

What is AI, really?

Definitions, the Turing test, the Chinese Room, and 80 years of summers and winters.

01 COLD OPEN 0-10 min

AI or not? You decide

AI or not? Vote as a class, then click. Score: 0 / 0 answered

A calculator app adding your CT marks

Google Maps rerouting you around Mirpur 10 traffic

Gmail pushing an email to spam

An AC that switches on above 26 °C

Deep Blue beating Kasparov at chess, 1997

Face unlock on your phone

An Excel sheet that turns marks into CGPA

A washing machine with a "fuzzy logic" sticker


A chatbot writing your assignment for you

02 CONCEPT 10-18 min

Four ways to define AI

Two axes: do we care about **thinking** or **acting**, and do we measure against **humans** or an ideal of **rationality**?

Four ways to define AI Click a quadrant.



Acting rationally

Do whatever maximises the expected outcome, given what you perceive and know. This is the approach of this course and of Russell & Norvig.

Examples: Self-driving cars, AlphaGo, route planners, insulin pumps.

On screen: JARVIS: never "human", always useful.



PRO TIP

Rational $\neq$ emotionless robot. It means doing the right thing given what you know. A rational agent can still be wrong if its information was wrong.

The imitation game

Human or bot? Which reply came from a machine? Vote, then click.

"What did you have for breakfast?"

Paratha and egg. The paratha was too oily, but I ate two anyway.

I don't eat, but a typical Bangladeshi breakfast might include paratha with egg or bhaji.

"Write one line about rain in Dhaka."

The monsoon rain washes the streets of Dhaka, bringing relief and renewal to the city.

Rain in Dhaka: twenty minutes of poetry, three hours of traffic.

"What is $7,493 \times 28$?"

uhh one sec... around 210k? 209,804?

209,804

1950: Turing swaps the fuzzy "Can machines think?" for a test. A judge chats by text with a hidden human and a hidden machine. Can't tell which is which? The machine passes.




POP CULTURE

Ex Machina is a Turing test with the stakes turned up. Its twist: passing the test says nothing about whether the machine shares your goals. Hold that thought for Week 12.

Passing isn't understanding?

The Chinese Room

The Chinese Room A note goes in, a perfect reply comes out. Does anyone inside understand Chinese?



THINK DEEPER

Today's chatbots arguably pass short Turing tests. Is an LLM just a very large Chinese Room?

Two defensible positions. **Yes (Searle):** the model maps input tokens to output tokens by learned rules; no understanding is required to do that. **No / it depends (systems reply):** understanding may be a property of the whole system, not any one part. The rational-agent view side-steps the debate: we ask what the system reliably *does*.

A STRONG ANSWER MENTIONS

- Searle: syntax (symbol shuffling) is not semantics (meaning).
- "Systems reply": the person doesn't understand, but maybe the whole room does.
- Practical view (this course): judge agents by their performance, and test them where it matters.

05 **CONCEPT** 31-35 min

AI is built from eight other subjects

PHILOSOPHY

Can reasoning be written down as rules? Where does knowledge come from?

[Logic & expert systems](#) · [W2](#)

MATHEMATICS

Logic, probability, what is computable and how fast.

[Search](#) · [W3-4](#), [Bayes](#) · [W8](#)

ECONOMICS

Utility, decisions under uncertainty, game theory.

[Utility agents](#) · [W1](#), [Minimax](#) · [W5](#)

NEUROSCIENCE

How do networks of simple neurons produce thought?

[Neural networks](#) · [W11](#)

PSYCHOLOGY

How do humans perceive, learn and decide?

COMPUTER ENGINEERING

Hardware to run it all: vacuum tubes to GPUs.

CONTROL THEORY

Feedback loops that keep a system on target.

Agents · next class

LINGUISTICS

How does language carry meaning?


NLP & LLMs · W10, W12

06 **LIVE LAB** 35-44 min

80 years in one scroll

80 years of AI, summers and winters

Click a dot, or use the buttons. The curve shows the mood of the field, not measured data.



2024 AI wins Nobel prizes

Physics to Hopfield and Hinton for neural networks; Chemistry partly to Hassabis and Jumper for AlphaFold.

07 **CONCEPT** 44-47 min

Narrow AI, general AI

Narrow AI

Superhuman at one task, useless outside it. A Go engine cannot fold laundry. Every deployed system you use today is narrow, including very capable chatbots (text in, text out).

General AI (AGI)

Human-level ability across almost any task, learning new ones quickly. Whether and when it arrives is one of the most argued questions in the field. We revisit it after we understand how LLMs work.

Worked example: is a smart traffic light AI?

GIVEN

An adaptive traffic signal at Farmgate, Dhaka.

FIND

Is it an AI system? Which of the four definitions fits? Use the four-question test.

Method: Russell & Norvig Ch. 1-2 (definitions, agents)

SOLUTION

1. **Does it perceive its environment?** Yes: cameras or loop sensors count the cars waiting in each direction.
2. **Does it choose between actions?** Yes: which direction gets green, and for how long. A fixed 90-second timer would not choose.
3. **Is the choice driven by a goal or performance measure?** Yes: minimise total waiting time, perhaps weighted toward buses and ambulances.
4. **Does it reason, plan or learn beyond a fixed lookup?** "More cars → longer green" alone is a simple reflex agent: AI in the textbook sense, barely. Predicting Mirpur traffic 20 minutes ahead and planning the whole corridor is clearly AI.
5. **Which quadrant?** Acting rationally: nobody cares whether it thinks like a traffic sergeant, only whether traffic moves.

| This four-question test is a safe structure for any "Is X an AI system? Justify." CQ.

ANSWER

It is an AI agent in the acting-rationally sense; how much AI depends on whether it plans ahead or only reacts.

