

Foundations of Artificial Intelligence

Chapter 01: Introduction to A.I.

Roberto Sebastiani

DISI, Università di Trento, Italy – `roberto.sebastiani@unitn.it`

`https://disi.unitn.it/rseba/DIDATTICA/fai\_2026/`

Teaching assistant: **Nicola Debole**, `nicola.debole@unitn.it`,

M.S. Course “Artificial Intelligence Systems”, academic year 2026-2027

Last update: Thursday 10th September, 2026, 17:41

Copyright notice: Most examples and images displayed in the slides of this course are taken from [Russell & Norvig, “Artificial Intelligence, a Modern Approach”, 3rd and 4th ed., Pearson], including explicitly figures from the above-mentioned books, so that their copyright is detained by the authors. A few other material (text, figures, examples) is authored by (in alphabetical order): Pieter Abbeel, Bonnie J. Dorr, Anca Dragan, Dan Klein, Nikita Kitaev, Tom Lenaerts, Michela Milano, Dana Nau, Maria Simi, who detain its copyright.

These slides cannot be displayed in public without the permission of the author.

Outline

- 1 AI: Fiction vs. Reality
- 2 What is AI?
- 3 AI: State of the Art

Outline

1 AI: Fiction vs. Reality

2 What is AI?


3 AI: State of the Art

There is plenty of AI in fiction ...



“Metropolis”, 1927, by Fritz Lang

AI in Fiction



"2001, Space Odyssey", 1968, by Stanley Kubrick

AI in Fiction




“Star Wars”, 1977, by George Lucas

AI in Fiction



“Blade Runner”, 1982, by Ridley Scott

AI in Fiction




“Terminator”, 1984, by James Cameron

AI in Fiction



“A.I., Artificial Intelligence”, 2001, by Steven Spielberg

AI in Fiction




“I, Robot”, 2004, by Alex Proyas

AI in Fiction



"Wall-E", 2008, by Andrew Stanton

AI in Fiction



“Ex Machina”, 2015, by Alex Garland

AI in Fiction



“Blade Runner, 2049”, 2017, by Denis Villeneuve

... and many others ...

(see, e.g.,

<https://www.looper.com/198685/the-stunning-evolution-of-ai-in-movies/>)

Many AI fantasies from fiction have become or are becoming reality ...

... autonomous vacuum cleaners, ...



... self-driving cars, ...



... autonomous service vehicles, ...



... autonomous flying vehicles, ...



AI in Reality

... soccer-playing robots, ...



AI in Reality

.. acrobatic humanoid robots, ...



AI in Reality

... autonomous trading bots, ...



..., vocal assistants, ...



AI in Reality

... image/face recognition tools, ...



AI in Reality

... world-champion beating chess players, ...



... world-champion beating go players, ...




AI in Reality

... AI fighter pilots, ...



AI in Reality

... LLM-based chatbots, ...















... Image/Video Synthesis ...



Tilly Norwood, the actress who does not exist (©RollingStone, 2025)

... Automated Code Synthesis ...

```
> /context
└─ 
     Context Usage
     claude-opus-4-1-20250805 • 120k/200k tokens (60%)
    
      System prompt: 3.1k tokens (1.5%)
      System tools: 11.9k tokens (6.0%)
      MCP tools: 3.4k tokens (1.7%)
      Memory files: 1.1k tokens (0.6%)
      Messages: 100.6k tokens (50.3%)
      Free space: 79.9k (40.0%)

MCP tools • /mcp
└─ mcp__postgres__query (postgres): 373 tokens
└─ mcp__puppeteer__puppeteer_navigate (puppeteer): 490 tokens
└─ mcp__puppeteer__puppeteer_screenshot (puppeteer): 506 tokens
└─ mcp__puppeteer__puppeteer_click (puppeteer): 395 tokens
└─ mcp__puppeteer__puppeteer_fill (puppeteer): 414 tokens
└─ mcp__puppeteer__puppeteer_select (puppeteer): 419 tokens
└─ mcp__puppeteer__puppeteer_hover (puppeteer): 396 tokens
└─ mcp__puppeteer__puppeteer_evaluate (puppeteer): 393 tokens

Memory files • /memory
└─ User (/Users/jngiam/.claude/CLAUDE.md): 32 tokens
└─ Project (/Users/jngiam/mcp-playground/CLAUDE.md): 1.1k tokens
```

... and many others ...

