

Fundamentals of Artificial Intelligence

Teaching Demonstration · 10 Minutes

July 2026



A guided tour — 5 stops, ~2 minutes each

Xubin Wang

wangxubin@ieee.org wangxubin.site

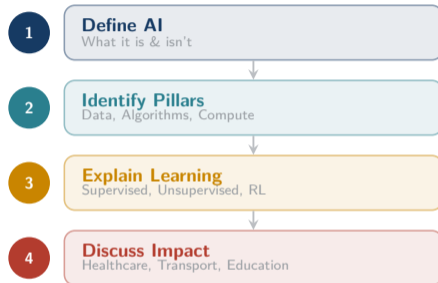
What We Will Learn Today

By the end of this session, you will be able to:

1. **Define** AI and distinguish it from human intelligence
2. **Identify** the three pillars that drive modern AI
3. **Explain** how machines “learn” from data with a real example
4. **Discuss** one AI application and its broader impact

▷ *Four measurable verbs —no memorization, just understanding.*

A Learning Journey



Teaching approach: Each step builds on the last —define → identify → explain → discuss.

Let's start with #1: what actually is AI?

Xubin Wang

What Is Artificial Intelligence?

AI = creating machines that perceive, reason, learn, and act in ways that mimic intelligent human behavior.

↑ *four verbs — use these to frame every question about AI*

Key distinction: AI is *not* about copying the brain.

It's about achieving **rational outcomes** through computation.

When I ask you “is this AI?” — I'm really asking:
Can it perceive, reason, learn, and act to achieve a goal? ▷ **Think:** ChatGPT — intelligent?

Functionally it perceives, reasons, learns, generates.


But does it truly *understand* anything?
That's the philosophical question.

Four ways to think about AI — a 2×2 human-like vs. rational, thinking vs. acting

| | | | | |
|--|----------|---|---|--------|
| | | Human-like | | |
| | Thinking | Thinking Humanly Cognitive models — replicate the mind | Acting Humanly Turing Test — pass for human | Acting |
| | | Thinking Rationally Formal logic, optimization | Acting Rationally Rational agents —act optimally ↑ <i>today's AI lives here</i> | |
| | | Rational / Ideal | | |

Russell & Norvig, AIMA

The AI Journey: 1950 – Today



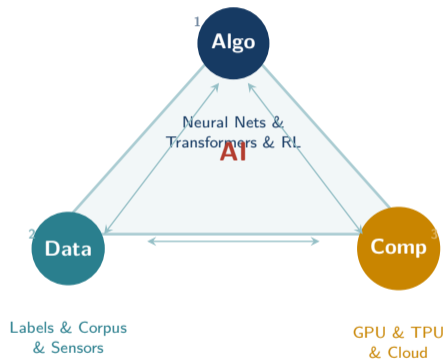
Computing power + data scale + algorithmic breakthroughs → accelerating progress

Notice the red bar —the AI Winter. Expectations outpaced reality, funding dried up. Then statistical ML revived the field, and 2012 triggered an explosion.

▷ **Think:** Biggest leap? My vote: 2012.
It wasn't just progress —it was a paradigm shift:
hand-crafted features → end-to-end learning.

Key pattern: Each era driven by a **new resource**:
theory → knowledge → data → compute → **scale**.

The Three Pillars of Modern AI



Three pillars. One engine. A virtuous cycle:

Bigger data → demands **better algorithms**

Better algorithms → need **more compute**

More compute → enables **even more data**

► Each pillar amplifies the next —a self-reinforcing loop.

Metaphor —a three-engine rocket: data = **fuel**, algorithms = **engine**, compute = **thrust**.

All three must fire together, or you stay on the ground.

▷ **Think:** Which pillar mattered most since 2012? *All three —ImageNet (data) + GPUs (compute) + AlexNet (algorithm).*

How Do Machines “Learn”?

Machine learning = finding patterns in data.

Not magic —it’s mathematics and statistics. Depending on your data, you pick one of three approaches:

Supervised

You have labeled
input → output pairs

Model learns
the mapping

e.g. spam detection
medical diagnosis

“here’s the answer”

Unsupervised

You have only
input, no labels

Model finds
hidden structure

e.g. customer segments
gene clustering

“find the pattern”

Reinforcement

Agent acts in
an environment
Gets reward or
penalty, adapts

e.g. AlphaGo
robot control

“try, fail, improve”


▼ *all three converge to the same engine* ▼

Key insight: Data in → model predicts → feedback returns → model improves.

That single loop is the engine of modern AI.

How does a computer recognize a cat?

▷ walk through the four steps —left to right ▷




CNN = Convolutional Neural Network. "Convolution" = scan image with small filters to detect edges, textures, shapes —same way our visual cortex works.

Input → [CNN Layers] → Features → Prediction: "Cat (97%)"

▷ **Think:** What if all training images are indoor cats? (*Hint: model fails outdoors —data bias*)

▷ *Training data must reflect reality —or your model will fail. A fundamental challenge.*

We've moved from “What is this?” to “Create this!”



*The underlying algorithm is the same — neural networks trained on data.
What changed is the **target** and the **scale**.*

Scale: billions of parameters, terabytes of data —turned “label this” into “generate that”.

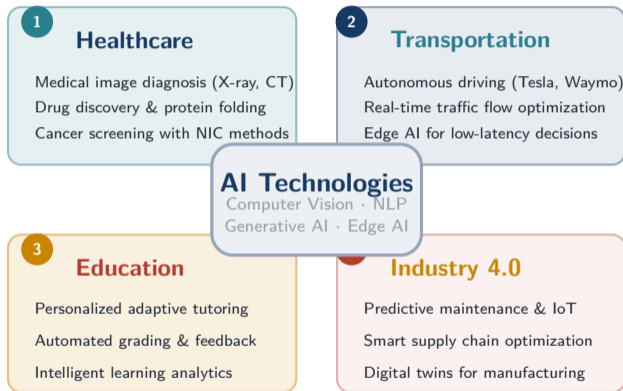
▷ **Think:** Which AI have you used this week? Traditional (Netflix, search) or generative (ChatGPT, Copilot)? Probably both.

▷ *Now: where are these technologies applied in the real world?*

Related: Wang et al., Demonstration Selection for In-Context Learning via Reinforcement Learning, ICML 2025 (20 cites)

AI in the Real World

AI technologies —from computer vision to LLMs— are transforming every sector.
Four key domains (speak clockwise):



▷ **Think:** Which area excites you most? *For me —healthcare. AI-assisted diagnosis saves lives. Every domain here is being transformed right now.*

Summary: What We Learned Today

1. **AI** is machines that perceive, reason, learn, and act

Not copying the brain —achieving rational outcomes through computation

2. These are powered by **three pillars** advancing together

Data, algorithms, and compute form a virtuous cycle —each amplifies the others

3. **Machine learning** is simply finding patterns in data

Three paradigms —supervised, unsupervised, RL —all share one loop: data → predict → feedback → improve

4. We've shifted from **recognizing** things to **creating** them

Same neural networks at massive scale —output changed from labels to content (DALL-E, ChatGPT, Copilot)

5. AI is already transforming **our world** right now

Healthcare, transportation, education, industry —from diagnosis to autonomous driving

Bottom line: AI is not magic —it's mathematics, data, and engineering working together.

That's our quick tour of AI fundamentals. Questions? Let's discuss how AI connects to **your** field of interest.

Thank You

Questions & Discussion

Xubin Wang · wangxubin@ieee.org · www.wangxubin.site

Q&A topics: teaching philosophy · engaging students · NIC / feature selection / edge AI · AI across disciplines