

Introduction to Artificial Intelligence

Outline

I. What is AI?

II. A brief history

III. The state of the art

I. What Is Intelligence?

A wish-list of general characteristics of intelligence

- **Perception:** manipulation & interpretation of data provided by sensors

I. What Is Intelligence?

A wish-list of general characteristics of intelligence

- **Perception:** manipulation & interpretation of data provided by sensors
- **Action:** control and use of effectors to accomplish a variety of tasks

I. What Is Intelligence?

A wish-list of general characteristics of intelligence

- **Perception:** manipulation & interpretation of data provided by sensors
- **Action:** control and use of effectors to accomplish a variety of tasks
- **Reasoning:** deductive (logical) inference, inductive inference

I. What Is Intelligence?

A wish-list of general characteristics of intelligence

- **Perception:** manipulation & interpretation of data provided by sensors
- **Action:** control and use of effectors to accomplish a variety of tasks
- **Reasoning:** deductive (logical) inference, inductive inference
- **Learning:** adapting behavior to better cope with changing environments, discovery of patterns, learning to reason, plan, and act.

I. What Is Intelligence?

A wish-list of general characteristics of intelligence

- **Perception:** manipulation & interpretation of data provided by sensors
- **Action:** control and use of effectors to accomplish a variety of tasks
- **Reasoning:** deductive (logical) inference, inductive inference
- **Learning:** adapting behavior to better cope with changing environments, discovery of patterns, learning to reason, plan, and act.
- **Communication:** with other intelligent agents including humans using signals, signs, icons, ...

I. What Is Intelligence?

A wish-list of general characteristics of intelligence

- **Perception:** manipulation & interpretation of data provided by sensors
- **Action:** control and use of effectors to accomplish a variety of tasks
- **Reasoning:** deductive (logical) inference, inductive inference
- **Learning:** adapting behavior to better cope with changing environments, discovery of patterns, learning to reason, plan, and act.
- **Communication:** with other intelligent agents including humans using signals, signs, icons, ...
- **Planning:** formulation of plans – sequences or agenda of actions to accomplish externally or internally determined goals

What is AI?

Thinking Humanly

“The exciting new effort to make computers think ... *machines with minds*, 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)

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)

Thinking Rationally

“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985)

“The study of computations that make it possible to perceive, reason, and act.” (Winston, 1992)

Acting Rationally


“Computational Intelligence is the study of the design of intelligent agents.” (Poole *et al*, 1998)

“AI ... is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)

Acting Humanly: Turing Test

Alan Turing (1950)

Operational definition of intelligence



<https://www.britannica.com/biography/Alan-Turing>




- A human interrogator poses some written questions to another human and a computer (or “robot”).
- The computer passes the test if the interrogator cannot tell whether the written responses come from the human responder or the computer.

<https://wsimag.com/science-and-technology/36961-no-turing-test-for-consciousne1ss>

Acting Humanly: Turing Test

Alan Turing (1950)

Operational definition of intelligence



<https://www.britannica.com/biography/Alan-Turing>



- A human interrogator poses some written questions to another human and a computer (or “robot”).
- The computer passes the test if the interrogator cannot tell whether the written responses come from the human responder or the computer.

<https://wsimag.com/science-and-technology/36961-no-turing-test-for-consciousne1ss>

Annual Loebner prize competition (since 1990): the first prize of \$100,000 to be awarded to the first program that passes the "unrestricted" Turing test.

Six Disciplines of AI

To pass the Turing test, the following capabilities are required:

- ◆ Natural language processing (NLP) to communicate in a human language;
- ◆ Knowledge representation to store what it knows;
- ◆ Automated reasoning to answer questions and draw new conclusions;
- ◆ Machine learning (ML) to adapt to new circumstances and to detect patterns.

Six Disciplines of AI

To pass the Turing test, the following capabilities are required:

- ◆ **Natural language processing (NLP)** to communicate in a human language;
- ◆ **Knowledge representation** to store what it knows;
- ◆ **Automated reasoning** to answer questions and draw new conclusions;
- ◆ **Machine learning (ML)** to adapt to new circumstances and to detect patterns.

Total Turing test requires interaction with the real world.