Outline

1 AI: Fiction vs. Reality

2 What is AI?

3 AI: State of the Art

What is AI for the Media

- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- LLMs (ChatGPT, Gemini, Claude, ...)
- ...

Intelligence vs. Artificial Intelligence

Intelligence

For thousands of years, we have tried to **understand** how we think:

- how can a “handful of matter” **perceive**, **understand**, **predict**, and **manipulate** a world far larger and more complicated than itself?
- involves many disciplines, including **logic**, **psychology**, **cognitive science**, **neuroscience**, **philosophy**, **ethics**, **linguistics**, ...

Artificial Intelligence

The field of **Artificial Intelligence (AI)** goes further still:

- it attempts not just to **understand**, but also to **build** intelligent entities
- involves all the above disciplines, but also **mathematics**, **computer science**, **engineering**, **economics**, **control theory & cybernetics**, **electronics**, ...



What is Intelligence?

Intelligence (from Wikipedia)

“(...) *it can be described as the ability to **perceive** or **infer information**, and to **retain** it as **knowledge** to be **applied** towards **adaptive behaviors** within an **environment** or **context**.* (...)”

What is Intelligence? [cont.]

Example: simple puzzle



(Courtesy of Michela Milano, UniBO)

- What is the solution of this puzzle?
⇒ (I'd say) **B**: result of column-by-column clockwise rotation
- What have you done for solving it?
 - 1 read & recognize figures ⇒ perceive information
 - 2 recognize patterns, problem and candidate solutions ⇒ retain knowledge
 - 3 choose solution ⇒ infer other knowledge

What is Artificial Intelligence?

Different definitions due to different criteria

Historically, four approaches, along two orthogonal dimensions:

- thought processes & reasoning (“thinking”)

vs.

behavior & action (“acting”)

- Success according to human standards

vs.

success according to an ideal concept of intelligence: rationality.

- **human-centered approach:** involves observations and hypotheses about human behavior
- **rationalist approach** involves a combination of mathematics and engineering.

The four groups have both disparaged and helped each other.

What is Artificial Intelligence? [cont.]

Thinking Humanly “The exciting new effort to make computers think . . . <i>machines with minds</i>, in the full and literal sense.” (Haugeland, 1985) “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)	Thinking Rationally “The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)
Acting Humanly “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)	Acting Rationally “Computational Intelligence is the study of the design of intelligent agents.” (Poole <i>et al.</i>, 1998) “AI . . . is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)

Thinking Humanly: The cognitive modeling approach

Problem: How do humans think?

- Idea: develop a **theory of the mind**
 - ⇒ express the theory as computer programs
 - e.g. Newell & Simon's **General Problem Solver** (1961)
 - Requires scientific theories of brain activities (**cognitive model**)
 - Inter-disciplinary field: **Cognitive Science**
 - combines computer models from AI and experimental techniques from psychology
 - construct precise and testable theories of the human mind
 - AI and Cognitive Science nowadays distinct
 - A.I: find an algorithm performing well on a task
 - C.S: find a good model of human performance
- although they fertilize each other (e.g. in computer vision)

Acting Humanly: The Turing Test Approach

Problem: When does a system behave intelligently?

The Turing Test

- Alan Turing “Computing Machinery and Intelligence” (1950)
- Operational test of intelligence (aka “The Imitation game”):
 - A human, a computer, an interrogator in a different room.
 - The interrogator should classify the human and the machine.
 - Can the computer mislead the interrogator and be classified as a human?
- “behave intelligently” $\iff$ “behave humanly”

Acting Humanly: The Turing Test Approach

Problem: When does a system behave intelligently?

