

Artificial Intelligence: A Natural Pursuit

Shivaram Kalyanakrishnan

Department of Computer Science and Engineering
Indian Institute of Technology Bombay
`shivaram@cse.iitb.ac.in`

February 2023



https://upload.wikimedia.org/wikipedia/commons/d/d3/Waymo_Chrysler_Pacifica_in_Los_Altos%2C_2017.jpg. CC image courtesy of Dllu on WikiCommons licensed under CC BY-SA 4.0.

2/29

DALL-E

DALL-E

“Teddy bears working on new AI research underwater with 1990s technology.”

DALL-E

“Teddy bears working on new AI research underwater with 1990s technology.”



[https://en.wikipedia.org/wiki/DALL-E#/media/File:
DALL-E_2_artificial_intelligence_digital_image_generated_photo.jpg](https://en.wikipedia.org/wiki/DALL-E#/media/File:DALL-E_2_artificial_intelligence_digital_image_generated_photo.jpg).

3/29

ChatGPT

ChatGPT

“Write a short rhyming poem explaining Einstein’s theory of general relativity in easy but accurate terms”

ChatGPT

“Write a short rhyming poem explaining Einstein’s theory of general relativity in easy but accurate terms”

*If you want to understand Einstein's thought
It's not that hard if you give it a shot
General Relativity is the name of the game
Where space and time cannot remain the same
Mass affects the curvature of space
Which affects the flow of time's race
An object's motion will be affected
By the distortion that is detected
The closer you are to a large mass
The slower time will seem to pass
The farther away you may be
Time will speed up for you to see*

Source: <https://arstechnica.com/information-technology/2022/11/openai-conquers-rhyming-poetry-with-new-gpt-3-update/>.

Retrieved January 4, 2023..

4/29

AlphaZero

AlphaZero



<https://www.flickr.com/photos/prachatai/25708381781>. CC image courtesy of Prachatai on Flickr licensed under CC BY-NC-ND 2.0.

5/29

This Lecture

- What is AI?
- What recently happened to AI?
- Topics in AI
- AI and society

This Lecture

- What is AI?
- What recently happened to AI?
- Topics in AI
- AI and society

AI: Definitions

*"It may even be proposed, as a rule of thumb, that any **activity** computers are able to perform and **people once performed** should be counted as an instance of intelligence."*

Artificial Intelligence and Life in 2030, Peter Stone, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA, September 2016.

AI: Definitions

*"It may even be proposed, as a rule of thumb, that any **activity** computers are able to perform and **people once performed** should be counted as an instance of intelligence."*

Artificial Intelligence and Life in 2030, Peter Stone, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA, September 2016.


*"Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function **appropriately and with foresight in its environment**."*

The Quest for Artificial Intelligence: A History of Ideas and Achievements, Nils J. Nilsson, Cambridge University Press, 2010.

The Urge to Replicate Human Behaviour and Thought



Automaton, Swiss CIMA Museum^[1]




Babbage's Difference Engine (1830s)^[2]

[1] https://upload.wikimedia.org/wikipedia/commons/9/9b/CIMA_mg_8333.jpg. CC image courtesy of Marku1988 on WikiCommons licensed under CC-BY-SA-2.0-FR.

[2] https://upload.wikimedia.org/wikipedia/commons/8/8b/Babbage_Difference_Engine.jpg. CC image courtesy of Geni on WikiCommons licensed under CC-BY-SA-4.0,3.0,2.5,2.0,1.0.


Foundations

Computing



Turing (1912–1954)^[1]

Digital logic



Boole (1815–1864)^[2]

Probability



Bayes (1702–1761)^[3]

Statistics



Mahalanobis (1893–1972)^[4]

[1] https://upload.wikimedia.org/wikipedia/commons/a/a1/Alan_Turing_Aged_16.jpg.

[2] https://upload.wikimedia.org/wikipedia/commons/7/73/PSM_V17_D740_George_Boole.jpg.

[3] https://en.wikipedia.org/wiki/Thomas_Bayes#/media/File:Thomas_Bayes.gif.

[4] <https://upload.wikimedia.org/wikipedia/en/c/ca/PCMahalanobis.png>.

Dartmouth Summer Research Project on Artificial Intelligence (1956)



John McCarthy (1927–2011)^[1]

Workshop attendees: Marvin Minsky, Julian Bigelow, D. M. Mackay, Ray Solomonoff, John Holland, John McCarthy, Claude Shannon, Nathaniel Rochester, Oliver Selfridge, Allen Newell, Herbert Simon.

[1] https://upload.wikimedia.org/wikipedia/commons/4/49/John_McCarthy_Stanford.jpg. CC image courtesy of WikiPedant on WikiCommons licensed under CC-BY-SA-2.0.