- ◆ **Computer vision (CV) & speech recognition** to perceive the world;
- ◆ **Robotics** to move around in the world and to manipulate objects.

Six Disciplines of AI

To pass the Turing test, the following capabilities are required:

- ◆ **Natural language processing (NLP)** to communicate in a human language;
- ◆ **Knowledge representation** to store what it knows;
- ◆ **Automated reasoning** to answer questions and draw new conclusions;
- ◆ **Machine learning (ML)** to adapt to new circumstances and to detect patterns.

Total Turing test requires interaction with the real world.

- ◆ **Computer vision (CV) & speech recognition** to perceive the world;
- ◆ **Robotics** to move around in the world and to manipulate objects.

Most vibrant and in-demand areas in AI today: ML, CV, NLP, and Robotics.

Thinking Humanly: Cognitive Modeling

To say a program thinks like a human, we must know humans think first.

Learn about human thought in three ways:

- **introspection** – catching thoughts as they go by;
- **psychological experiments** – observing a person in action;
- **brain imaging** – observing brain activities.

Thinking Humanly: Cognitive Modeling

To say a program thinks like a human, we must know humans think first.

Learn about human thought in three ways:

- **introspection** – catching thoughts as they go by;
- **psychological experiments** – observing a person in action;
- **brain imaging** – observing brain activities.

General problem solver (GPS) by Newell & Simon (1961)

Also compares reasoning steps to those in the human solution.

Thinking Humanly: Cognitive Modeling

To say a program thinks like a human, we must know humans think first.

Learn about human thought in three ways:

- **introspection** – catching thoughts as they go by;
- **psychological experiments** – observing a person in action;
- **brain imaging** – observing brain activities.

General problem solver (GPS) by Newell & Simon (1961)

Also compares reasoning steps to those in the human solution.

Cognitive science constructs theories of the human mind by combining

- ◆ computer models from AI
- ◆ experimental techniques from psychology

Thinking Rationally: “Laws of Thought”

Represent problems using logic and build programs to create intelligent systems.

Circumscription by John McCarthy (1980)

Drawbacks:

- ♣ Not easy to state informal (and often uncertain) knowledge in the formal logical notation.
- ♣ Big difference between solving a problem “in principle” and solving it in practices.

Acting Rationally: Rational Agent


This course is about designing rational agents.

Rational behavior: doing the right thing

↑
expected to maximize goal achievement,
given the available information

- An agent is an entity that perceives and acts.
- A rational agent is one that acts to achieve the best outcome.

Foundations of AI



Philosophy

Mathematics

Economics

Neuroscience

Psychology

Computer
Engineering

Control
Theory

Linguistics

II. Brief History of AI

- 1943 McCulloch & Pitts: model of artificial neurons
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 McCarthy, Minsky, Newell, Simon, Shannon, Nash et al. Dartmouth workshop: birth of "Artificial Intelligence"
- 1952-69 Early enthusiasm, great expectations, optimism fueled by early success on some problems thought to be hard
- 1966-73 Collapse in AI research: Progress was slower than expected. Unrealistic predictions, Herbert Simon (1957) AI discovers computational complexity.

Brief History of AI (cont'd)

- 1969-86 Expert systems
- 1980- AI becomes an industry: expert systems booms, then busts (1988-93): “AI Winter”
- 1986- Neural networks regain popularity
- 1987- Probabilistic reasoning and machine learning

Brief History of AI (cont'd)

- 1995- Emergence of intelligent agents
 - ♣ AI technologies continue to find applications in
 - ◆ information retrieval
 - ◆ data mining and knowledge discovery
 - ◆ customizable software systems
 - ◆ smart devices (e.g., home, automobiles)
 - ◆ agile manufacturing systems
 - ◆ autonomous vehicles
 - ◆ bioinformatics
 - ◆ internet tools: search engines, recommender systems
 - ◆ ...
 - ♣ Steady progress on fundamental AI research problems continues.

