LEVEL 1 | FOUNDATIONS OF AI & COMPUTING

Introduction to Artificial Intelligence

MODULE 1 | Complete Study Guide

What is AI · History · Categories · AI vs Humans · Myths vs Reality

TABLE OF CONTENTS

This module is the starting point for understanding AI. No prior technical knowledge is required.

01 WHAT IS ARTIFICIAL INTELLIGENCE?

Defining AI

Artificial Intelligence is the field of computer science focused on building machines that can do things that normally require human thinking. That includes understanding language, recognising images, making decisions, and learning from experience.

A simple way to think about it: AI is a set of instructions (called algorithms) that lets a computer look at data, find patterns, and make decisions or predictions based on those patterns.

SIMPLE DEFINITION

AI = getting computers to simulate human-like thinking and problem-solving.

What AI Actually Does

AI Can...

- Understands and generates human language (e.g. chatbots, translation)
- Recognises images and video (e.g. face unlock, medical scans)
- Makes decisions and predictions (e.g. credit scoring, recommendations)
- Controls physical systems (e.g. self-driving cars, robots)

Real Examples

- Plays complex games better than humans (e.g. Chess, Go)
- Writes code, essays, and creative content
- Diagnoses diseases from scan data
- Detects fraud in real time across millions of transactions

$200B+

Global AI market size in 2023

70%

Companies using AI in at least one function (McKinsey 2023)

300M

Jobs potentially affected by AI automation (Goldman Sachs 2023)

02 HISTORY OF AI

Key Milestones

YEAR	EVENT	WHY IT MATTERS
1943	McCulloch and Pitts publish a model of artificial neurons	First mathematical model of how a brain cell could work in a computer
1950	Alan Turing publishes 'Computing Machinery and Intelligence'	Proposed the Turing Test: if a machine can fool a human, it can be called intelligent
1956	Dartmouth Conference coins the term 'Artificial Intelligence'	The official birth of AI as a field of research

YEAR	EVENT	WHY IT MATTERS
1957	Frank Rosenblatt builds the Perceptron	First learning machine that could classify simple patterns
1966	ELIZA chatbot created at MIT	Early natural language program; showed machines could simulate conversation
1969-79	First AI Winter: funding and interest collapse	Early promises unmet; expectations far exceeded what hardware could do
1980s	Expert Systems become popular in industry	Rule-based AI used in medicine, finance, and engineering
1986	Backpropagation algorithm popularised by Rumelhart, Hinton, Williams	Made training multi-layer neural networks practical for the first time
1987-93	Second AI Winter	Expert systems proved too brittle and expensive to maintain
1997	IBM Deep Blue defeats world chess champion Garry Kasparov	Showed AI could beat humans at complex strategic games
2006	Geoffrey Hinton coins 'Deep Learning'	Revived neural networks with many layers; foundation of modern AI
2012	AlexNet wins ImageNet competition by a huge margin	Proved deep learning could process images far better than any previous method
2016	Google DeepMind's AlphaGo defeats world Go champion Lee Sedol	Go has more possible board positions than atoms in the observable universe
2017	Transformer architecture introduced (paper: Attention Is All You Need)	The foundation of GPT, BERT, and all modern large language models
2020	GPT-3 released with 175 billion parameters	Showed that very large language models could generate human-quality text
2022	ChatGPT reaches 100 million users in 2 months	Fastest consumer product adoption in history; brought AI to mainstream public
2023-24	GPT-4, Gemini, Claude, Llama compete as frontier models	AI embeds into search, productivity tools, healthcare, and education

KEY PATTERN

AI history follows cycles: excitement builds, reality disappoints, funding drops (an 'AI Winter'), then a breakthrough restarts the cycle. The current boom is the longest and most commercially successful yet.

From Calculators to Thinking Machines

Computers started as very fast calculators. They could only do exactly what a programmer told them, step by step. AI changed that by giving machines the ability to learn from data rather than following fixed rules. Here is how the capability of machines has grown over time.

