



CSC460/660

Artificial Intelligence: Principles & Techniques

Faisal N. Abu-Khzam

Department of Computer Science & Mathematics

Lebanese American University

Beirut, Lebanon



Overview

- Course policy
- Topics covered
- Introductory material
- Topics for Presentations & Projects

Course policy

- **Attendance is mandatory**
- **Electronic devices OFF or silent+hidden**
- **Grades are based on:**
 - Exams I and II (65-70%)
 - Topic presentations (grads only, 10%)
 - Projects (20-25%)
 - HW & participation (5%)
- **Topic/project selection should start now**
 - A possible list of topics will be given in the sequel
- **Office hours:** mainly by appointment; suggested slots: Monday 1:30 – 3:30pm
 - Other possible time slots: Tuesday/Wednesday 12-3pm (for appointments, sometimes via zoom)

Course material

- Part I - Discussion
 - From history to state of the art
 - What's next?
 - Effect on other disciplines/education/jobs/society
 - The use of HPC
 - Mandatory? The **Shuffle & Check** phenomenon
 - Is lightweight AI possible?
- The first few chapters of Russel & Norvig
 - Intelligent Agents
 - Search problems
 - Local search and optimization
 - Adversarial Search and Games
 - Constraint Satisfaction Problems
 - Logical Agents
 - First-Order Logic and Inference Algorithms
- Learning from examples (mainly decision trees)
- LLMs

Topics for presentations & projects

- GNNs
- Fuzzy logic
- Knowledge graphs
- Game Theory and AI
- Explainable AI (XAI)
- Reinforcement Learning
- Speech Recognition
- Robotics and Autonomous Navigation
- Image classification & Computer vision
- NLP topic (e.g. sentiment analysis / opinion mining)

The Turing Test

- Turing (1950): replace “Can machines think?” with an operational test of human-like conversational behavior.
- A question for 2026:

Is passing a conversational test still a convincing test of intelligence?

What is Artificial Intelligence?

- Several definitions...
- Four classical views:
 - Thinking humanly
 - Acting humanly
 - Thinking rationally
 - Acting rationally
- Which one best describes what we call AI today?

What is Artificial Intelligence?

- Four classical views:
 - Thinking humanly
 - Acting humanly
 - Thinking rationally
 - Acting rationally
- Which one best describes what we call AI today?
- In 2016 I would have strongly argued for “Acting rationally.” What would you choose in 2026?

What is Artificial Intelligence?

Why “thinking rationally” matters more now?

- LLMs have largely **changed the bottleneck**.
 - Producing impressively human-like behavior is no longer the extraordinary challenge it once seemed.
 - A model can converse, write, translate, summarize and imitate human styles extremely convincingly.
- But it can also produce a beautifully written **wrong answer!**

What is Artificial Intelligence?

- So we have moved from asking:
 - *Can a machine produce intelligent-looking behavior?*
- Toward asking:
 - *Can it actually reason reliably enough that we can trust how it arrives at conclusions?*

What is Artificial Intelligence?

- *Can it actually reason reliably enough that we can trust how it arrives at conclusions?*

That is why current work emphasizes **reasoning models**, inference-time reasoning/computation, chain-of-thought-like internal deliberation, verification, tool use, planning, neuro-symbolic approaches, and formal reasoning.

Acting rationally remains the broader objective. But as AI systems become more autonomous and are entrusted with increasingly consequential tasks, rational action may require reliable rational reasoning.



Before Dartmouth – The Work of McCulloch and Pitts (1943)

- The paper connected three previously separate ideas: neuroscience, mathematical logic and computation.
- It showed how reasoning-like logical operations could arise from networks of extremely simple elements.
- This was an early bridge from “the mind” as an abstract object to intelligence as something that could be represented and computed.

One of the first clear demonstrations that reasoning could be treated as a mechanical process rather than something mystical, which set up the core AI idea that thinking can be computed.

Before Dartmouth – Wiener connects feedback, control & communication (1948)

- Norbert Wiener brought together control, communication and feedback under the idea of cybernetics.
- An intelligent system need not simply compute once: it can observe the consequences of its actions and correct itself.
- Feedback turns behavior into a loop: sense → compare with a goal → act → sense again.


Anticipates the modern AI agent interacting continuously with an environment.