1950's–1980's


Theorem proving: **Logic Theorist** (Newell and Simon).

Mobile robotics: **Shakey** (Rosen).

Pattern recognition: **Pandemonium** (Selfridge).

Speech processing: **Spoken language systems** (Reddy).

Expert systems: **Dendral** (Feigenbaum).



Shakey^[1]

[1] https://upload.wikimedia.org/wikipedia/commons/0/0c/SRI_Shakey_with_callouts.jpg. CC image courtesy of Disavian on WikiCommons licensed under CC-BY-SA-3.0.

1950's–1980's


Theorem proving: **Logic Theorist** (Newell and Simon).

Mobile robotics: **Shakey** (Rosen).

Pattern recognition: **Pandemonium** (Selfridge).

Speech processing: **Spoken language systems** (Reddy).

Expert systems: **Dendral** (Feigenbaum).



Shakey^[1]


1980's: AI Winter!

[1] https://upload.wikimedia.org/wikipedia/commons/0/0c/SRI_Shakey_with_callouts.jpg. CC image courtesy of Disavian on WikiCommons licensed under CC-BY-SA-3.0.

This Lecture

- What is AI?
- What recently happened to AI?
- Topics in AI
- AI and society

AI in Life Today



[1]



[2]



[3]



[4]

[1] https://upload.wikimedia.org/wikipedia/commons/c/cf/Waymo_self-driving_car_front_view.gk.jpg. CC image courtesy of Piyush maru on WikiCommons licensed under CC-BY-SA-4.0.


[2] https://upload.wikimedia.org/wikipedia/commons/d/d3/Waymo_Chrysler_Pacifica_in_Los_Altos%2C_2017.jpg. CC image courtesy of Dllu on WikiCommons licensed under CC BY-SA 4.0.

[3] https://upload.wikimedia.org/wikipedia/commons/5/51/IBM_Watson_w_Jeopardy.jpg. CC image courtesy of Rosemaryetoufee on WikiCommons licensed under CC-BY-SA-2.0.

[4] <https://pixabay.com/illustrations/add-to-cart-cart-shopping-buy-add-1747164/>.

13/29

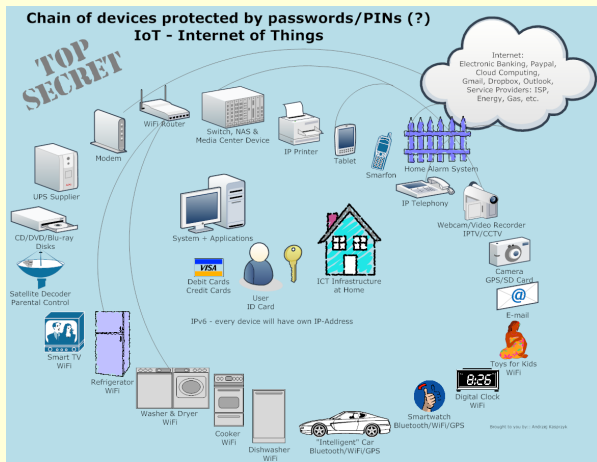
Internet



[1]

[1] <https://www.flickr.com/photos/curiouslee/3485479724>. CC image courtesy of Mike Lee on Flickr licensed under CC BY 2.0. 14/29

Cheaper Hardware and Sensors

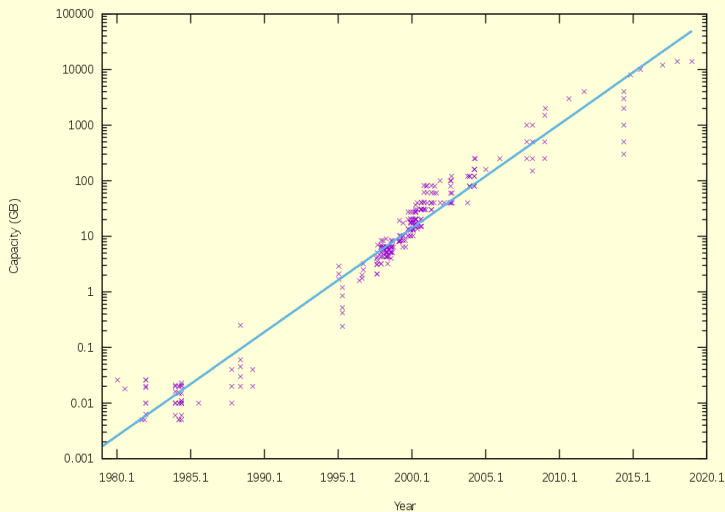


[1]

[1] [https://commons.wikimedia.org/wiki/File:Chain_of_home_devices_\(including_IoT\)_with_passwords_or_pin.png](https://commons.wikimedia.org/wiki/File:Chain_of_home_devices_(including_IoT)_with_passwords_or_pin.png).
CC image courtesy of AKasprzyk on WikiCommons licensed under CC-BY-3.0.