ERA	WHAT MACHINES COULD DO	EXAMPLE
1940s-50s Basic Computing	Follow exact instructions only. No learning. No flexibility.	ENIAC calculated artillery trajectories using punched cards
1960s-70s Symbolic AI	Follow complex logic rules written by experts. Could reason but only within strict limits.	DENDRAL identified chemical compounds using expert-written rules
1980s-90s Expert Systems	Apply thousands of if-then rules from domain experts to make decisions.	MYCIN diagnosed blood infections with 65% accuracy, comparable to doctors
1990s-2000s Statistical ML	Learn from labelled data using statistics. Adapt to new examples.	Spam filters learned to classify email without hard-coded rules
2010s Deep Learning	Learn complex patterns from raw data using many layers of artificial neurons.	AlexNet recognised 1,000 image categories with near-human accuracy
2017+ Foundation Models	Learn from massive datasets and generalise to many different tasks with minimal fine-tuning.	GPT-4 writes code, explains science, and translates languages, all from one model

Hardware Made It Possible

The jump to deep learning in the 2010s was not just a software breakthrough. Three hardware trends made large-scale training affordable.

Graphics cards designed for games turned out to be perfect for training neural networks. A modern GPU performs thousands of operations in parallel.	Companies can now rent enormous computing power by the hour. A model that would have taken decades to train in 1990 can be trained in days.	The internet generated billions of labelled examples (images, text, audio) that AI systems could learn from at scale.

1,000x

Cost reduction in compute per FLOP since 2012 (OpenAI 2020)

3.4M x

More compute used in largest AI models vs 2012 (AI Index 2023)

2.3B

Internet users generating training data daily

04 AI REVOLUTIONS AND THE MODERN AI BOOM

Three Waves of AI

Researchers often describe AI as going through distinct waves or revolutions, each with a different approach to making machines intelligent.

Wave 1: Symbolic AI 1950s to 1980s	Wave 2: Statistical ML 1990s to 2010s	Wave 3: Deep Learning 2010 to present
Machines followed rules written by humans. Knowledge was hand-coded. Could handle narrow, well-defined problems only.	Machines learned patterns from labelled examples. Less hand-coding needed. Worked well when data was limited and structured.	Machines learn directly from raw data through many layers of processing. Can handle images, speech, text, and code at human or super-human levels.
Strength: transparent and explainable	*Strength: practical and reliable*	*Strength: flexible and powerful*
Weakness: breaks outside its rules	*Weakness: needs clean labelled data*	*Weakness: hard to explain decisions*

Why the Current Boom is Different

Previous AI booms faded because the technology hit real limits. The current boom is different for several reasons that make it more likely to persist.

--> **Scale works.** Larger models trained on more data consistently perform better. This is called the scaling law. Doubling parameters and data reliably improves performance across almost all tasks.

--> **General purpose.** Earlier AI was single-purpose. A chess AI could not play checkers. Foundation models today handle writing, coding, analysis, and image generation all in one system.

--> **Commercial pull.** In 2023, Microsoft invested $13B in OpenAI. Google, Amazon, and Meta each spent billions on AI infrastructure. This capital means research does not depend on government grants.

--> **Ecosystem maturity.** Open source models (Meta's Llama), accessible APIs, and cloud platforms mean any developer in the world can build AI products without building the underlying model.

$91.9B

Global AI investment in 2022 (CB Insights)

$500B

US Stargate AI infrastructure project announced 2025

$1.8T

Projected AI market value by 2030 (Grand View Research)

05 AI MYTHS VS REALITY

Common Misconceptions Explained

MYTH: AI is basically magic

Reality: AI is math and statistics. At its core, it finds patterns in numbers. The outputs can seem impressive, but they come from matrix multiplication and probability, not anything mysterious.

MYTH: AI thinks and understands like a human

Reality: AI processes patterns in data but has no understanding, beliefs, or intentions. A language model predicts the next likely word based on training data. It does not 'know' what the words mean the way a person does.

MYTH: AI will immediately take all jobs

Reality: AI automates specific tasks, not entire jobs. A radiologist's job involves talking to patients, ethical decisions, and physical procedures, not just reading scans. Most research finds AI augments workers more than it replaces them, at least in the near term. Oxford Economics (2023) projects net job growth in AI-adjacent roles.

MYTH: AI is always right

Reality: AI systems make mistakes, sometimes confidently. Language models 'hallucinate': they produce text that sounds correct but is factually wrong. AI image classifiers fail on unusual lighting or angles. Always verify AI outputs for important decisions.

MYTH: AI is conscious or could become conscious

Reality: No current AI system has any form of consciousness or inner experience. Consciousness is not well-understood by science and there is no evidence that scaling up a language model produces it. This remains an open philosophical question.