Brief History of AI (cont'd)

- 2001- Big data
 - e.g., ImageNet
- 2011- Deep learning
 - ♣ Successful large-scale real-world applications in
 - ♦ image recognition
 - ♦ natural language processing
 - ♦ speech recognition
 - ♦ machine translation
 - ♦ ...
 - ♣ Convolutional neural networks (CNNs)
 - ♣ Chat Generative Pre-trained Transformer (ChatGPT)

Turing Award Winners in AI

1969 Marvin Minsky (MIT)

1971 John McCarthy (Stanford)
National Medal of Science (1990)

1975 Allen Newell* (Carnegie Mellon)
National Medal of Science (1992)

Herbert Simon* (Carnegie Mellon)
Nobel Prize in Economics (1978)
National Medal of Science (1986)

1994 Edward Feigenbaum (Stanford)
Raj Reddy (Carnegie Mellon)

2011 Judea Pearl (UCLA)

2018 Yoshua Bengio (U Montreal)
Geoffrey Hinton (U Toronto)
Yann Lecun (NYU & Facebook)

*omitted from the textbook on p. 17.
Big mistake!

Turing Award Winners in AI

1969 Marvin Minsky (MIT)

1971 John McCarthy (Stanford)
National Medal of Science (1990)

1975 Allen Newell* (Carnegie Mellon)
National Medal of Science (1992)

Herbert Simon* (Carnegie Mellon)
Nobel Prize in Economics (1978)
National Medal of Science (1986)

1994 Edward Feigenbaum (Stanford)
Raj Reddy (Carnegie Mellon)

2011 Judea Pearl (UCLA)

2018 Yoshua Bengio (U Montreal)
Geoffrey Hinton (U Toronto)
Yann Lecun (NYU & Facebook)



Founding
Fathers of AI

*omitted from the textbook on p. 17.
Big mistake!

Turing Award Winners in AI

1969 Marvin Minsky (MIT)

1971 John McCarthy (Stanford)
National Medal of Science (1990)

1975 Allen Newell* (Carnegie Mellon)
National Medal of Science (1992)

Herbert Simon* (Carnegie Mellon)
Nobel Prize in Economics (1978)
National Medal of Science (1986)

1994 Edward Feigenbaum (Stanford)
Raj Reddy (Carnegie Mellon)

2011 Judea Pearl (UCLA)

2018 Yoshua Bengio (U Montreal)
Geoffrey Hinton (U Toronto)
Yann Lecun (NYU & Facebook)



[US News rankings of CS programs in AI \(2024\):](#)

1. Carnegie Mellon
2. MIT
3. Berkeley
4. Stanford
- ...

*omitted from the textbook on p. 17.
Big mistake!


Academic Genealogy

Solomon Lefschetz

Albert Tucker



Herbert Simon



John McCarthy



Marvin Minsky



Allen Newell



Edward Feigenbaum



Raj Reddy

III. The State of the Art

- Game playing
 - ◆ Chinook defeated human checkers champions (1994).
 - ◆ Deep Blue (IBM) defeated the reigning world chess champion Garry Kasparov (1997).
 - ◆ Supercomputer Watson (IBM) beat human champions on “Jeopardy” (2011).
 - ◆ AlphaGo (Google) beat the world’s No.1 ranked player Ke Jie in Go (2017).
 - ◆ Libratus (Carnegie Mellon) defeated the world’s best Texas Hold ‘em poker players head-to-head (2017), six-player (2019).
 - ◆ “AlphaStar” AI (Google) defeated human pros at StarCraft II (2019)

(cont'd)

- Robotic vehicles

- ◆ Stanley (Stanford) won the DARPA Grand Challenge (2005)
- ◆ Boss (Carnegie Mellon) won the DARPA's Urban Challenge (2007)
- ◆ Self-driving cars by Tesla, Google, etc.



Stanley



Boss



Google car



Tesla car crash
Florida, 2016

(cont'd)

- Robotic vehicles

- ◆ Stanley (Stanford) won the DARPA Grand Challenge (2005)
- ◆ Boss (Carnegie Mellon) won the DARPA's Urban Challenge (2007)
- ◆ Self-driving cars by Tesla, Google, etc.



Stanley



Boss



Google car



Tesla car crash
Florida, 2016

- Legged robots (Boston Dynamics)



SpotMini

(cont'd)

- Robotic vehicles

- ◆ Stanley (Stanford) won the DARPA Grand Challenge (2005)
- ◆ Boss (Carnegie Mellon) won the DARPA's Urban Challenge (2007)
- ◆ Self-driving cars by Tesla, Google, etc.