15/29


Growth of Data




[1]

[1] https://upload.wikimedia.org/wikipedia/commons/thumb/9/90/Hard_drive_capacity_over_time.svg/1000px-Hard_drive_capacity_over_time.svg.png. CC image courtesy of JoKalliauer on WikiCommons licensed under CC-BY-SA-3.0,2.5,2.0,1.0.


Machine Learning




Machine Learning



Machine Learning



Machine Learning




x_1	x_2	Label
12	1	-
-4	20	+
-15	-2	+
-4	-4	-
15	-6	-

Learn a model

16	-7	?
----	----	---

Machine Learning




x_1	x_2	Label
12	1	-
-4	20	+
-15	-2	+
-4	-4	-
15	-6	-

Learn a model

16	-7	?
----	----	---

Supervised learning. Points can be high-dimensional (even millions).

Machine Learning



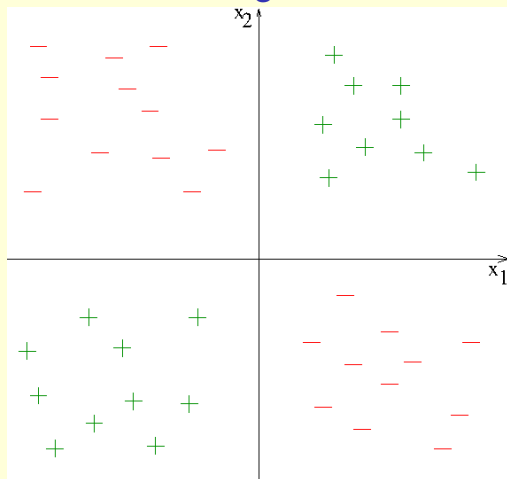
x_1	x_2	Label
12	1	-
-4	20	+
-15	-2	+
-4	-4	-
15	-6	-

Learn a model

16	-7	?
----	----	---

Supervised learning. Points can be high-dimensional (even millions).
Applications: Face recognition, Credit fraud detection, Sentiment analysis, . . .

Machine Learning



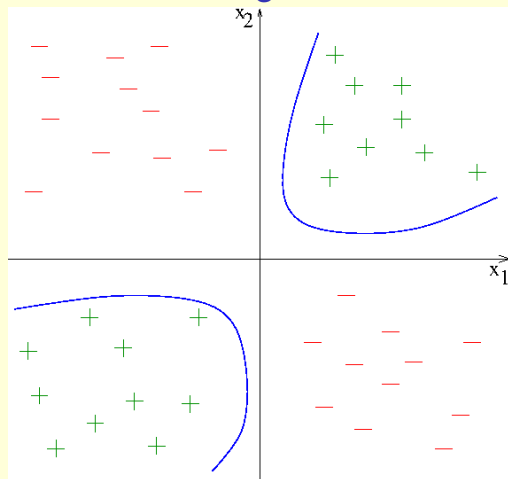
x_1	x_2	Label
12	1	-
-4	20	+
-15	-2	+
-4	-4	-
15	-6	-

Learn a model

16	-7	?
----	----	---

Supervised learning. Points can be high-dimensional (even millions).
Applications: Face recognition, Credit fraud detection, Sentiment analysis, . . .

Machine Learning



x_1	x_2	Label
12	1	-
-4	20	+
-15	-2	+
-4	-4	-
15	-6	-

Learn a model

16	-7	?
----	----	---

Supervised learning. Points can be high-dimensional (even millions).

Applications: Face recognition, Credit fraud detection, Sentiment analysis, . . .

Deep learning can find highly non-linear patterns in visual, audio, text data.

This Lecture

- What is AI?
- What recently happened to AI?
- Topics in AI
- AI and society

Machine Learning: Supervised Learning

Given labeled data, produce model to predict labels for unseen data.




Figure 2. CTG Signal over 10 Minutes, Including Four Contraction-Deceleration Pairs.

Top: UP signal with contraction onsets (C) indicated. Bottom: FHR signal with deceleration onsets (D) indicated.

[1]

1. A Machine Learning Approach to the Detection of Fetal Hypoxia during Labor and Delivery, Philip A. Warrick, Emily F. Hamilton, Robert E. Kearney, Doina Precup, AI Magazine, 33(2):79–90, AAAI Press, 2012.

18/29

Machine Learning: Unsupervised Learning

Given unlabeled data, produce model to assign to clusters.



[1]

[1] ESA/Hubble: https://upload.wikimedia.org/wikipedia/commons/6/61/Globular_Star_Cluster_NGC_1466.png. CC image courtesy of Killarnee on WikiCommons licensed under CC-BY-4.0.