The Turing Test

- Alan Turing “Computing Machinery and Intelligence” (1950)
- Operational test of intelligence (aka “The Imitation game”):
 - A human, a computer, an interrogator in a different room.
 - The interrogator should classify the human and the machine.
 - Can the computer mislead the interrogator and be classified as a human?
- “behave intelligently” $\iff$ “behave humanly”




Acting Humanly: The Turing Test Approach

Problem: When does a system behave intelligently?

The Turing Test

- **Alan Turing** “Computing Machinery and Intelligence” (1950)
- Operational test of intelligence (aka “**The Imitation game**”):
 - A human, a computer, an interrogator in a different room.
 - The interrogator should classify the human and the machine.
 - Can the computer mislead the interrogator and be classified as a human?
- “behave intelligently” $\iff$ “behave humanly”




Acting Humanly: The Turing Test Approach

Problem: When does a system behave intelligently?

The Turing Test

- Alan Turing “Computing Machinery and Intelligence” (1950)
- Operational test of intelligence (aka “The Imitation game”):
 - A human, a computer, an interrogator in a different room.
 - The interrogator should classify the human and the machine.
 - Can the computer mislead the interrogator and be classified as a human?
- “behave intelligently” $\iff$ “behave humanly”



Acting Humanly: The Turing Test Approach [cont.]

Capabilities for passing the Turing Test

- **natural language processing** to enable it to communicate successfully in English (or other)
- **knowledge representation** to store what it knows or hears
- **automated reasoning** to answer questions and to draw new conclusions
- **machine learning** to adapt to new circumstances and to detect and extrapolate patterns

For **Total Turing test** (with physical interaction wrt. interrogator):

- **computer vision** to perceive objects
- **computer speech** to communicate orally
- **robotics** to manipulate objects and move about

- These disciplines compose most of AI
- Turing Test is still relevant in AI (although not fundamental)

Acting Humanly: The Turing Test Approach [cont.]

Some successes with Turing test

- (2014) a chatbot by Eugene Goostman, mimicking the answer of a 13 years old boy, has succeeded the test.
 - chatbots are now frequently available
 - recent boost with LLM-based chatbots (e.g. ChatGPT)
- vocal assistants are now of common use
 - e.g. Alexa (Amazon), Siri (Apple), Cortana (Microsoft), ...
- Currently LLM-generated text/images/videos hard to distinguish from real ones


Limitations of Turing Test

- Not reproducible, constructive or amenable to mathematical analysis
- AI researchers devoted little effort to make systems pass the Turing Test
- Do humans always pass the Turing test?
- Should we really emulate humans to achieve intelligence?
- Shouldn't we study the underlying principles of intelligence instead?

Acting Humanly: The Turing Test Approach [cont.]

Metaphorical Example

Successful flight machines have not been developed by imitating birds, rather by studying engines and aerodynamics.



(see e.g. [this video](#)).

Thinking Rationally: The “Laws of Thought” Approach

Problem: Can we capture the laws of thought?

- Aristotle: What are **correct** argument and thought processes?
 - codify “right thinking” i.e. irrefutable reasoning processes (**sylogisms**):
(e.g. “all men are mortal; Socrates is a man; therefore, Socrates is mortal”)
⇒ **Logic** and **Logical inference**
- The **Logicist tradition** in AI hopes to create intelligent systems using logic-based inference systems
 - “**algorithm = logic + control**”
 - logic programming, automated-deduction systems, ...
 - logics: propositional, first-order, modal & decription, temporal, ...
- Two main limitations:
 - not easy to state informal knowledge into the formal terms of logic
 - problems undecidable or computationally very hard (NP-hard)
- **Logical reasoning** is currently part of many fields of AI
 - **problem solving, knowledge representation & reasoning, planning,**
 - **does not exhaustively cover AI**

Acting Rationally: The Rational-Agent Approach

Problem: Can we make systems “do the right thing”?

