

Artificial Intelligence: An Introduction

Wei Wang(王伟)

Engineering Research Centre of Applied Technology on Machine Translation and Artificial Intelligence , Macao
Polytechnic University

weiwang@mpu.edu.mo; 匯智樓 (WUI CHI)-5/F, N56

Jan. 12, 2026

Outline

- **About the Course**

About the Course

- When: 2025/2026, Mon, 14:30AM-17:30PM
(**Make-up**: Apr. 11; Apr. 27)
- Where: E905
- Course Code: COMP421
- Professor: Wei Wang
- Office Hours: Tue., 9:30- 12:00 AM
- TA: TBD

Final Grade:

The final grade will be determined as follows:

- 25% Assignments
- 25% Midterm Test (TBD)
- 50% Examination (**Apr. 30 19:00 - 22:00**)

Assignments:

The course will focus on **understanding** concepts and algorithms rather than using AI for solving problems than programming, although there will be two projects involving some **low-code programming**.

Assignment 1: from week 2. (**AI as Research**)

Assignment 2: from week 7/8. (**AI as Tool**)

Students with an overall score of **less than 35 in the coursework** must take the re-sit examination even if the overall score for the module is 50 or above.

Students with a score of **less than 35 in the final examination** must take the re-sit examination even if the overall score for the module is 50 or above.

Students with an overall final grade of **less than 35** are NOT allowed to take the re-sit examination.

Prompt attendance at class is required from the University

Attendance will be taken during each class session, either through **roll call** or via **assignments submitted during class**.

About the Course

This learning module will systematically introduce the theory and practice of artificial intelligence (AI). You will master the **fundamental principles** driving these applications and gain hands-on experience implementing certain systems. The course covers topics including search heuristics, Markov decision processes, Bayesian inference, and machine learning, comprehensively addressing core domains of modern AI—from deterministic reasoning to uncertain reasoning, and from rule-based systems to learning systems. The core objective of this course is to equip you with the tools to tackle novel AI challenges you may encounter in real-world scenarios.

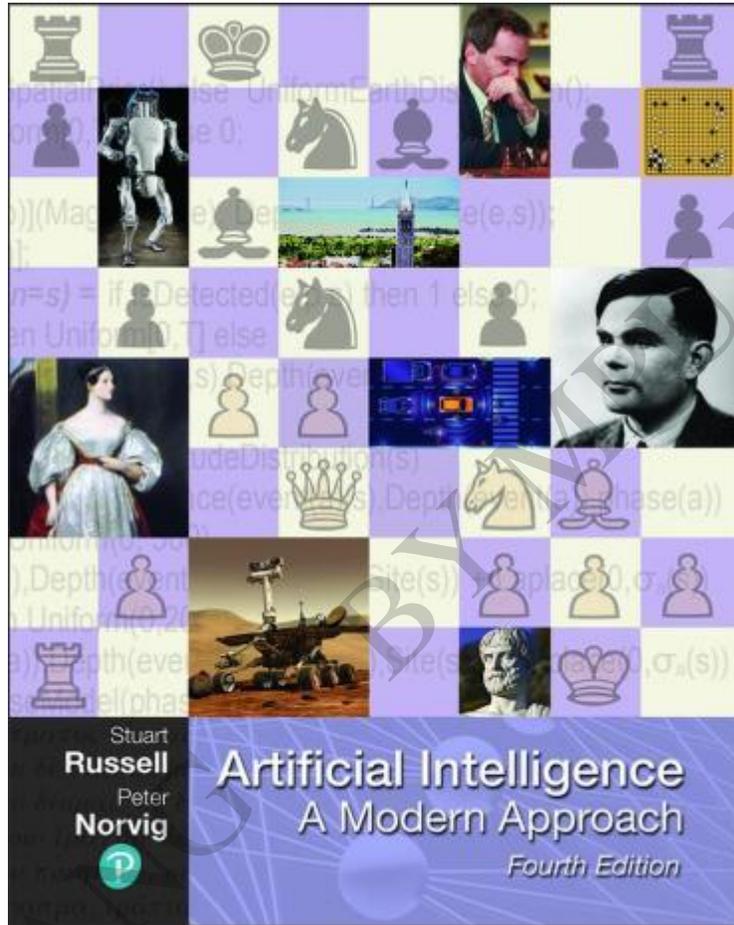
Know AI Hot Topics

Agentic AI & Multi-Agent Systems
World Models & Next-State Prediction
Embodied Intelligence
AI for Science (AI4S) & AI Scientists
AI Safety & Trustworthy AI

Know Popular AI Tools

ChatGPT (GPT-5)
DeepSeek
Ddoubao
Nano-Banana
Google Gemini
Tongyi Qwen
Tongyi Wanxiang

About the Course



**Python code for the book Artificial Intelligence:
A Modern Approach:**

<https://github.com/aimacode/aima-python>

Harvard University:

CS50,s Introduction to Artificial Intelligence with Python

<https://cs50.harvard.edu/ai/2024/>

Stanford University:

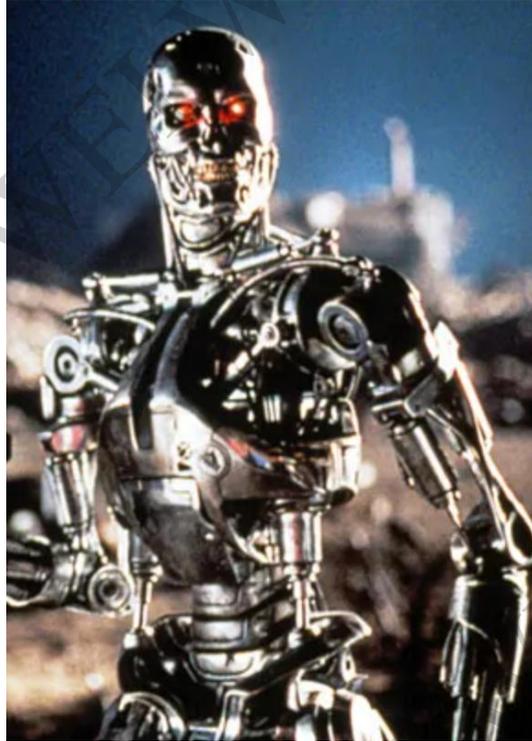
CS221: Artificial Intelligence: Principles and Techniques

<https://stanford-cs221.github.io/spring2024/>

Outline

- **What is AI?**

AI in the Movie





WORKFORCE & HR AI

January 9, 2026

From cloud to factory – humanoid robots coming to workplaces



HEALTHCARE & WELLNESS AI

January 8, 2026

“Dr AI, am I healthy?” 59% of Brits rely on AI for self-diagnosis



AI MARKET TRENDS

January 8, 2026

2026 to be the year of the agentic AI intern



MANUFACTURING & ENGINEERING

January 8, 2026

Bosch's €2.9 billion AI investment and shifting manufacturing priorities

L'Oréal brings AI into everyday digital advertising production

AI IN ACTION

January 5, 2026

Disney is embedding generative AI into its operating model

ENTERTAINMENT & MEDIA

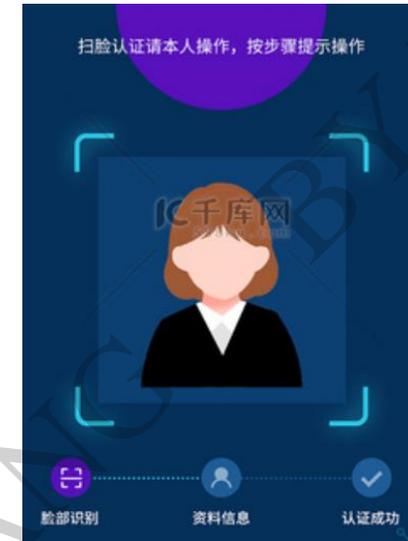
December 24, 2025

Zara's use of AI shows how retail workflows are quietly changing

AI in Daily Life



AI in Daily Life



What is Intelligence?

The word “intelligence” is surprisingly recent. Ancients used it to mean “the universal mind.”

Early moderns (e.g., Bacon, Hobbes; 1500s) ridiculed it, and stopped using it. It was then repurposed to its current meaning by psychologists and eugenicists in the early 20th century.

What is Intelligence?

Spearman's correlation matrix for six measures of school performance. All the correlations are positive, the *positive manifold* phenomenon. The bottom row shows the *g* loadings of each performance measure.^[7]

	Classics	French	English	Math	Pitch	Music
Classics	–					
French	.83	–				
English	.78	.67	–			
Math	.70	.67	.64	–		
Pitch discrimination	.66	.65	.54	.45	–	
Music	.63	.57	.51	.51	.40	–
<i>g</i>	.958	.882	.803	.750	.673	.646

[https://en.wikipedia.org/wiki/G_factor_\(psychometrics\)](https://en.wikipedia.org/wiki/G_factor_(psychometrics))

Charles Spearman popularized the modern definition in his paper “General intelligence objectively determined and measured,” *American Journal of Psychology* 15(2):201-292.

He showed that test scores are correlated across many subjects and proposed “general intelligence” as the faculty that unifies them.

What is AI?

The term was invented in (**John McCarthy**, Marvin Minsky, Nathaniel Rochester, and **Claude Shannon**, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” **August 1956**):

“We propose that a 2-month, 10-man study of artificial intelligence be carried out during the summer ... An attempt will be made to find how to make machines

1. use language,
2. form abstractions and concepts,
3. solve kinds of problems now reserved for humans
4. improve themselves.”



“An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

We think that a significant advance can be made if we work on it together for a summer.”

What is AI?

Russell & Norvig's "four approaches to AI" : Intelligence means...

Artificial Intelligence, A Modern Approach: Third Edition by Russell & Norvig

1. Thinking like
a Human

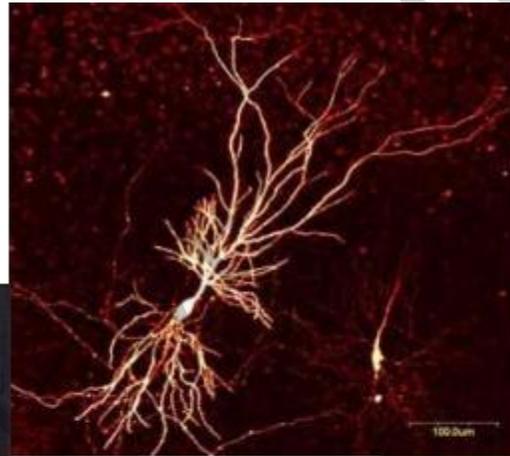
2. Acting like a
Human

3. Thinking
Rationally

4. Acting
Rationally

Thinking like a Human

An English novelist best known for *Frankenstein*, a foundational work of science fiction that explores the ethical limits of scientific ambition



Should we create human-like intelligence?

Mary Shelley, author of *Frankenstein: The Modern Prometheus*; Neuron, showing branching of the dendrites; EEG cap; Cortical connectivity map, computed using diffusion tensor MRI

Can we simulate a human brain?

How many binary computations per second can the brain perform?

- Spatial scale: there are 100 trillion neurons (10^{14}).
- Numerical precision: each neuron either generates an action potential or doesn't (binary!).
- Temporal scale: Other neurons are sensitive to timing with a resolution of perhaps roughly 1 millisecond (1000 bits/second).

Answer: if each neuron performs 1000 binary computations/second, then the brain performs up to $(100 \text{ trillion}) \times (1000) = 10^{17}$ binary computations/second (100 Peta-ops: about 100,000 GPUs)



Then why can't we simulate a human brain?

How many brain computations can we IMAGE?

- Temporal scale: no problem, EEG(electroencephalography) gets ~ 5000 samples/second
- Spatial scale is the problem:
 - EEG: 100 pixels/brain
 - fMRI and ECOG: 1mm scale ($\sim 10^5$ voxels/brain)
 - Versus 10^{14} neurons/brain



Then why can't we simulate a human brain?

- The short answer: we can't find out what computations a living human brain is performing, because any current imaging modality that precise would kill it.
- ...and we are currently about 9 orders of magnitude (10^9) away from the necessary level of precision (volume).
- MRI improved by roughly 2 orders of magnitude per decade from 1970 to 2000, then slowed significantly, has improved perhaps 1 o.o.m. per two decades since then. So perhaps this approach will be possible in 180 years.

Acting like a Human

Russell & Norvig's "four approaches to AI" : Intelligence means...

Artificial Intelligence, A Modern Approach: Third Edition by Russell & Norvig

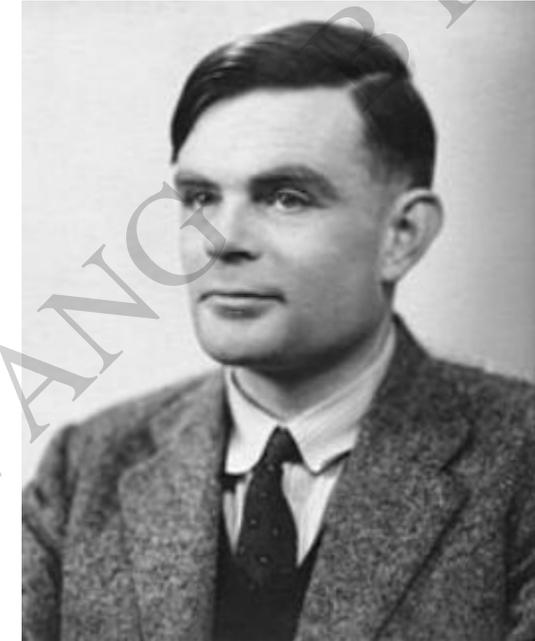
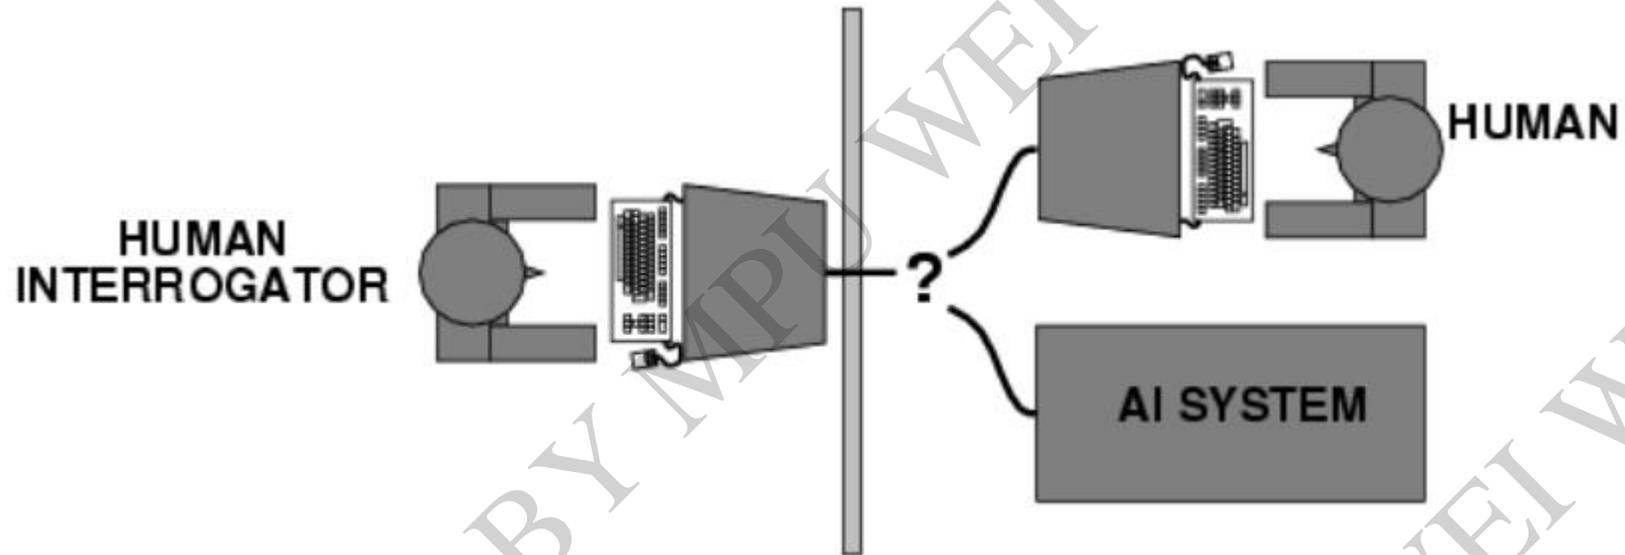
1. Thinking like a Human

2. Acting like a Human

3. Thinking Rationally

4. Acting Rationally

Acting like a Human



Schematic of the Turing test; Alan Turing

The Turing Test

- Alan Turing, “Intelligent Machinery,” 1947:

“Now get three men as subjects for the experiment. A, B and C. A and C are to be rather poor chess players, B is the operator who works the paper machine... a game is played between C and either A or the paper machine.