Machine Learning: Reinforcement Learning

How does a child learn to walk?



[1]

[1] <https://www.publicdomainpictures.net/pictures/120000/velka/crawling-baby.jpg>.

Machine Learning: Reinforcement Learning

How does a child learn to walk?

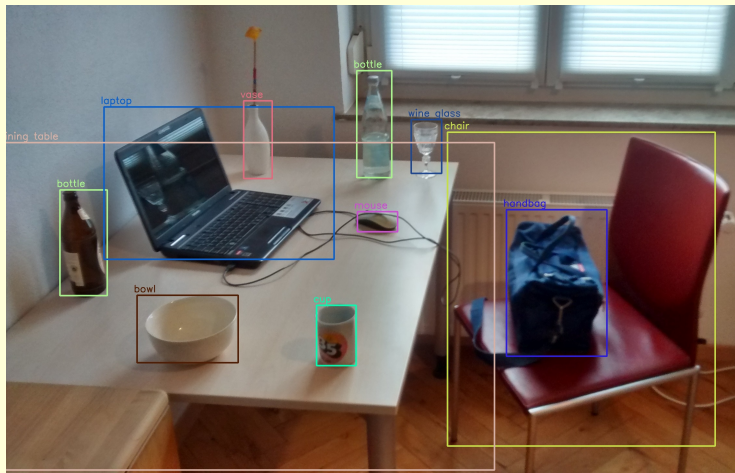


[1]

Learning by **trial and error** to take successful **sequences** of actions.

[1] <https://www.publicdomainpictures.net/pictures/120000/velka/crawling-baby.jpg>.


Computer Vision



[1]

1. <https://upload.wikimedia.org/wikipedia/commons/3/38/Detected-with-YOLO--Schreibtisch-mit-Objekten.jpg>. CC image courtesy of MTheiler on WikiCommons licensed under CC-BY-SA-4.0.

Speech and Natural Language Processing



[1]

Topics: Text summarisation, Sentiment analysis, Machine translation, etc.

[1] https://upload.wikimedia.org/wikipedia/commons/6/6c/Indian_language_word_cloud_depicting_language_neutrality.png.

CC image courtesy of Psubhashish on WikiCommons licensed under CC-BY-SA-4.0.

23/29

Heuristic Search

How to go from IIT Bombay to Carter Road?



[1]

[1] <https://www.flickr.com/photos/nat507/16088993607>. CC image courtesy of Nathan Hughes Hamilton on Flickr licensed under CC BY 2.0.


Heuristic Search

How to go from IIT Bombay to Carter Road?

How to solve Sudoku?



[1]



[2]

[1] <https://www.flickr.com/photos/nat507/16088993607>. CC image courtesy of Nathan Hughes Hamilton on Flickr licensed under CC BY 2.0.

[2] https://upload.wikimedia.org/wikipedia/commons/e/eb/Sudoku_Puzzle_%28a_symmetrical_puzzle_with_17_clues%29.png. CC image courtesy of LithiumFlash on WikiCommons licensed under CC-BY-SA-4.0.

Multiagent Systems: Game Theory



[1]

Security and game theory: algorithms, deployed systems, lessons learned,
Milind Tambe, Cambridge University Press, 2012.

[1] https://upload.wikimedia.org/wikipedia/commons/1/1b/2016_04_19_Airport_Security-1_%2826140081053%29.jpg. 25/29

Knowledge Representation and Reasoning



[1]

[1] https://cdn.pixabay.com/photo/2013/07/18/15/06/books-164530_960_720.jpg.

Planning and Scheduling



[1]

[1] https://upload.wikimedia.org/wikipedia/commons/0/0e/Three_different_generations_of_Indian_railway_locomotives.jpg.
CC image courtesy of AuthIM3847 or on WikiCommons licensed under CC-BY-SA-4.0.

This Lecture

- What is AI?
- What recently happened to AI?
- Topics in AI
- AI and society

AI and Society

- Opportunities

- ▶ Scale, efficiency, reliability: advantages over humans.
- ▶ New applications, enhancements to quality of life.

- Challenges

- ▶ Loss of jobs.
- ▶ Data privacy/security.
- ▶ Inequalities in access to AI technology.
- ▶ Biases in data \implies biases in AI's decision making.
- ▶ Trust, liability.

AI and Society

- Opportunities

- ▶ Scale, efficiency, reliability: advantages over humans.
- ▶ New applications, enhancements to quality of life.

- Challenges

- ▶ Loss of jobs.
- ▶ Data privacy/security.
- ▶ Inequalities in access to AI technology.
- ▶ Biases in data \implies biases in AI's decision making.
- ▶ Trust, liability.

Thank you!