Rational Agents

- An **agent** is an entity that **perceives** and **acts**
 - persists over a prolonged time period
- A **rational agent** acts so as
 - **to achieve the best outcome** (maximize goal achievement),
or
 - **to achieve the best expected outcome** (under uncertainty)
- Rational agents need all skills needed for the Turing Test!
- Thinking rationally is **most often (but not always)** part of being a rational agent
 - e.g. planning an action
 - sometimes action without thinking (e.g. reflexes)
- Two advantages over previous approaches:
 - **More general than law of thoughts approach**
(correct inference is just one of several possible mechanisms for achieving rationality)
 - **More amenable to scientific development than human-emulation approaches**
(rationality is mathematically well defined & general)

Acting Rationally: The Rational-Agent Approach [c.]

This course concentrates on **general principles of rational agents** and on the **components for constructing them**. (Following AIMA book.)

Remark

- achieving **perfect rationality** is not feasible in complex environments
 - computational demands too high
 - however, good working hypothesis and starting point for analysis

⇒ dealing with **limited rationality**

- acting appropriately when not enough time to do all computations

AI Systems Classification

Weak vs. Strong AI

- **Weak AI**: Is it possible to build systems that **act as if they were intelligent**?
- **Strong AI**: Is it possible to build systems that **are intelligent**?
(i.e., that have conscious minds, wills and sentiments?)

General AI vs. Narrow AI

- **Artificial General Intelligence (AGI)** refers to systems able to cope with **any generalized task** which is asked of it, much like a human.
- **Artificial Narrow Intelligence (ANI)** refers to systems able to handle one particular task.
⇒ AI system displays a certain degree of intelligence only in a particular narrow field to perform **highly specialized tasks**

AI Systems Classification [cont.]

Symbolic Approach vs. Connectionist Approach

- **Top-down, or Symbolic Approach:**
 - Symbolic representation of knowledge
 - Logics, ontologies, rule based systems, declarative architecture
 - Human-understandable models
- **Bottom up, or Connectionist Approach:**
 - Based on Neural networks.
 - Knowledge is not symbolic and it is “encoded” into connections between neurons.
 - Concepts are learned by examples
 - Non understandable by humans

Current trend in AI research

Combine Symbolic with Connectionist Approaches (**Neuro-Symbolic AI**)

⇒ try to get the best out of the two approaches

If interested, see the section “History of AI” in AIMA book.

The Foundations of Artificial Intelligence

Overall, different fields have contributed to AI in the form of **ideas**, **viewpoints** and **techniques**

- **Philosophy**: Logic, reasoning, mind as a physical system, foundations of learning, language and rationality
- **Mathematics**: Formal representation and proof, computation, (un)decidability, (in)tractability, probability
- **Economics**: formal theory of rational decisions, game theory
- **Neuroscience**: physical substrate for mental activities
- **Psychology**: adaptation, phenomena of perception and motor control
- **Computer Science & Engineering**: algorithms, data structures, efficient implementations
- **Control Theory & Cybernetics**: homeostatic systems, stability, optimal agent design
- **Linguistics**: knowledge representation, grammar

Main AI Research Venues

- Major AI Journals

- Artificial Intelligence
- Computational Intelligence
- Journal of Artificial Intelligence Research
- IEEE Transactions on Pattern Analysis and Machine Intelligence
- IEEE Intelligent Systems
- [area-specific journals]

- Main AI Conferences

- International Joint Conference on AI (IJCAI)
- AAAI Conference on AI (AAAI)
- European Conference on AI (ECAI)
- [area-specific conferences]

- Main professional societies for AI

- The Association for the Advancement of Artificial Intelligence (AAAI)
- European Association for Artificial Intelligence (EurAI)
- Asia-Pacific Artificial Intelligence Association (AAIA)
- ACM Special Interest Group in Artificial Intelligence (SIGART)
- Society for Artificial Intelligence and Simulation of Behaviour (AISB)

Outline

- 1 AI: Fiction vs. Reality
- 2 What is AI?
- 3 AI: State of the Art

AI is everywhere ...