C may find it quite difficult to tell which he is playing...

These questions replace our original, ‘Can machines think?’”

Practical Problems with the Turing Test

- Can't be automated (you need human judges).
- Human judges can be fooled by misdirection, e.g., by a chatbot that pretends to be a paranoid schizophrenic (<https://en.wikipedia.org/wiki/PARRY>) or a 13-year-old Ukrainian boy (https://en.wikipedia.org/wiki/Eugene_Goostman)

Winograd Schema

- Winograd schema (H. Levesque, [On our best behaviour](#), IJCAI 2013) attempts to solve the practical problems with the Turing test
- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”
- Always arranged in pairs:

The trophy would not fit in the brown suitcase because it was so small, big. What was so small/big?

- *The trophy*
- *The brown suitcase*

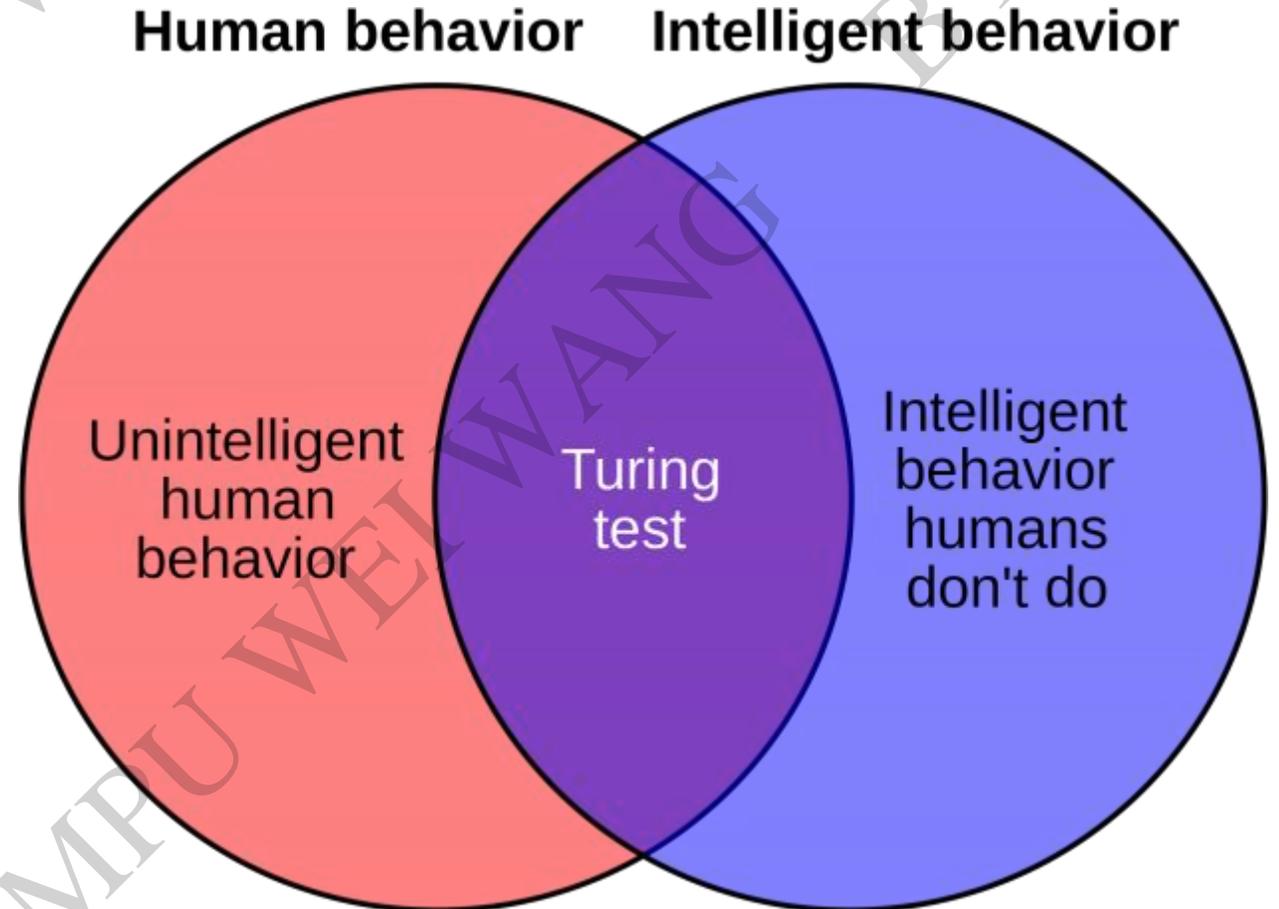
A theoretical problem with the Turing test

Why is human behavior the standard?

Red Zone (Unintelligent human behavior): Represents non-intelligent human actions (such as logical errors, incoherent expressions, etc.);

Blue Zone (Intelligent behavior humans don't do): Represents intelligent actions humans cannot perform (such as flawless logical reasoning, ultra-large-scale data computation, etc.);

Purple Intersection Zone (Turing test): Its core lies in determining whether machine behavior simultaneously satisfies the boundary between “non-human unintelligent actions” and “human intelligent actions”, that is, judging whether a machine possesses “intelligence” through behavioral indistinguishability.



<https://commons.wikimedia.org/w/index.php?curid=3457053>

Acting like a Human

Tech

Turing Test breakthrough as super-computer becomes first to convince us it's human

Eugene Goostman, a computer programme pretending to be a young Ukrainian boy, successfully duped enough humans to pass the iconic test

Andrew Griffin

Monday 09 June 2014 12:55 BST



研究表明GPT-4通过图灵测试



第一财经

2024-06-17 11:37 | 第一财经官方账号

关注

来自加利福尼亚大学圣迭戈分校（University of California in San Diego）认知科学家本杰明·伯根（Benjamin Bergen）和卡梅隆·琼斯（Cameron Jones）的最新研究结果表明，越来越多的人难以在图灵测试中区分GPT-4和人类。

在上述科学家所做的一项实验中，500名人类与四种AI语言模型进行了5分钟的对话，其中GPT-4在54%的时间里被误认为是人类，这个比例超过了此前版本GPT-3.5的相应比例（50%）。

这一结果表明，GPT-4已通过图灵测试。图灵测试是用来判断机器是否具有模拟人类思维能力的测试。

Russell & Norvig's "four approaches to AI" : Intelligence means...

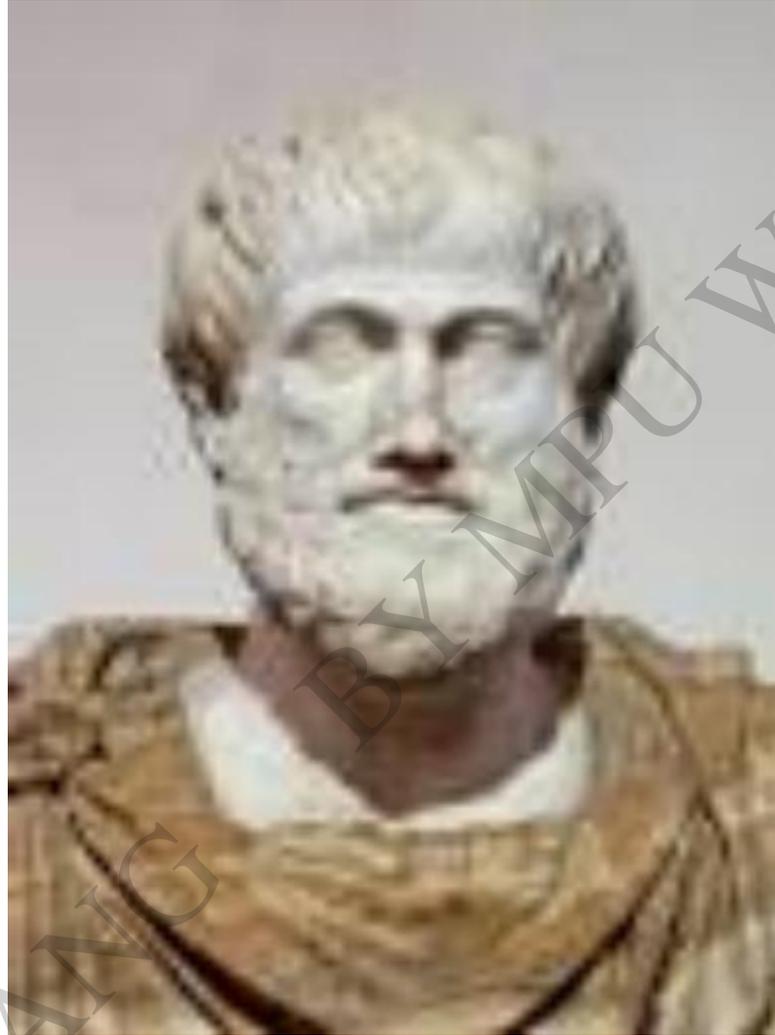
Artificial Intelligence, A Modern Approach: Third Edition by Russell & Norvig

1. Thinking like a
Human

2. Acting like a
Human

3. Thinking
Rationally

4. Acting
Rationally



Aristotle (384–322 BC)

The Nature of Intelligence

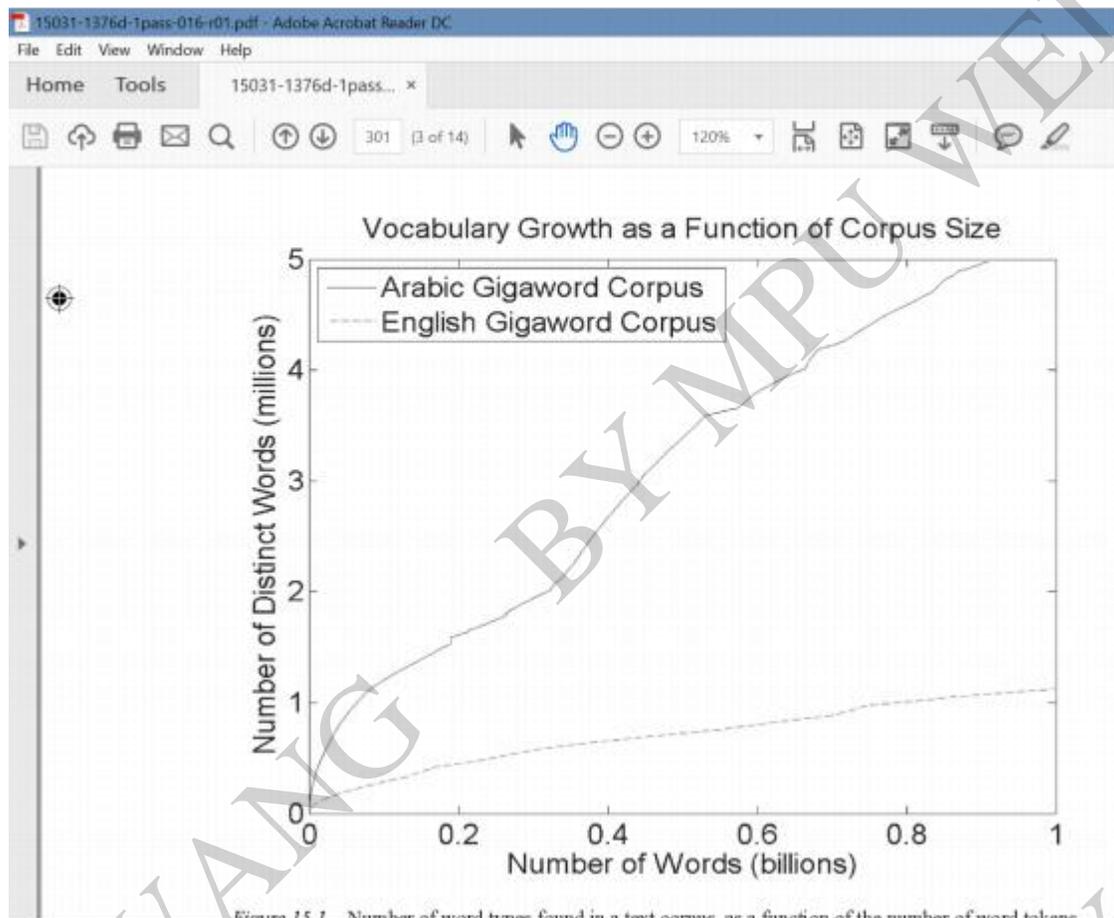
- Logic as structured reasoning
 - Intelligence as goal-directed activity
- Why intelligence can be modeled

- Idealized or “right” way of thinking
- **Logic:** patterns of argument that always yield correct conclusions when supplied with correct premises
 - *“Socrates is a man; all men are mortal; therefore Socrates is mortal.”*
- **Logicist approach to AI:** describe problem in formal logical notation and apply general deduction procedures to solve it

Successes of Logician Approach: **Expert Systems**

- Expert system = (knowledge base) + (logical rules)
 - Knowledge base = easy to collect from human judges and/or encyclopedia
 - Logical rules = easy to deduce from examples, and easy to verify by asking human judges
 - Combination of the two: able to analyze never-before-seen examples of complicated problems, and generate the correct answer
- Example: speed control system of the https://en.wikipedia.org/wiki/Sendai_Subway_Namboku_Line. “This system (developed by Hitachi) accounts for the relative smoothness of the starts and stops when compared to other trains, and is 10% more energy efficient than human-controlled acceleration.”