Stanley



Boss




Google car




Tesla car crash
Florida, 2016

- Legged robots (Boston Dynamics)



SpotMini



Atlas



(cont'd)

- Robotic vehicles

- ◆ Stanley (Stanford) won the DARPA Grand Challenge (2005)
- ◆ Boss (Carnegie Mellon) won the DARPA's Urban Challenge (2007)
- ◆ Self-driving cars by Tesla, Google, etc.




Stanley



Boss




Google car




Tesla car crash
Florida, 2016

- Legged robots (Boston Dynamics)



SpotMini



Atlas



Handle

(cont'd)

- Autonomous planning
 - ◆ NASA's Remote Agent program controlled the scheduling of operations for a space craft (2000).
 - ◆ Uber and Google Maps plan optimal routes for hundreds of millions of users.
- Speech recognition
 - ◆ Flight booking via conversation with an automated system
 - ◆ Real-time speech-to-speech translation
- Image understanding
 - ◆ ImageNet object recognition
 - ◆ Image captioning
 - ◆ Face recognition (boarding airplanes, bullet trains, etc.)

} Extensive use of convolutional neural networks (CNNs)
- Natural language processing
 - ◆ ChatGPT by OpenAI allows user to steer and refine a conversation (2022)

Precautions, Downsides & Failures

- Has artificial intelligence become alchemy?

Science, 4 May 2018, Vol 360, Issue 6388, p. 478

- ♠ Lack of rigor
- ♠ No explanation why some algorithms work and why others don't.
- ♠ Inability to replicate results by others

Precautions, Downsides & Failures

- Has artificial intelligence become alchemy?

Science, 4 May 2018, Vol 360, Issue 6388, p. 478

- ♠ Lack of rigor
- ♠ No explanation why some algorithms work and why others don't.
- ♠ Inability to replicate results by others

- 2021's Top Stories About AI > Spoiler: A lot of them talked about what's wrong with machine learning today

IEEE Spectrum, 27 Dec 2021

Precautions, Downsides & Failures

- Has artificial intelligence become alchemy?

Science, 4 May 2018, Vol 360, Issue 6388, p. 478

- ♠ Lack of rigor
- ♠ No explanation why some algorithms work and why others don't.
- ♠ Inability to replicate results by others

- 2021's Top Stories About AI > Spoiler: A lot of them talked about what's wrong with machine learning today

IEEE Spectrum, 27 Dec 2021

“2021 was the year in which the wonders of artificial intelligence stopped being a story.”

Precautions, Downsides & Failures

- Has artificial intelligence become alchemy?

Science, 4 May 2018, Vol 360, Issue 6388, p. 478

- ♠ Lack of rigor
- ♠ No explanation why some algorithms work and why others don't.
- ♠ Inability to replicate results by others

- 2021's Top Stories About AI > Spoiler: A lot of them talked about what's wrong with machine learning today

IEEE Spectrum, 27 Dec 2021

“2021 was the year in which the wonders of artificial intelligence stopped being a story.”

- ♠ Huge computational and energy costs of training deep learning systems
- ♠ AI's turbulent past and uncertain future (its boom-and-bust cycle)
- ♠ Is Deep Learning deep enough?
- ♠ AI failures: brittleness, embedded bias, catastrophic forgetting, explainability, common sense etc.

Precautions, Downsides & Failures

- Has artificial intelligence become alchemy?

Science, 4 May 2018, Vol 360, Issue 6388, p. 478

- ♠ Lack of rigor
- ♠ No explanation why some algorithms work and why others don't.
- ♠ Inability to replicate results by others

- 2021's Top Stories About AI > Spoiler: A lot of them talked about what's wrong with machine learning today

IEEE Spectrum, 27 Dec 2021

“2021 was the year in which the wonders of artificial intelligence stopped being a story.”

- ♠ Huge computational and energy costs of training deep learning systems
- ♠ AI's turbulent past and uncertain future (its boom-and-bust cycle)
- ♠ Is Deep Learning deep enough?
- ♠ AI failures: briittleness, embedded bias, catastrophic forgetting, explainability, common sense etc.

AI is not a science. But it is changing our lives.