Before Dartmouth – Turing studies machine intelligence through conversation (1950)

- Turing replaced the vague question “Can machines think?” with an operational question based on observable behavior.
- But the paper went much further: it discussed learning machines, objections to machine intelligence, and how intelligent behavior might be produced computationally.
- The question was no longer only philosophical: it became something that could be investigated experimentally.

The imitation game became what we now call the Turing Test.

Before Dartmouth – Digital Computers (by 1955-56)



The Birth of Artificial Intelligence

- By the early 1950s, several communities were already attacking pieces of what we now call AI: neural networks, automata, cybernetics, information theory, game playing, mathematical logic, and learning.
- In 1955, John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon wrote a funding proposal for a summer research project to be held at Dartmouth in 1956.
- The proposal was titled *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*.

Their central conjecture:

“Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”

Dartmouth's aftermath

- The workshop gave a research field a name and a community
- Its institutional influence exceeded its immediate results
- 1958: McCarthy introduced Lisp for symbolic computation
- 1960s: expert systems and language programs developed

Short History

- 1950 — Turing and the imitation game
 - 1956 — Dartmouth: Artificial Intelligence becomes a field
 - 1970s–80s — Symbolic AI, expert systems — and AI winters
 - 1997 — Deep Blue defeats the world chess champion
 - 2012 — Deep learning breakthrough (AlexNet)
 - 2016 — AlphaGo defeats Lee Sedol
 - 2017 — Transformer architecture
 - 2022–Generative AI and large language models
- 

What Can AI Do Today?

- Understand and generate language
- Recognize and generate images, audio and video
- Write and analyze computer programs
- Play complex games and support planning/optimization
- Assist scientific and medical research
- Control increasingly autonomous systems

- But capability is not the same as understanding...

But Is This Intelligence?

- A modern AI system may:
 - answer difficult questions
 - write convincing text
 - solve many programming problems
 - explain mathematical ideas
- ...and still produce a confident, completely wrong answer.
- What exactly do we mean by “intelligent”?

AI = Algorithms + Data + Computing Power?

- Three major drivers of modern AI:
 - Better algorithms and models
 - Enormous amounts of data
 - Massive computing power
- How much of recent progress comes from each?
- Can brute computational power sometimes look like intelligence?

Why Study Classical AI in 2026?

- Search → choosing among possible actions
 - Heuristics → using knowledge to search intelligently
 - Games → reasoning against another agent
 - Constraints → reasoning under restrictions
 - Logic → representing knowledge and deriving conclusions
 - Learning → acquiring patterns from examples
-
- These are fundamental ways of solving intelligent tasks — not merely “old AI.”

The Road Through This Course

- INTELLIGENT AGENTS



- SEARCH & HEURISTICS



- OPTIMIZATION • GAMES • CONSTRAINTS



- LOGIC & INFERENCE



- LEARNING FROM EXAMPLES



- LARGE LANGUAGE MODELS

Before We Go Further...


- Two questions to keep in mind throughout the course:
 - Does producing an intelligent-looking result imply an intelligent method?
 - If enough computing power can compensate for a poor method, should we still call the resulting system intelligent?
- This brings us to algorithms, HPC, and:
- The “Shuffle & Check” phenomenon...

After the break: agents and search

- Chapter 2: Intelligent Agents
- Chapter 3: Solving Problems by Searching
- Formulating problems and tracing search algorithms
- Comparing guarantees, memory costs and heuristic guidance

Agents and rationality

- An agent perceives its environment and selects actions
- A performance measure defines success
- Rational action maximizes expected performance given the evidence
- Rationality depends on information and available actions
- Important: rationality not the same thing as omniscience



PEAS

- Performance measure, Environment, Actuators, Sensors
- Must first specify the setting for intelligent agent design
- Consider, e.g., the task of designing an automated taxi driver:
 - Performance measure
 - Environment
 - Actuators
 - Sensors


PEAS

- Performance measure, Environment, Actuators, Sensors
- Must first specify the setting for intelligent agent design
- Consider, e.g., the task of designing an automated taxi driver:
 - Performance measure: Safe, fast, legal, comfortable trip, maximize profits
 - Environment: Roads, other traffic, pedestrians, customers
 - Actuators: Steering wheel, accelerator, brake, signal, horn
 - Sensors: Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard


Reading

- Read about various types of agent environment as classified by Russel & Norvig
- Read about notion of an agent function and programs
- Read the nice note about table-lookup agents
- Read about types of agents (summarized next...)


Simple Reflex Agent




Model Based Agent



Goal Based Agent



Utility Based Agent



Learning Agent