Failures of Logicist Approach: Robust AI



- Humans commonly believe that there are a finite number of facts that must be entered into a knowledge base. Evidence suggests that this is **incorrect**.
- Example (Hasegawa-Johnson, Elmahdy & Mustafawi, "Arabic Speech and Language Technology," 2017): the number of distinct words in any corpus of text is linearly proportional to the number of words. In English, a never-before-seen word occurs \sim once/1000 words; in Arabic, \sim once/180 words.

Russell & Norvig's "four approaches to AI" : Intelligence means...

Artificial Intelligence, A Modern Approach: Third Edition by Russell & Norvig

1. Thinking like a
Human

2. Acting like a
Human

3. Thinking
Rationally

4. Acting
Rationally

- **Rational agent = acts to optimally achieve its goals**
 - Goals are application-dependent and are expressed in terms of the **utility of outcomes**
 - Being rational means **maximizing your (expected) utility**
- This definition of rationality only concerns the decisions/actions that are made, not the cognitive process behind them
- An unexpected step: rational agent theory was originally developed in the field of economics
 - Norvig and Russell: “most people think Economists study money. Economists think that what they study is the behavior of rational actors seeking to maximize their own happiness.”

Russell & Norvig's "four approaches to AI"

Intelligence is...

1. Thinking like a human

- Sometimes called "grounded AI" – create an AI with neurons like ours

2. Acting like a human

- Turing's definition of AI; ignores the underlying thought process
- Might include acting irrationally

3. Thinking rationally

- Logician AI: must use a rational/logical thought process

4. Acting rationally

- Utilitarianism: act in order to maximize your own benefit, regardless of the thought process you use

What is AI?



Mary Shelley (1797–1851)

The Question of Creation

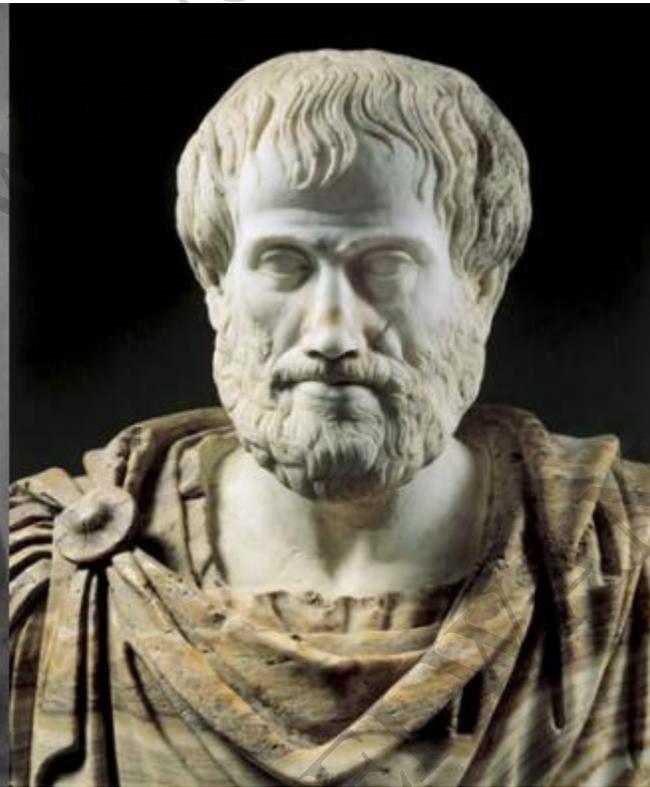
- Can humans create intelligence?
 - What responsibilities follow creation?
- Why AI ethics exists



Alan Turing (1912–1954)

The Machine Hypothesis

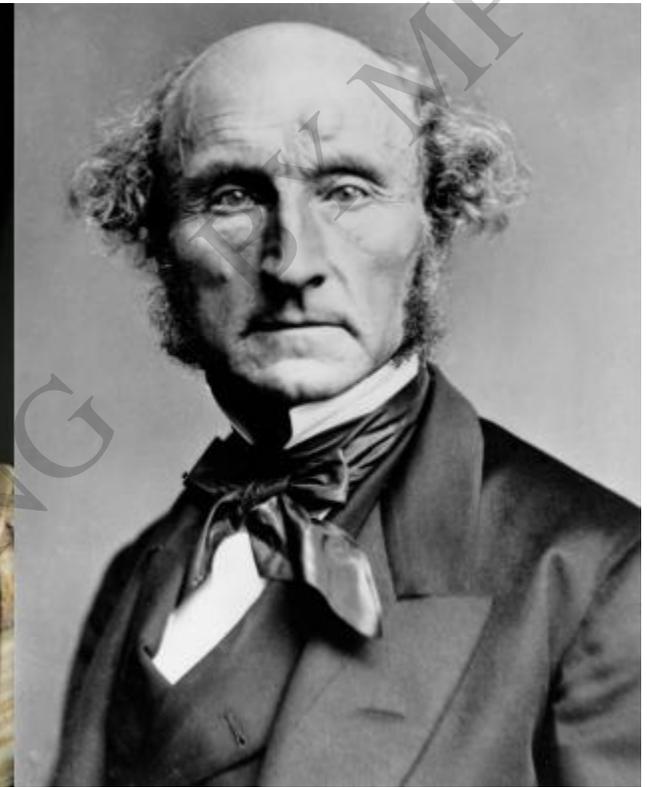
- Intelligence as computable behavior
 - Turing Machine & Turing Test
- Why AI can be built



Aristotle (384–322 BC)

The Nature of Intelligence

- Logic as structured reasoning
 - Intelligence as goal-directed activity
- Why intelligence can be modeled



John Stuart Mill (1806–1873)

Ethics & Governance

- Freedom and moral limits
 - Utility vs. individual rights
- Why AI must be governed

AI is Developing Fast

Short video on Shakey

<https://www.youtube.com/watch?v=7bsEN8mwUB8>

Short video on Unitree Humanoid Robot Training

<https://www.youtube.com/watch?v=JZlfrHRc4g>

Short video on Unitree Humanoid Robot Dancing

<https://www.youtube.com/watch?v=QOQP1dyASls>

Breaking Records in User Adoption

Technology	Time to 100M Users
Telephone	75 Years
Mobile Phone	16 Years
World Wide Web	7 Years
Facebook	4.5 Years
TikTok	9 Months
ChatGPT	2 Months

This is the classic "Hockey Stick" graph.

In the past, technology took decades to permeate society.

AI (specifically Generative AI) has seen the fastest consumer adoption in human history.

This means the feedback loop (users -> data -> better models) is tighter than ever.

AI is Developing Fast

It's Not Just Chips, It's the Demand

Moore's Law: Computing power doubles roughly every 2 years.

• **Deep Learning Era:** Since 2012, the amount of compute used in the largest AI training runs has doubled every 3.4 months.

"We are increasing the compute used to train these models by **a factor of 10** every year."

Key Driver:

• **Parallelization:** Unlike traditional software that runs sequentially, Neural Networks can be split across thousands of GPUs running simultaneously.

• Look at the disparity. Hardware gets better every **2** years, but our usage of that hardware for AI is outpacing Moore's Law by a massive margin.

• This is why **NVIDIA and GPU** clusters have become the most valuable resource in tech.



The Explosion of Model Size

Parameter Counts (The "Brain Cells" of AI)

Year	Model	Parameters (Approx)	Capability
2018	GPT-1	117 Million	Basic sentences
2019	GPT-2	1.5 Billion	Coherent paragraphs
2020	GPT-3	175 Billion	Coding, Translation, Creative Writing
2023	GPT-4	Trillions (MoE)*	Bar exam pass (Top 10%), Medical diagnosis

A "parameter" is roughly equivalent to a synapse in the brain—a connection where information is stored.

In just 5 years, we went from 100 million to trillions.

Notice how "Emergent Capabilities" appear: The model wasn't explicitly taught to code, but at a certain scale, it learned the logic of code from the data.

Note: Recent models focus on "efficiency" (Small Language Models) rather than just size, but the capability jump remains exponential.

AI is Developing Fast

Visual Evolution (Computer Vision)

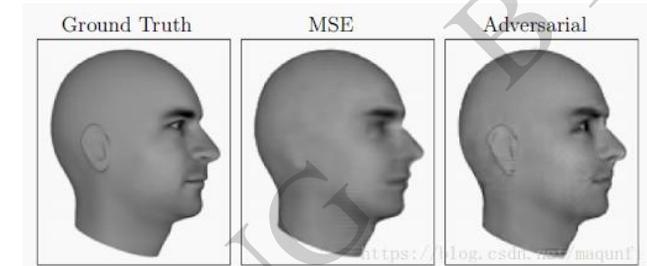
From "Pixelated Blobs" to "Hyper-Realism"

- 2014 (GANs introduced): Generated faces looked like black-and-white blurry smears.
- 2018: Faces looked realistic but had artifacts (weird ears, glasses).
- 2022: Midjourney V4 – Artistic, high definition.
- 2024: Sora/Video Generation – High-definition video with consistent physics.

<https://www.youtube.com/watch?v=1PaoWKvcJP0>

Concept: **The "Half-Life" of State-of-the-Art**

- In AI, a "cutting edge" paper is often obsolete within **6** months.



AI is Developing Fast

Why is it accelerating?

The Flywheel Effect

1. **Better Hardware creates** →

2. **Better AI Tools** (AI helps design better chips and write better code) which creates →

3. **More Efficiency** (AI helps researchers read papers and summarize data) which leads to →

4. **Faster Breakthroughs** (Cycle repeats).

Recursion: We are using AI to build the next generation of AI. (e.g., Using Github Copilot to write PyTorch code).

This is unique to this field. A steam engine couldn't design a better steam engine. But an AI can optimize the code for the next AI. This suggests we haven't hit the ceiling yet.

The Impact of Speed

Disruption Timeline

Traditional Software Cycle: Updates every 1-2 years.

- **AI Cycle:** Updates weekly/monthly.

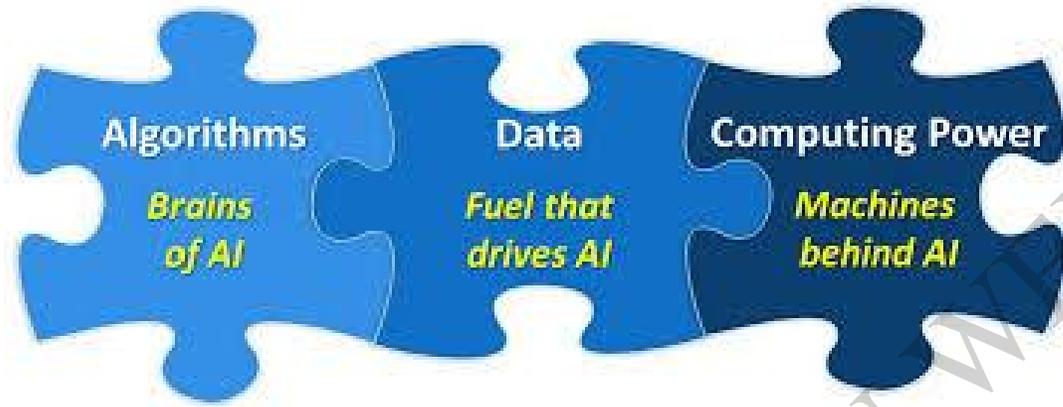
Consequences:

- **Regulatory Lag:** Laws cannot keep up with the tech (Copyright, Privacy).
- **Education Gap:** Universities struggle to update curriculum fast enough.
- **Workforce Adaptation:** Employees must "upskill" constantly.

The speed is exciting, but also destabilizing.

Governments are debating "How do we regulate something that changes every week?"

AI is Developing Fast



**THE ENERGY
CONSTRAINT**



Discussion

Is there a limit?

The Data Wall: Are we running out of high-quality internet data to train on?

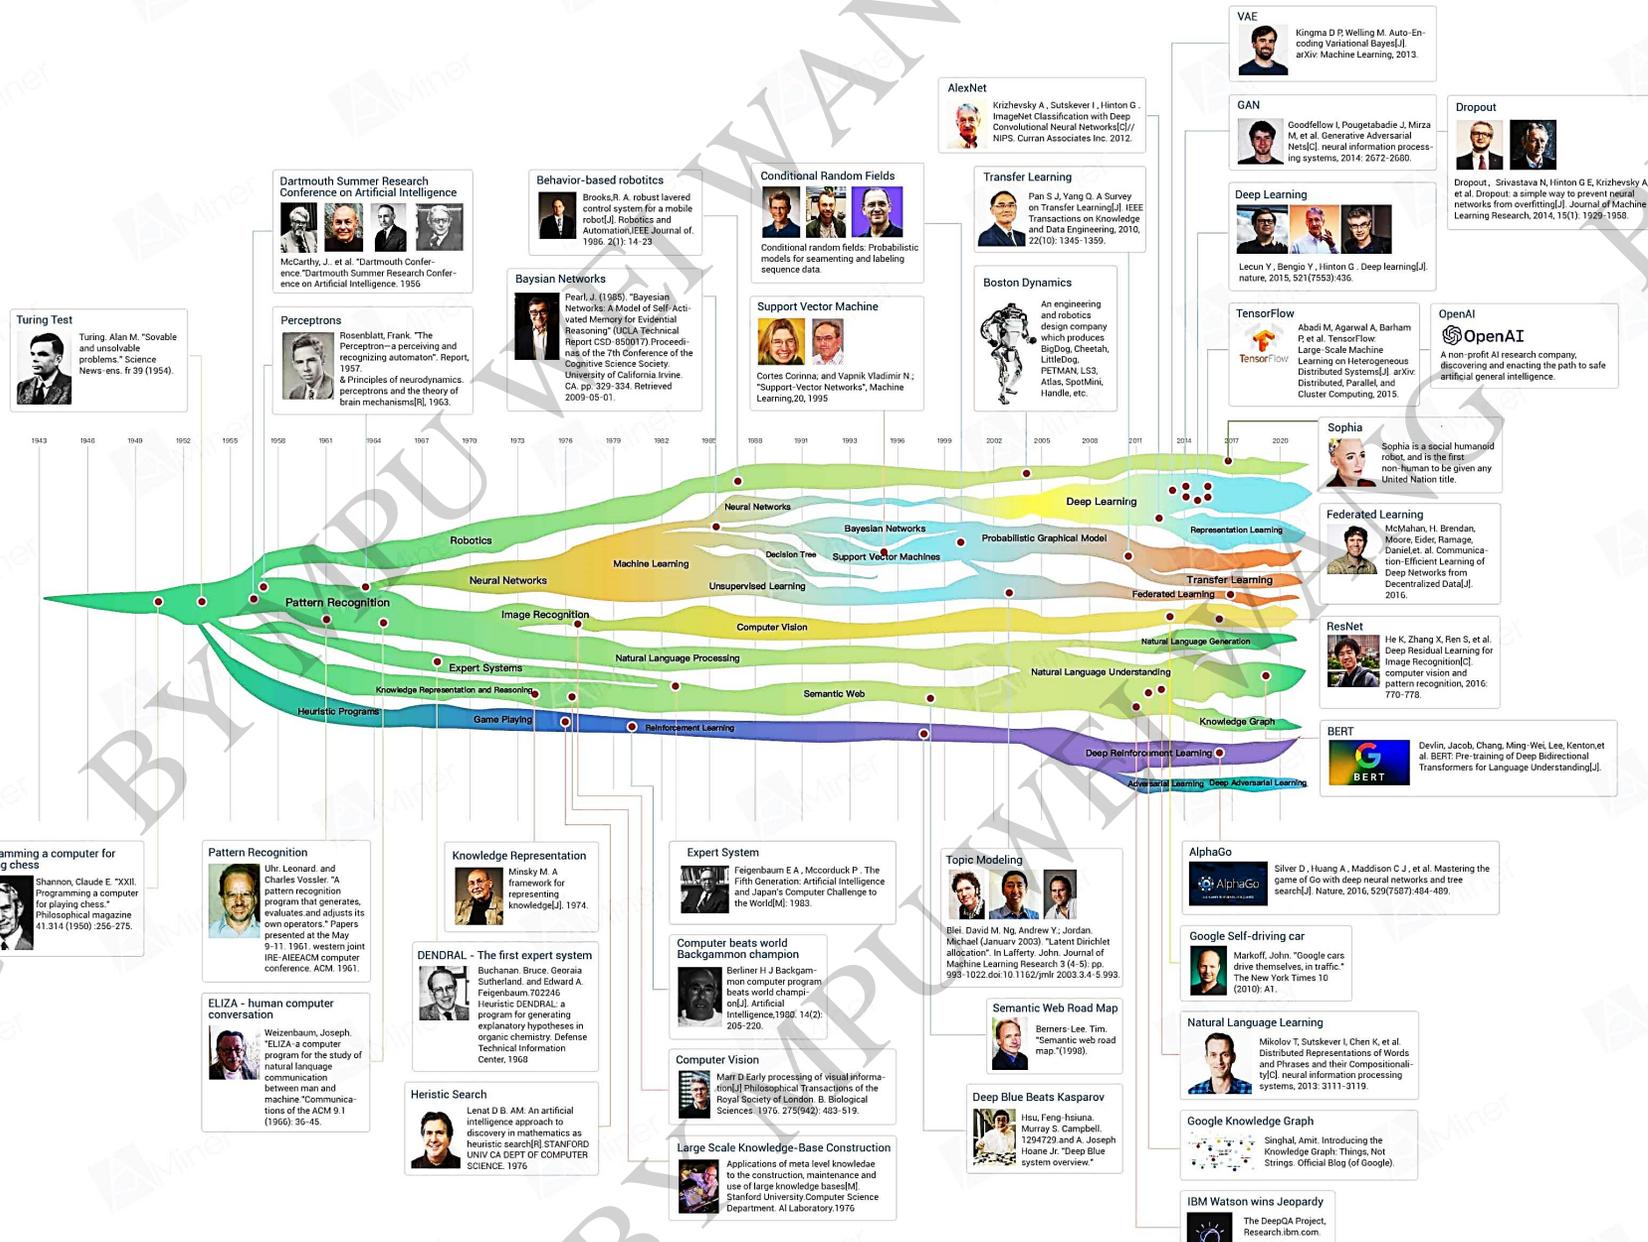
Energy Constraints: AI data centers consume massive amounts of electricity.