Your turn • 1: a Bangla voice assistant

1 EXAM CQ

A Bangla voice assistant on your phone sets alarms and answers questions.

Is it AI? Justify with the four-question test and name its quadrant.

SOLUTION

ANSWER

Perceives (microphone), chooses (which action or answer), has a goal (satisfy the request), learns (speech models). Acting humanly for the conversation, acting rationally for the tasks. Clearly AI.

Your turn • 2: two quick verdicts

2a WARM-UP

A washing machine that always runs a 40-minute cycle when you press Start.

SOLUTION

ANSWER

No perception of the load, no choice, fixed sequence: **not AI**. A fuzzy controller that senses load and dirt and picks the time would be (Week 5).

2b WARM-UP

A diagnosis assistant that reads symptoms and lab values and ranks dengue, malaria and chikungunya by probability.

SOLUTION

ANSWER

Perceives patient data, chooses a ranked differential, goal = correct diagnosis, reasons with probability. Thinking rationally (Bayes) in service of acting rationally (useful advice). A real system: see the case slide.

Critical thinking • the AI effect

THINK DEEPER

In 1997 beating the chess champion was 'AI'. Today a free phone app does it and nobody calls it AI. Why does every solved problem stop counting?

This is the **AI effect**. Understanding the mechanism removes the mystery, and "intelligence" is often used to mean "whatever is still mysterious". That is why this course uses a *precise* definition (rational agents) instead of a vibe.

A STRONG ANSWER MENTIONS

- Once we know *how* it works (search, rules), it feels like "just computation".
- The goalposts move to whatever machines still can't do (Rich & Knight's definition bakes this in).
- Consequence: judge by capability and method, not by the label.

THINK DEEPER

Your friend says: 'If it doesn't use machine learning, it isn't real AI.' Give two counter-examples from today.

Deep Blue (search), MYCIN (rules), Google Maps' routing core (A* search on a graph). ML is one branch of AI, not the whole tree.

A STRONG ANSWER MENTIONS

- Deep Blue: search + hand-tuned evaluation, almost no learning.
- MYCIN / expert systems: hand-written rules.
- Route planners (A*), schedulers, theorem provers.

Critical thinking • design

α Turing test for 2026

3

DESIGN

Chat-only tests are now easy to pass. Propose **three changes** to the Turing test so it measures something closer to intelligence in 2026, and one risk your new test creates.

SOLUTION

ANSWER

Strong answers change *what* is tested, not just length: e.g. (1) tasks needing persistent memory over days; (2) physical-world grounding (describe what is on your desk right now, verified by camera); (3) consistency under adversarial follow-ups; (4) admitting uncertainty honestly. And they note the danger: a test that rewards *imitating* humans may reward deception.

13 REAL WORLD 67–72 min

From our lab: AI for tropical-fever diagnosis



FROM OUR LAB

Dengue, malaria and chikungunya look alike in the first days: fever, aches, fatigue. In low-resource clinics lab tests are slow or unavailable. At BUET's mHealth Lab, your lecturer's team built **mLabLLM**: a fine-tuned LLaMA 3.2 language model combined with Bayesian reasoning, reaching 82.8% top-3 accuracy for these three diseases (AAAI 2025).

Why it matters today: a plain chatbot is "acting humanly" and may confidently invent answers. Adding explicit probability moves it toward "thinking rationally".

You will learn both halves in this course: Bayes in Week 8, LLMs in Week 12.

14 POP QUIZ 72–78 min

Pop quiz

1. Which approach to AI does this course (and Russell & Norvig) mainly follow?

- A. Thinking humanly
- B. Acting humanly
- C. Thinking rationally

D. Acting rationally ✓

The rational-agent view: do the action that maximises the expected performance measure.

2. The Turing test is a test of...

- A. Thinking humanly
- B. Acting humanly ✓**
- C. Thinking rationally
- D. Acting rationally

It only checks behaviour: can the machine pass as a human in conversation?

3. Searle's Chinese Room argues that...

- A. Machines can never pass the Turing test
- B. Producing the right symbols is not the same as understanding them ✓**
- C. Chinese is too hard for computers
- D. Rulebooks are always incomplete

The room behaves perfectly yet nobody inside understands Chinese. Behaviour (syntax) is not meaning (semantics), says Searle.

4. Deep Blue (1997) mostly used...

- A. Deep learning
- B. Search with a hand-made evaluation function ✓**
- C. Large language models
- D. Genetic algorithms

Massive game-tree search plus expert-tuned evaluation. AI is not only machine learning.

5. MYCIN (1970s) was an early...

- A. Chess program
- B. Chatbot
- C. Rule-based medical expert system ✓**
- D. Self-driving car

It diagnosed blood infections using hundreds of if-then rules with certainty factors. We build a mini version next week.

15 **RECAP** 78-80 min

Three things to remember

01

Four definitions: thinking or acting, humanly or rationally. This course takes the rational-agent view.

02

Behaving intelligently $\neq$ understanding (Chinese Room). AI is bigger than ML: search, logic, probability, planning and learning all count.

03

History runs in hype cycles. Next class: what exactly an agent is, and PEAS.

HW

BEFORE NEXT CLASS

Find one AI system used in Bangladesh (an app, a bank, a hospital, a farm). In five lines: what it perceives, what it chooses, its goal, and which quadrant it belongs to. Bring it: we use your examples next class.

SOLUTION

ANSWER

Any real system works (bKash fraud detection, Pathao dispatch, a hospital triage tool, an AgriTech crop-disease app). Full marks need all four parts *and* a quadrant with a one-line justification.

S. M. Sakeef Sani

Adjunct Lecturer, CSE · Northern University Bangladesh

smsakeefsani.hbitcorp.com