MYTH: AGI (human-level AI) is just around the corner

Reality: Researchers disagree sharply. Some put AGI arrival within 10 years; others say 100+ years or never. Current AI systems are narrow: extremely good at specific tasks but lacking flexible common sense and general reasoning.

06 NARROW AI vs AGI vs ARTIFICIAL SUPERINTELLIGENCE

One of the most important distinctions in AI is the difference between what AI can do now and what it might theoretically do in the future. These three categories help frame that.

	NARROW AI (ANI)	GENERAL AI (AGI)	SUPERINTELLIGENCE (ASI)
What it is	AI designed for one specific task or domain	AI that can learn and perform any intellectual task a human can	AI that surpasses human intelligence across all domains
Status	Exists now. Widely deployed.	Does not exist. Active research goal.	Does not exist. Theoretical concept.
Example	Netflix recommendations, Google Translate, medical diagnosis AI	A system that could learn medicine, law, programming, and art with the same ease a human could	A system that designs better AI than humans can, then continuously improves itself
Key limit	Cannot transfer learning outside its training domain. A chess AI cannot play checkers.	Would need common sense, physical understanding, and robust reasoning across all contexts.	By definition would outpace human ability to control or understand it.
Risk level	Low to moderate (bias, errors, misuse)	High (alignment is an open problem)	Existential risk if misaligned (Bostrom, Russell, Tegmark)

WHERE WE ARE TODAY

Everything you use right now (ChatGPT, Gemini, Claude, image generators, recommendation systems) is Narrow AI. AGI does not exist. ASI is purely theoretical. The gap between where we are and AGI is enormous, though researchers disagree on how long it will take to close.

07 WEAK AI vs STRONG AI

This distinction comes from philosophy, not engineering. It asks: does the AI actually think, or does it just simulate thinking?

WEAK AI (what exists)

- Also called Narrow AI in practice
- Simulates intelligent behaviour without any genuine understanding
- Has no consciousness, intentions, or awareness
- The machine acts as if it understands but it is processing symbols and statistics
- Every AI product that exists today is Weak AI
- A language model that writes poetry is not 'experiencing' creativity

STRONG AI (theoretical)

- A hypothetical AI that actually thinks and understands
- Would have genuine mental states, beliefs, and possibly consciousness
- Would understand language, not just process it statistically
- Currently impossible with known technology and methods
- Would pass not just the Turing Test but any deep test of reasoning
- Philosopher John Searle argued this may never be achievable with computers

SEARLE'S CHINESE ROOM (1980)

John Searle argued that a computer following rules to respond in Chinese does not understand Chinese, just as a person following a rulebook in a locked room does not understand the language either. This is the core argument that all current AI is Weak AI regardless of how convincing it seems.

08 AI vs HUMAN INTELLIGENCE

Where Each Excels

CAPABILITY	AI	HUMAN
Speed	Processes millions of data points per second	Limited by biological processing speed
Memory	Perfect recall of everything in training data	Forgetting is normal and often useful
Consistency	Same output for same input every time (if deterministic)	Mood, fatigue, and context change responses
Creativity	Remixes and recombines patterns from training data	Can generate truly novel ideas with no prior example
Common Sense	Weak. Can fail on simple physical or social reasoning	Strong. Developed through lived experience from childhood
Physical World	Cannot experience the physical world (no body in standard AI)	Embodied. Understands weight, texture, and temperature through experience
Social & Emotional	No genuine empathy or emotion. Can simulate social responses.	Deep social reasoning, empathy, and emotional intelligence
Learning Efficiency	Needs millions of examples to learn a task	Can learn from 1 to 5 examples (one-shot or few-shot learning)
Multi-task Flexibility	Narrow AI fails outside its training domain	Easily transfers skills across unrelated domains
Specialised Tasks	Superhuman in chess, Go, image recognition, protein folding	Human-level or below in narrow specialised benchmarks

Real Benchmarks: Where AI Stands Today

Standardised tests give a clearer picture of AI capability. Here are results from well-known benchmarks.