Discussion Question: Do you think the progress will plateau (Smoothly) or keep accelerating (Singularity)?

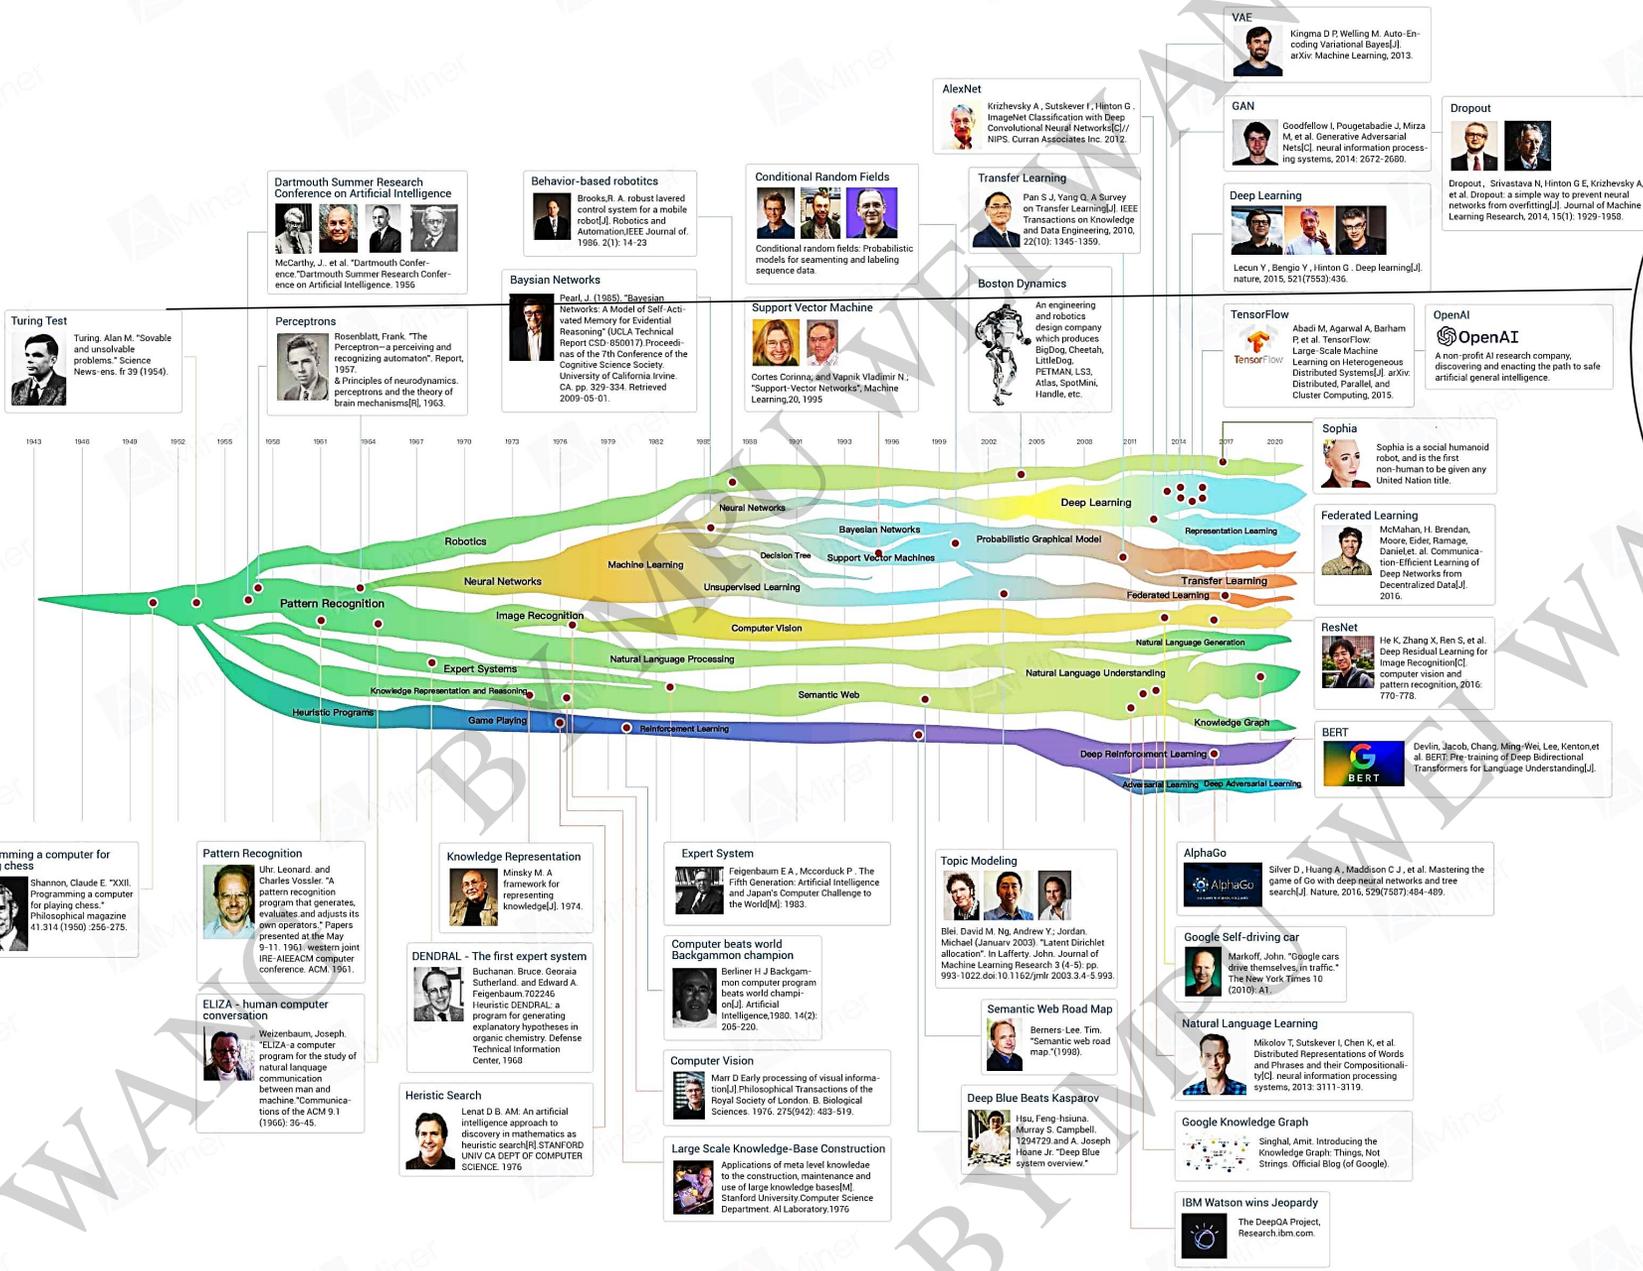
Outline

- **Brief History of AI**

Brief History of AI



Brief History of AI



Turing Test

Turing, Alan M. "Solvable and unsolvable problems." Science News-ens. fr 39 (1954).

Can machines think?

MIND
A QUARTERLY REVIEW
OF
PSYCHOLOGY AND PHILOSOPHY

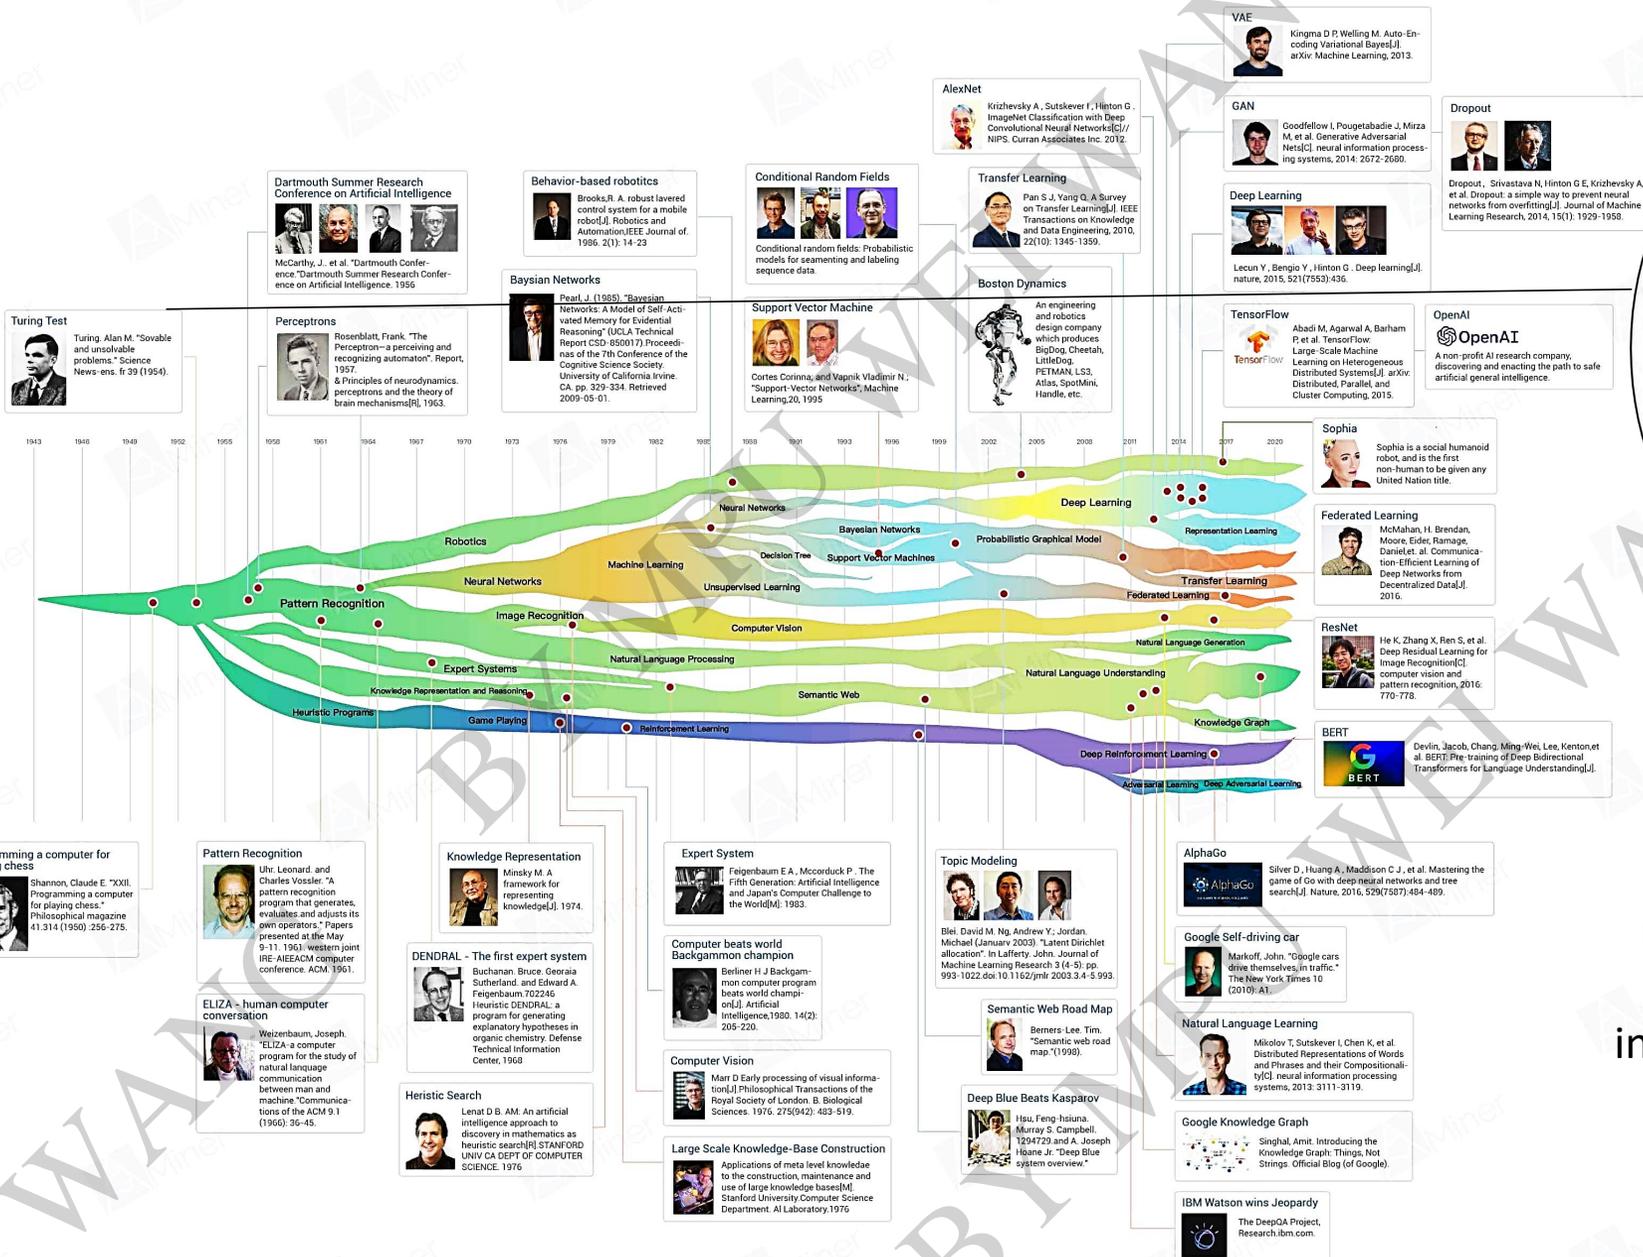
I.—COMPUTING MACHINERY AND INTELLIGENCE

By A. M. TURING

1. *The Imitation Game.*

PROPOSE to consider the question, 'Can machines think?' should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as

Brief History of AI





Turing Test

Turing, Alan M. "Solvable and unsolvable problems." *Science News-ens. fr* 39 (1954).