- Search engines
- Route planning (e.g. maps, traffic)
- Logistics (e.g. packages, inventory, airlines)
- Medical diagnosis, machine diagnosis
- Automated help desks
- Spam/fraud detection
- Smarter devices, e.g. cameras
- Product recommendations
- Assistants, smart homes
- Text generators
- Image & video generators
- ...

Lots more!

What can AI Systems Currently Do?

... classify incoming e-mails as spam (or not), ...



<http://www.resilientsystems.co.uk/>

What can AI Systems Currently Do?

... predict stock price evolution, ...

Apple Inc. (NASDAQ:AAPL)

Add to portfolio

More results

105.67 -0.46 (-0.43%)

Mar 24 - Close
NASDAQ real-time data - Disclaimer
Currency in USD

Range 104.89 - 106.25
52 week 92.00 - 134.54
Open 105.47
Vol / Avg 26.13M / 25.90M
Mkt cap 583.36B
P/E 11.22
Div/yield 0.52 / 1.97
EPS 9.41
Shares 5.54B
Beta 0.97
Inst. own 59%

G+I 9.3K

Dow Jones 17,515.73 0.08%
Nasdaq 4,773.50 0.10%
Technology 0.18%
AAPL 105.67 -0.43%



Compare: Enter ticker here Add Dow Jones Nasdaq SNDK MSFT SSNNF VZ HPQ IBM HTCKF

Zoom: 1d 5d 1m 6m 1y 10y All
Mar 24, 2016 - Mar 24, 2016 -0.46 (-0.43%)



Settings | Technicals | [Link to this view](#)

Volume delayed by 15 mins.
Prices are not from all markets.
Sources include SIX.

What can AI Systems Currently Do?

... understanding handwriting, ...

80322-4129 80206

40004 14310

37879 05153


~~33~~02 75216

35460 44209

[LeCun et al. 1989]

What can AI Systems Currently Do?


... learn to grab a cup, ...



<http://www.informatik.uni-bremen.de/>

What can AI Systems Currently Do?

... design a molecule with given properties, ...



<http://pande.stanford.edu/>

What can AI Systems Currently Do?

... translate text from Chinese to English, ...



©Google Inc.


What can AI Systems Currently Do?

... convert a voice into text, ...



What can AI Systems Currently Do?

... predict traffic trajectories, ...



What can AI Systems Currently Do?

... automatically writing the caption of a figure, ...



"man in black shirt is playing guitar."



"construction worker in orange safety vest is working on road."



"two young girls are playing with lego toy."



"boy is doing backflip on wakeboard."



"girl in pink dress is jumping in air."



"black and white dog jumps over bar."



"young girl in pink shirt is swinging on swing."



"man in blue wetsuit is surfing on wave."

[Karpathy & Fei-Fei, 2015; Donahue et al., 2015; Xu et al, 2015;...]

What can AI Systems Currently Do?

... driving autonomously, ...



©Google Inc.

What can AI Systems Currently Do?

... run & jump on two legs, ...



©Boston Dynamics


What can AI Systems Currently Do?

... beat a top-gun pilot in a simulated F16 dogfight, ...



What can AI Systems Currently Do?

... synthesize meaningful text automatically, ...



What can AI Systems Currently Do?

... synthesize realistic images or videos automatically, ...



Tilly Norwood, the actress who does not exist (©RollingStone, 2025)

Quiz: What can AI Systems Currently Do?

- | | |
|--|-----|
| • Play a decent game of Jeopardy? | YES |
| • Win against any human at chess? | YES |
| • Win against the best humans at Go? | YES |
| • Play a decent game of tennis? | YES |
| • Grab a particular cup and put it on a shelf? | YES |
| • Unload any dishwasher in any home? | NO |
| • Drive safely along the highway? | YES |
| • Drive safely in Naples' center on rush hour? | NO |
| • Buy groceries on the web? | YES |
| • Buy groceries at next corner shop? | NO |
| • Discover and prove a new mathematical theorem? | NO |
| • Perform a surgical operation? | NO |
| • Translate spoken Chinese into spoken English in real time? | YES |
| • Write an intentionally funny story? | NO |
| • ... | |