidently wrong and reflects biases in training data

Case Study: Marcus's Homework Shortcut

Marcus is a junior who used an AI chatbot to write a quick summary of a chapter for his history class. The AI produced a confident, well-organized paragraph — but when his teacher asked him to discuss the main argument of the chapter in class, Marcus had nothing. He had not read a word. Later that week, his school held an assembly where the principal explained the new AI policy: students could use AI tools to brainstorm or check their work, but had to demonstrate understanding themselves. Marcus felt defensive at first. 'The AI was right, though,' he told his friend. 'It gave me correct information.' His friend, who had actually read the chapter, pushed back: 'Was it correct? Did you check it?' Marcus had not. That evening, he ran the AI summary against his textbook and found two details the AI had gotten confidently wrong — names and a date that seemed plausible but were inaccurate. It was the first time Marcus had a real reason to care about what AI actually is and how it works.

1. Based on the reading, why was Marcus vulnerable to the AI's errors in a way that his friend was not? What does this tell us about the relationship between AI tools and human understanding?

2. The reading states that AI can be "confidently wrong." How does Marcus's experience illustrate this risk, and what habit should users develop based on this limitation?

Multiple Choice

1. According to the reading, what is the key distinction between traditional software and AI?

    a) Traditional software follows fixed rules while AI learns from examples

    b) AI requires the internet while traditional software does not

    c) Traditional software is more expensive to build than AI

    d) AI is faster than traditional software when processing most complex tasks

2. Which of the following best describes what current AI systems lack?

    a) The ability to process large amounts of data

    b) Consciousness, feelings, and genuine understanding

    c) The ability to perform tasks without human input

    d) The ability to recognize patterns within large data structures

3. What does the reading identify as one of AI's most important limitations?

    a) AI requires human supervision and review for every decision

    b) AI systems can only work in English

    c) AI systems can be confidently wrong and reflect data biases

    d) AI cannot be used in schools or educational settings

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ Artificial Intelligence	A) A subfield of machine learning using brain-inspired structures to handle complex tasks
2. ___ Machine Learning	B) Software that learns from data to perform tasks requiring human judgment
3. ___ Deep Learning	C) An AI approach where engineers write explicit if-then instructions for every scenario
4. ___ Rule-Based System	D) A brain-inspired computing structure used in deep learning to process complex patterns
5. ___ Neural Network	E) A type of AI where a system is trained on massive data rather than hand-coded rules

Fill-in-the-Blank

1. At its core, artificial intelligence refers to software that can perform tasks that would normally require human _____ or learning.

2. AI systems are incredibly powerful within _____ tasks, such as recognizing a face or translating a sentence, but they have no understanding or common sense.

3. One of the most important things to understand about AI is that these systems can be confidently _____, sometimes making spectacular errors.

Short Answer

Answer each question in 2-4 complete sentences.

1. In your own words, explain the difference between AI and traditional automation. Use a specific example from the reading to support your answer.

2. The reading argues that understanding the real definition of AI gives you a "serious advantage." What does it mean in practical terms, and why does the definition matter for how you use these tools?

3. Why do you think so many people — including adults — have a confused or inaccurate picture of what AI is? What role do movies, news, and marketing play in creating that confusion?