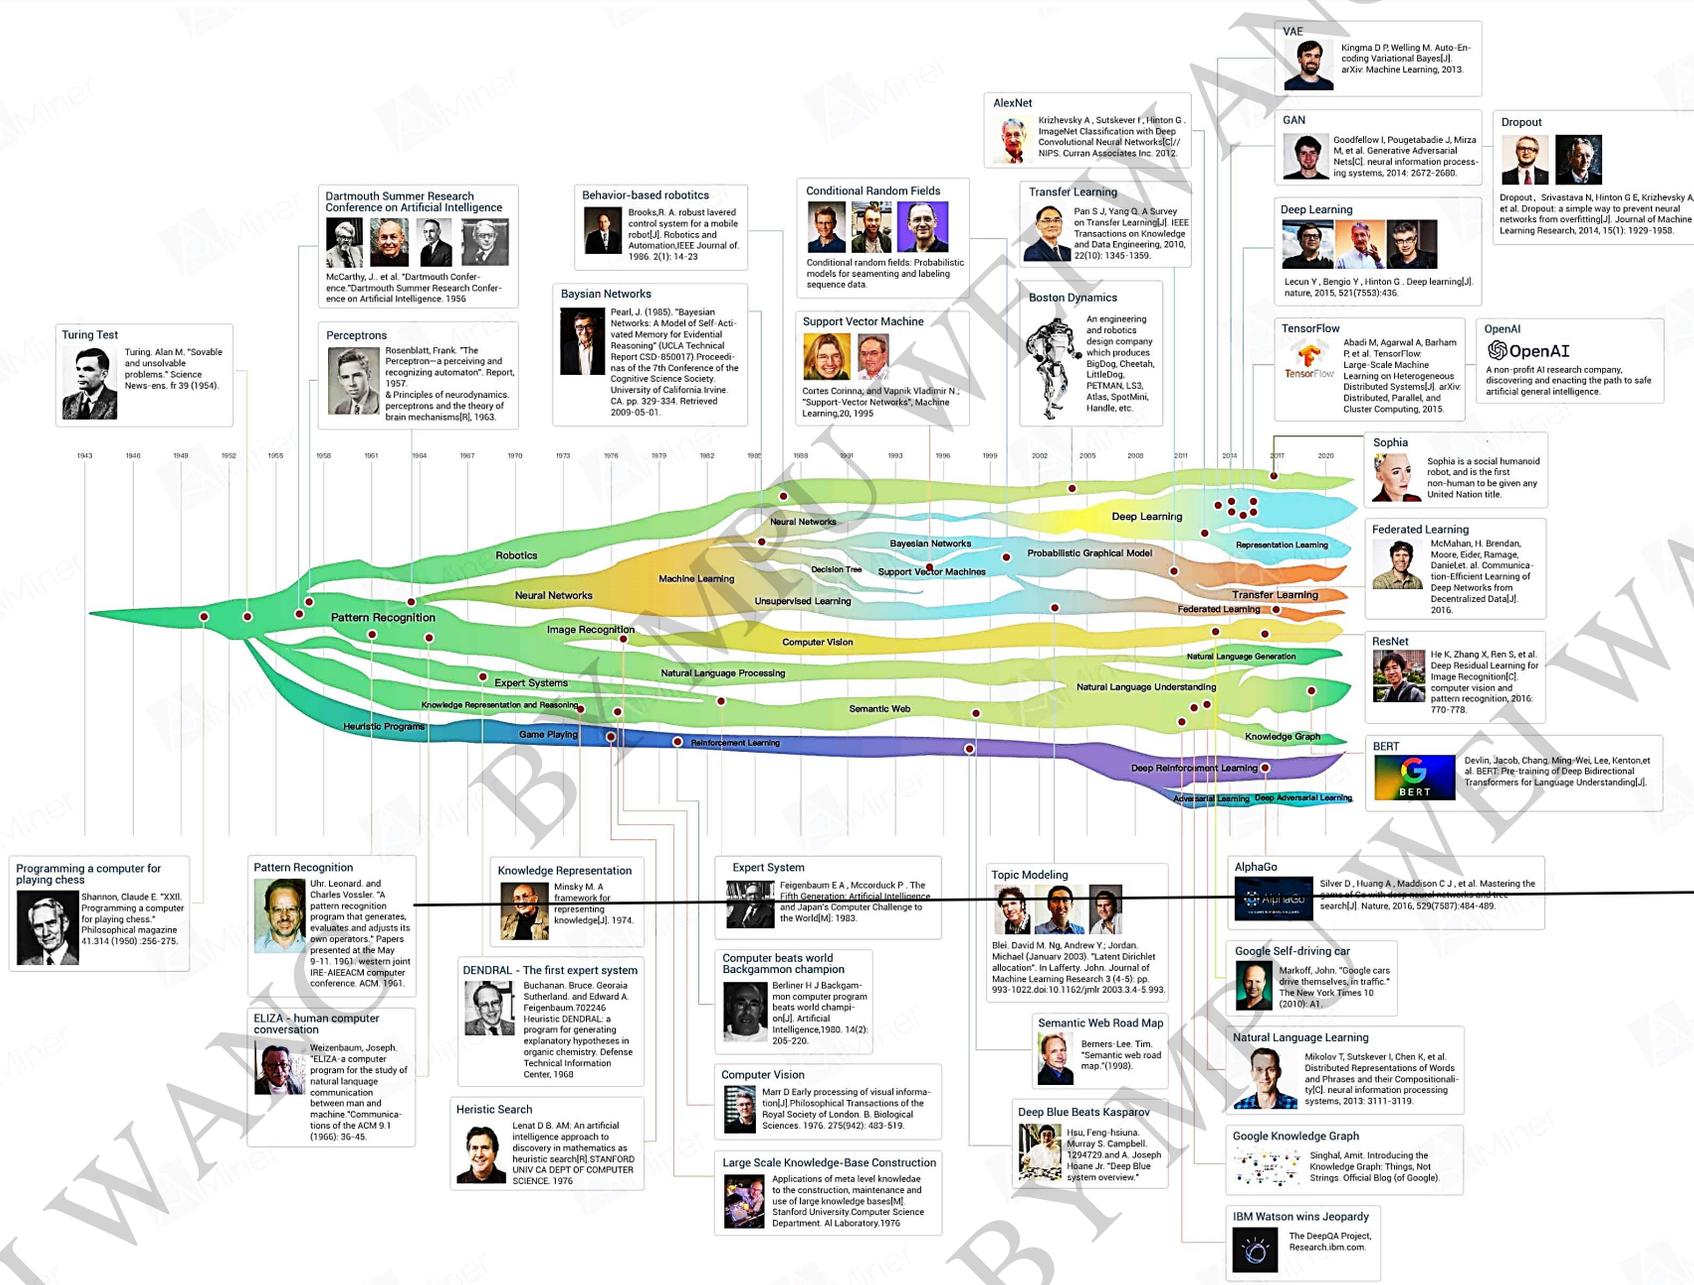
Can machines think?

1936 1946 1949

Turing Test

The Turing Test is a measure of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human.

Brief History of AI



Pattern recognition is a field within artificial intelligence and machine learning that focuses on the identification of patterns, regularities, and trends in data. The goal of pattern recognition is to classify data based on previous knowledge or statistical information derived from the data.

Pattern Recognition



Uhr, Leonard and Charles Vossler. "A pattern recognition program that generates, evaluates and adjusts its own operators." Papers presented at the May 9-11, 1961. western joint IRE-AIEEACM computer conference. ACM. 1961

Programming a computer for playing chess



Shannon, Claude E. "XXII. Programming a computer for playing chess." Science News-ens. fr 39 (1954): 41-314 (1950). 256-275.

Pattern Recognition



Uhr, Leonard and Charles Vossler. "A pattern recognition program that generates, evaluates and adjusts its own operators." Papers presented at the May 9-11, 1961. western joint IRE-AIEEACM computer conference. ACM. 1961.

ELIZA - human computer conversation



Weizenbaum, Joseph. "ELIZA - a computer program for the study of natural language communication between man and machine." Communications of the ACM 9.1 (1966). 36-45.

Knowledge Representation



Minsky, M. A framework for representing knowledge[J]. 1974.

DENDRAL - The first expert system



Buchanan, Bruce, George A. Sutherland, and Edward A. Feigenbaum 702246 Heuristic DENDRAL: a program for generating explanatory hypotheses in organic chemistry. Defense Technical Information Center, 1968

Heuristic Search



Lenat D.B. AM: An artificial intelligence approach to discovery in mathematics as heuristic search[R]. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE. 1976

Expert System



Feigenbaum E.A., McCorduck P. The Fifth Generation. Artificial Intelligence and Japan's Computer Challenge to the World[M]. 1983.

Computer beats world Backgammon champion



Berliner H.J. Backgammon computer program beats world champion[J]. Artificial Intelligence.1980. 14(2): 205-220.

Computer Vision



Marr D. Early processing of visual information[J]. Philosophical Transactions of the Royal Society of London. B. Biological Sciences. 1976. 275(942). 483-519.

Large Scale Knowledge-Base Construction



Applications of meta level knowledge to the construction, maintenance and use of large knowledge bases[M]. Stanford University Computer Science Department. AI Laboratory.1976

Topic Modeling



Blei, David M., Ng, Andrew Y., Jordan, Michael (January 2003). "Latent Dirichlet allocation". In Lafferty, John. Journal of Machine Learning Research 3 (4-5). pp. 993-1022.doi:10.1162/jmlr.2003.3.4-5.993.

Semantic Web Road Map



Berners-Lee, Tim. "Semantic web road map." (1998).

Deep Blue Beats Kasparov



Hsu, Feng-hsiang, Murray S. Campbell, 1294729 and A. Joseph Hoane Jr. "Deep Blue system overview."

AlphaGo



Silver D, Huang A, Maddison C J, et al. Mastering the game of Go with deep neural networks and tree search[J]. Nature, 2016, 529(7501):484-489.

Google Self-driving car



Markoff, John. "Google cars drive themselves in traffic." The New York Times 10 (2010): A1.

Natural Language Learning



Mikolov T, Sutskever I, Chen K, et al. Distributed Representations of Words and Phrases and their Compositionality[C]. neural information processing systems, 2013. 3111-3119.

Google Knowledge Graph



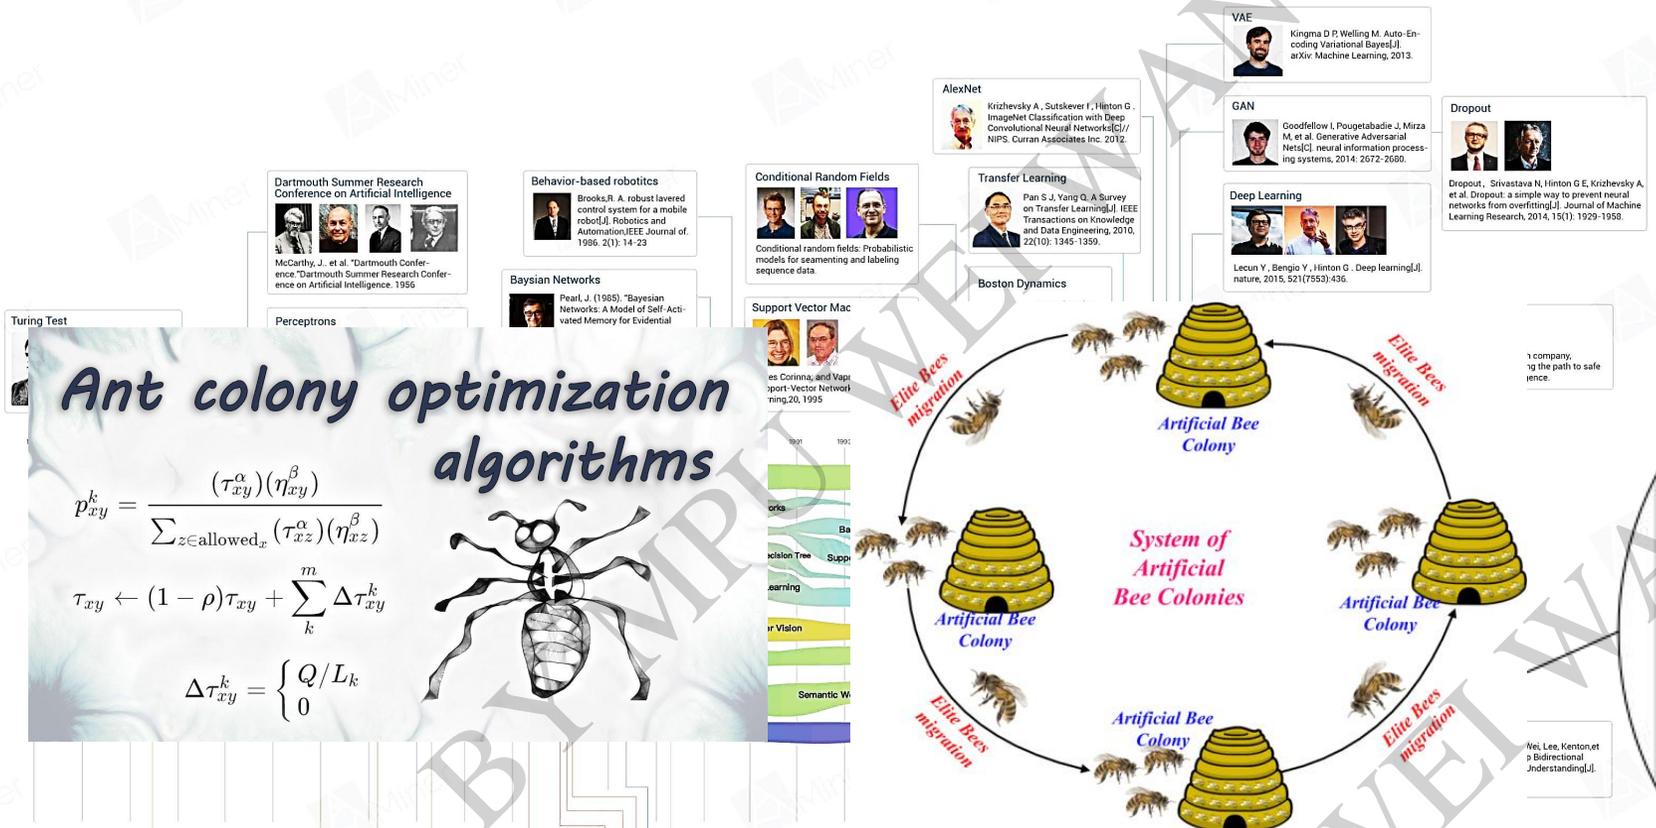
Singhal, Amit. Introducing the Knowledge Graph. Things, Not Strings. Official Blog (of Google)

IBM Watson wins Jeopardy



The DeepQA Project, Research.ibm.com.