BENCHMARK	WHAT IT TESTS	HUMAN LEVEL	BEST AI (2024)
MMLU	University-level knowledge across 57 subjects	~89%	GPT-4: 86.4%
HumanEval	Writing correct code from a description	~77%	GPT-4: 87%
ImageNet	Classifying photos into 1,000 categories	~95%	ViT-22B: 90.9%
HellaSwag	Completing sentences with common sense	~95%	GPT-4: 95.3%
AlphaFold	Predicting protein 3D structure from sequence	Years of lab work	Near perfect (2022)
AIME (Math Olympiad)	Advanced competition mathematics	~85%	o1-preview: 83% (2024)

The Main Sub-fields

AI is not one technology. It is an umbrella term covering many different approaches and sub-fields. Understanding these branches helps you see how different AI products actually work.

BRANCH	WHAT IT DOES	EXAMPLES
Machine Learning (ML)	The most important branch today. Instead of programming rules, you give the machine data and let it find the patterns. ML is the foundation of almost all modern AI products.	Spam filter, fraud detection, product recommendations
Deep Learning (DL)	A subset of ML that uses artificial neural networks with many layers (hence 'deep'). Excellent at processing raw data like images, audio, and text. Requires large datasets and computing power.	Image recognition, voice assistants, language models (GPT, Claude)
Natural Language Processing (NLP)	Getting computers to understand, generate, and translate human language. Includes tasks like summarisation, translation, sentiment analysis, and question answering.	Google Translate, ChatGPT, grammar checkers, chatbots
Computer Vision (CV)	Giving machines the ability to interpret visual information from images and video. Used in both static image analysis and real-time video understanding.	Face ID, self-driving cars, medical imaging, quality control in factories
Robotics and Embodied AI	Combining AI with physical machines that can sense and interact with the real world. Involves perception, planning, and motor control working together.	Warehouse robots (Amazon Robotics), surgical robots (da Vinci), Boston Dynamics
Expert Systems and Knowledge Graphs	Encoding human expert knowledge as structured rules or connected facts. Less dominant today but still used in regulated industries where transparency is required.	Medical diagnosis assistants, legal research tools, Google Knowledge Graph
Reinforcement Learning (RL)	An AI agent learns by trial and error, receiving rewards for good actions and penalties for bad ones. Particularly powerful for games, robotics, and optimisation problems.	AlphaGo, game-playing AI, data centre cooling (Google DeepMind cut energy use by 40%)
Generative AI	AI that creates new content rather than just classifying or predicting. Generates text, images, audio, video, and code that did not exist in the training data.	DALL-E, Midjourney, Stable Diffusion (images); GPT-4, Claude (text); Sora (video)
Explainable AI (XAI)	Methods and tools that help humans understand why an AI made a particular decision. Critical for high-stakes applications where accountability matters.	Loan approval explanations, medical diagnosis justifications, autonomous vehicle decision logs

10 LEARNING OUTCOMES

By the end of this module, students should be able to demonstrate the following competencies.

LO 1 Understand AI Fundamentals

- ✓ Define AI in plain terms without using jargon
- ✓ Explain the difference between AI as a tool and AI as a concept
- ✓ Describe what an algorithm is and how it relates to AI
- ✓ Give at least three real-world examples of AI in everyday life

LO 2 Explain the Evolution of AI

- ✓ Name at least five key milestones in AI history and explain why each mattered
- ✓ Explain what caused the two AI Winters and why the current boom is different
- ✓ Describe how improvements in hardware, data, and algorithms enabled deep learning
- ✓ Place AI developments in chronological order from the 1950s to the 2020s

LO 3 Understand Different AI Categories

- ✓ Distinguish between Narrow AI, AGI, and ASI with examples of each
- ✓ Explain the difference between Weak AI and Strong AI and why it matters
- ✓ Compare AI and human intelligence across at least five specific capabilities
- ✓ Name and briefly describe at least five branches of AI
- ✓ Identify which branch of AI is behind a given product or application

SELF-CHECK QUESTIONS

Q1. What is the difference between Narrow AI and AGI? Give one example of each.

Q2. Name two reasons why the AI boom of the 2020s is different from previous booms.

Q3. A friend says 'ChatGPT understands everything it says.' How would you respond?

Q4. What are the three hardware trends that made deep learning practical?

Q5. Which branch of AI is used in self-driving cars? Which is used in language models?

Q6. Explain the Weak AI vs Strong AI distinction in one or two sentences.

Q7. What is an AI Winter? What caused the first one?

Q8. Name one task where AI outperforms humans and one where humans clearly outperform AI.