Brief History of AI



Ant colony optimization algorithms

$$p_{xy}^k = \frac{(\tau_{xy}^\alpha)(\eta_{xy}^\beta)}{\sum_{z \in \text{allowed}_x} (\tau_{xz}^\alpha)(\eta_{xz}^\beta)}$$

$$\tau_{xy} \leftarrow (1 - \rho)\tau_{xy} + \sum_k \Delta\tau_{xy}^k$$

$$\Delta\tau_{xy}^k = \begin{cases} Q/L_k & \text{if } \text{visited} \\ 0 & \text{otherwise} \end{cases}$$

Heuristic Search

Lenat D B. AM: An artificial intelligence approach to discovery in mathematics as heuristic search [R]. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE. 1976

Technical Informatics Center, 1968

Programming a computer for playing chess
Shannon, Claude E. "XXII. Programming a computer for playing chess." Philosophical magazine 41.314 (1950): 236-275.

Pattern Recognition
Uhr, Leonard, and Charles Vossler. "A pattern recognition program that generates, evaluates and adjusts its own operators." Papers presented at the May 9-11, 1961 western joint IRE-AIEE-ACM computer conference. ACM, 1961.

ELIZA - human computer conversation
Weizenbaum, Joseph. "ELIZA - a computer program for the study of natural language communication between man and machine." Communications of the ACM 9.1 (1966): 36-45.

Knowledge Representation
Minsky, M. A. framework for representing knowledge [J]. 1974.

DENDRAL - The first expert system
Buchanan, Bruce, George A. Sutherland, and Edward A. Feigenbaum 702246. Heuristic DENDRAL: a program for generating explanatory hypotheses in organic chemistry. Defense Technical Information Center, 1968.

Heuristic Search
Lenat D B. AM: An artificial intelligence approach to discovery in mathematics as heuristic search [R]. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE. 1976.

Expert System
Feigenbaum E. A., McCorduck P. The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World [M]. 1983.

Computer beats world Backgammon champion
Berliner H. J. Backgammon computer program beats world champion [J]. Artificial Intelligence, 1980, 14(2): 205-220.

Computer Vision
Marr D. Early processing of visual information [J]. Philosophical Transactions of the Royal Society of London. B, Biological Sciences. 1976, 275(942): 483-519.

Large Scale Knowledge-Base Construction
Applications of meta level knowledge to the construction, maintenance and use of large knowledge bases [M]. Stanford University Computer Science Department. AI Laboratory, 1976.

Topic Modeling
Blei, David M., Ng, Andrew Y., Jordan, Michael (January 2003). "Latent Dirichlet allocation". In Lafferty, John. Journal of Machine Learning Research 3 (4-5). pp. 993-1022. doi:10.1162/mlr.2003.3.4-5.993.

Semantic Web Road Map
Berners-Lee, Tim. "Semantic web road map." (1998).

Deep Blue Beats Kasparov
Hsu, Feng, Hsuna, Murray S, Campbell, 1294729 and A. Joseph Hoane Jr. "Deep Blue system overview."

AlphaGo
Silver D, Huang A, Maddison C J, et al. Mastering the game of Go with deep neural networks and tree search [J]. Nature, 2016, 529(7507): 484-489.

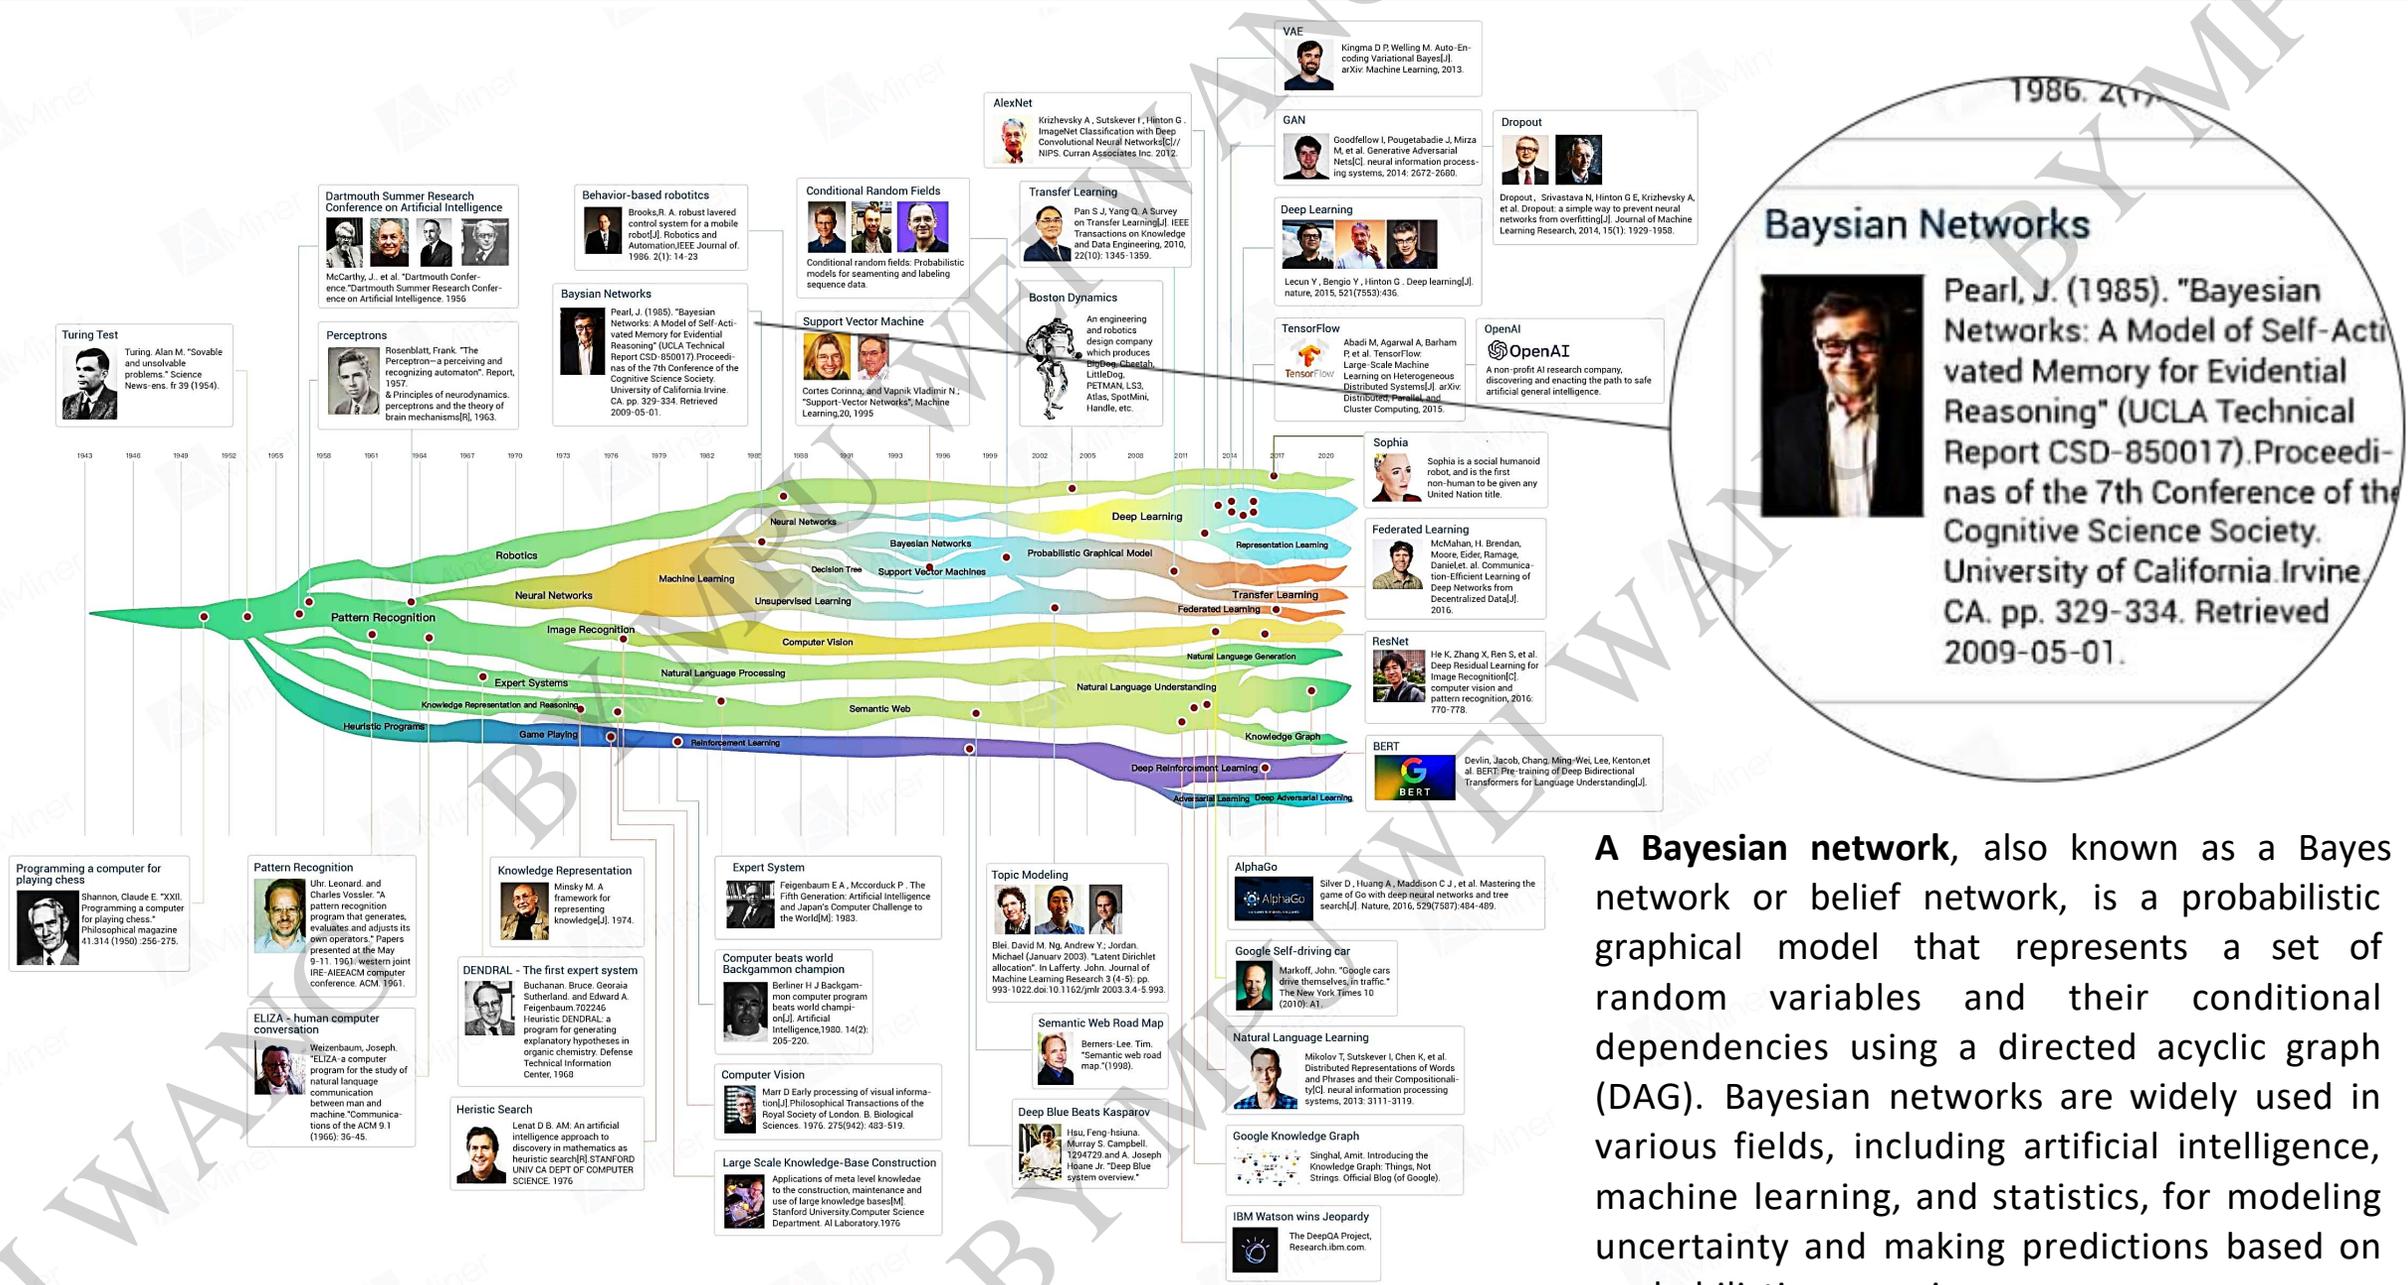
Google Self-driving car
Markoff, John. "Google cars drive themselves, in traffic." The New York Times 10 (2010): A1.

Natural Language Learning
Mikolov T, Sutskever I, Chen K, et al. Distributed Representations of Words and Phrases and their Compositionality [C]. neural information processing systems, 2013: 3111-3119.

Google Knowledge Graph
Singhal, Amit. Introducing the Knowledge Graph: Things, Not Strings. Official Blog (of Google).

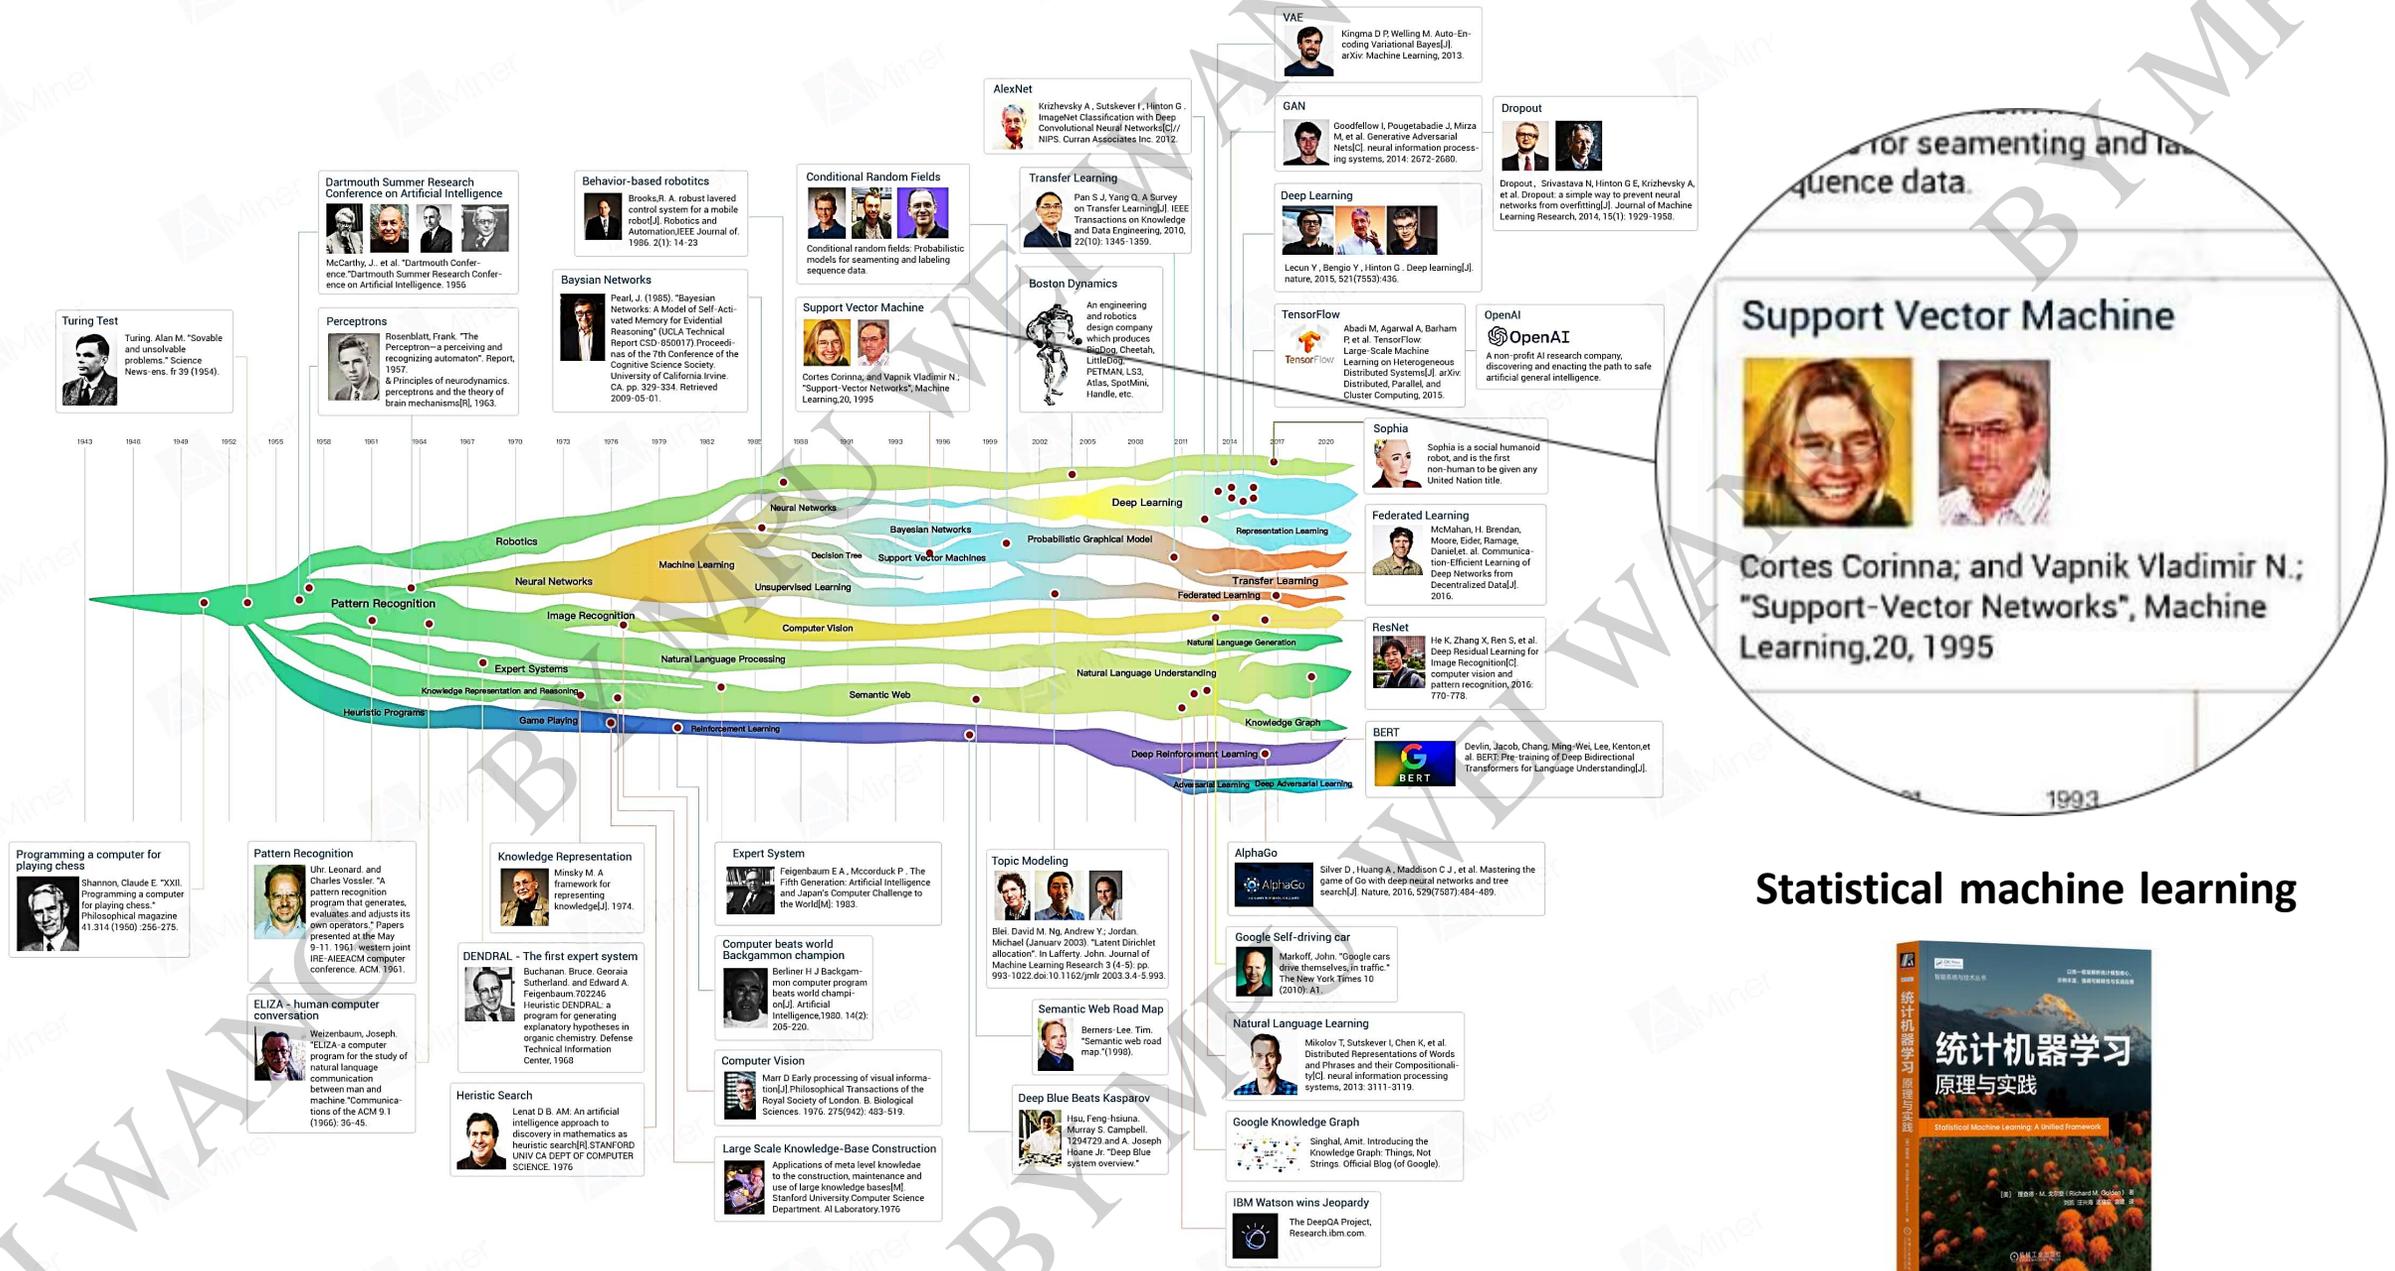
IBM Watson wins Jeopardy
The DeepQA Project, Research.ibm.com.

Brief History of AI



A Bayesian network, also known as a Bayes network or belief network, is a probabilistic graphical model that represents a set of random variables and their conditional dependencies using a directed acyclic graph (DAG). Bayesian networks are widely used in various fields, including artificial intelligence, machine learning, and statistics, for modeling uncertainty and making predictions based on probabilistic reasoning.

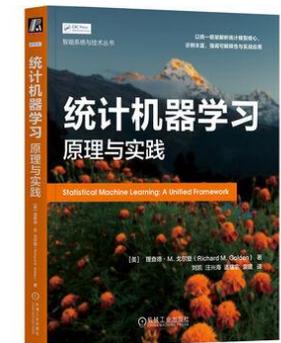
Brief History of AI



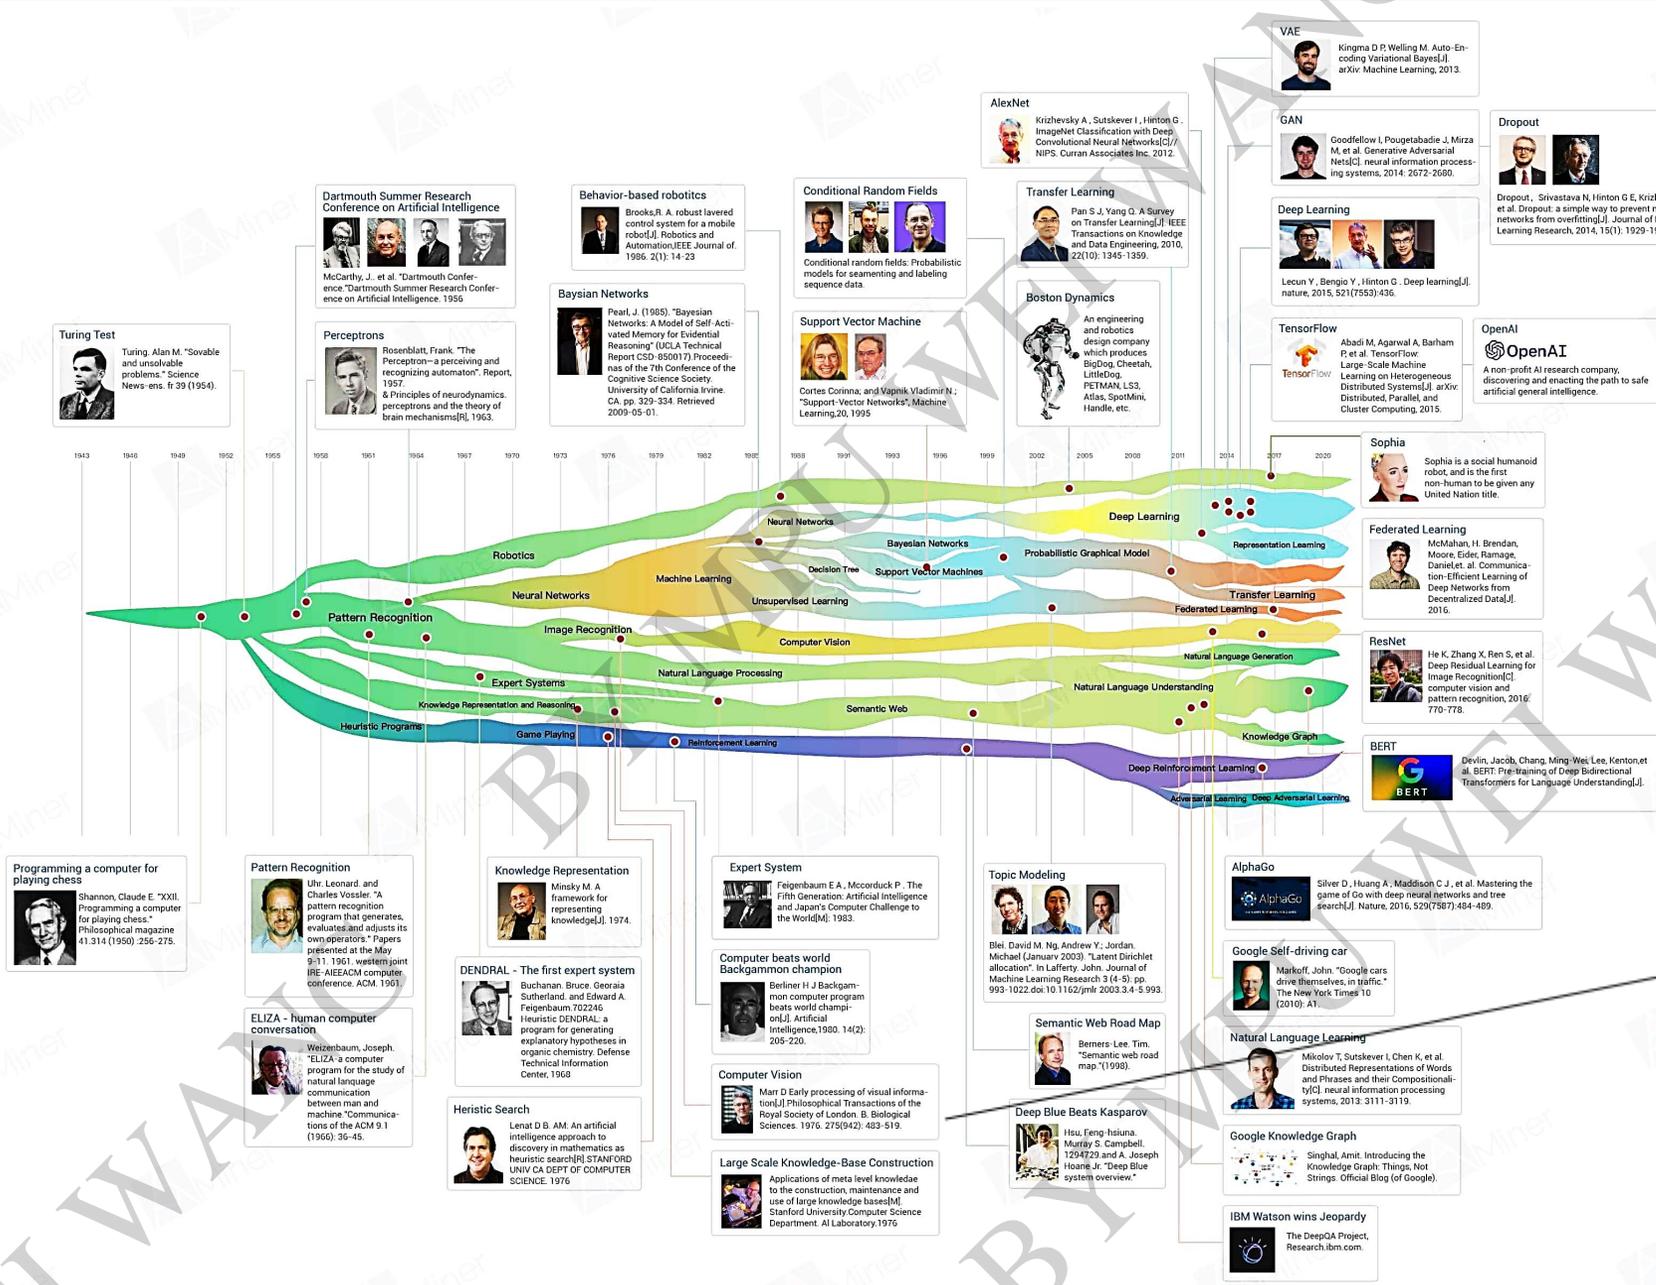
Support Vector Machine

Cortes Corinna; and Vapnik Vladimir N.; "Support-Vector Networks", Machine Learning, 20, 1995

Statistical machine learning



Brief History of AI



Computer vision is a field of artificial intelligence that enables machines to interpret and understand the visual world. It involves the development of algorithms and techniques that allow computers to extract meaningful information from digital images or videos, much like how humans interpret visual information.

on[J]. Artificial Intelligence, 1980. 14(2): 205-220.

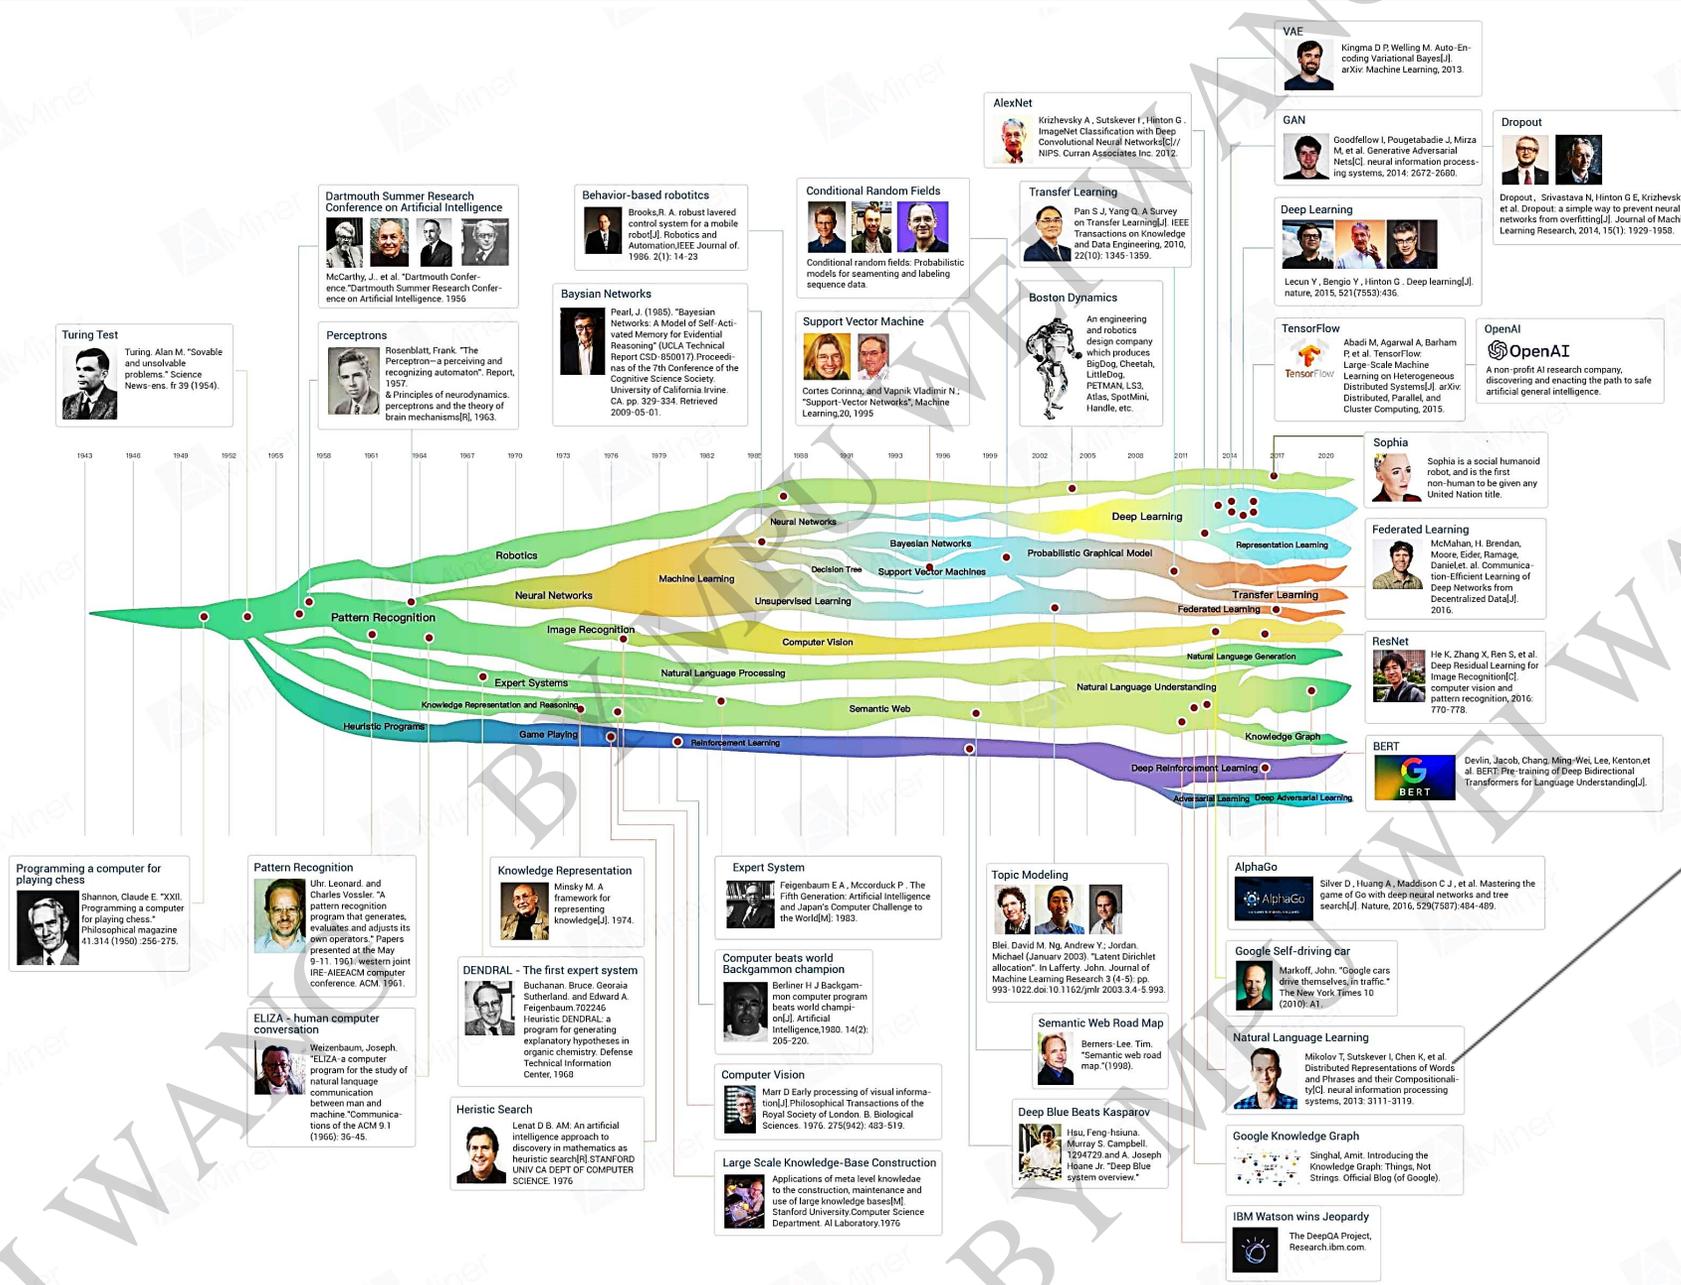
Computer Vision

Marr D Early processing of visual information[J]. Philosophical Transactions of the Royal Society of London. B. Biological Sciences. 1976. 275(942): 483-519.

Large Scale Knowledge-Base Construction

Applications of meta level knowledge to the construction, maintenance and use of large knowledge bases[M]. Stanford University Computer Science Department. AI Laboratory. 1976

Brief History of AI



It could potentially refer to the process of machines or AI systems learning natural language, which would fall under the broader category of Natural Language Processing (NLP) or Natural Language Understanding (NLU).

drive themselves, in traffic. The New York Times 10 (2010): A1.

Natural Language Learning

Mikolov T, Sutskever I, Chen K, Distributed Representations of Words and Phrases and their Compositionality[C]. neural information processing systems, 2013: 3111-3119.

Knowledge Graph

Singhal, Amit

VAE
Kingma D P, Welling M. Auto-encoding Variational Bayes[J]. arXiv: Machine Learning, 2013.

GAN
Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative Adversarial Nets[C]. neural information processing systems, 2014: 2672-2680.

Dropout
Dropout, Srivastava N, Hinton G E, Krizhevsky A, et al. Dropout: a simple way to prevent neural networks from overfitting[J]. Journal of Machine Learning Research, 2014, 15(1): 1929-1958.

TensorFlow
Abadi M, Agarwal A, Barham P, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems[J]. arXiv: Distributed, Parallel, and Cluster Computing, 2015.

OpenAI
OpenAI
A non-profit AI research company, discovering and enacting the path to safe artificial general intelligence.

Sophia
Sophia is a social humanoid robot, and is the first non-human to be given any United Nation title.

Federated Learning
McMahan, H. Brendan, Moore, Edgar, Ramage, Daniel, et al. Communication-Efficient Learning of Deep Networks from Decentralized Data[J], 2016.

ResNet
He K, Zhang X, Ren S, et al. Deep Residual Learning for Image Recognition[C]. computer vision and pattern recognition, 2016: 770-778.

BERT
Devlin, Jacob, Chang, Ming-Wei, Lee, Kenton, et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding[J].

AlphaGo
Silver D, Huang A, Maddison C J, et al. Mastering the game of Go with deep neural networks and tree search[J]. Nature, 2016, 529(7597): 484-489.

Google Self-driving car
Markoff, John. "Google cars drive themselves, in traffic." The New York Times 10 (2010): A1.

Natural Language Learning
Mikolov T, Sutskever I, Chen K, et al. Distributed Representations of Words and Phrases and their Compositionality[C]. neural information processing systems, 2013: 3111-3119.

Google Knowledge Graph
Singhal, Amit. Introducing the Knowledge Graph: Things, Not Strings. Official Blog (of Google).

AlexNet
Krizhevsky A, Sutskever I, Hinton G. ImageNet Classification with Deep Convolutional Neural Networks[C]. NIPS. Curran Associates Inc. 2012.

Transfer Learning
Pan S J, Yang G. A Survey on Transfer Learning[J]. IEEE Transactions on Knowledge and Data Engineering, 2010, 22(10): 1345-1359.

Conditional Random Fields
Conditional random fields: Probabilistic models for segmenting and labeling sequence data.

Support Vector Machine
Cortes Corinna, and Vapnik Vladimir N. "Support-Vector Networks". Machine Learning, 20, 1995.

Boston Dynamics
An engineering and robotics design company which produces BigDog, Cheetah, LittleDog, PETMAN, LS3, Atlas, SpotMini, Handle, etc.

Expert System
Feigenbaum E A, McCorduck P. The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World[M]. 1983.

Computer beats world Backgammon champion
Berliner H J. Backgammon computer program beats world champion[J]. Artificial Intelligence, 1980, 14(2): 205-220.

Computer Vision
Marr D. Early processing of visual information[J]. Philosophical Transactions of the Royal Society of London. B, Biological Sciences. 1976, 275(942): 483-519.

Large Scale Knowledge-Base Construction
Applications of meta-level knowledge to the construction, maintenance and use of large knowledge bases[M]. Stanford University Computer Science Department. AI Laboratory. 1976.

Knowledge Representation
Minsky M. A framework for representing knowledge[J]. 1974.

DENDRAL - The first expert system
Buchanan Bruce, Goelaia Sutherland, and Edward A. Feigenbaum 702246 Heuristic DENDRAL: a program for generating explanatory hypotheses in organic chemistry. Defense Technical Information Center, 1968.

Heuristic Search
Lenat D B. AM: An artificial intelligence approach to discovery in mathematics as heuristic search[R]. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE. 1976.

Pattern Recognition
Uhr, Leonard and Charles Vossler. "A pattern recognition program that generates, evaluates and adjusts its own operators." Papers presented at the May 9-11, 1961 western joint IRE-AIEACM computer conference. ACM, 1961.

ELIZA - human computer conversation
Weizenbaum, Joseph. "ELIZA-a computer program for the study of natural language communication between man and machine." Communications of the ACM 9.1 (1966): 36-45.

Perceptrons
Rosenblatt, Frank. "The Perceptron—a perceiving and recognizing automaton". Report, 1957

Bayesian Networks
Pearl, J. (1985). "Bayesian Networks: A Model of Self-Activated Memory for Evidential Reasoning" (UCLA Technical Report CSD-850017). Proceedings of the 7th Conference of the Cognitive Science Society, University of California Irvine. CA. pp. 329-334. Retrieved 2009-05-01.

Behavior-based robotics
Brooks, R. A. robust layered control system for a mobile robot[J]. Robotics and Automation, IEEE Journal of, 1996, 2(1): 14-23.

Programming a computer for playing chess
Shannon, Claude E. "XXII. Programming a computer for playing chess." Philosophical magazine 41.314 (1950): 256-275.

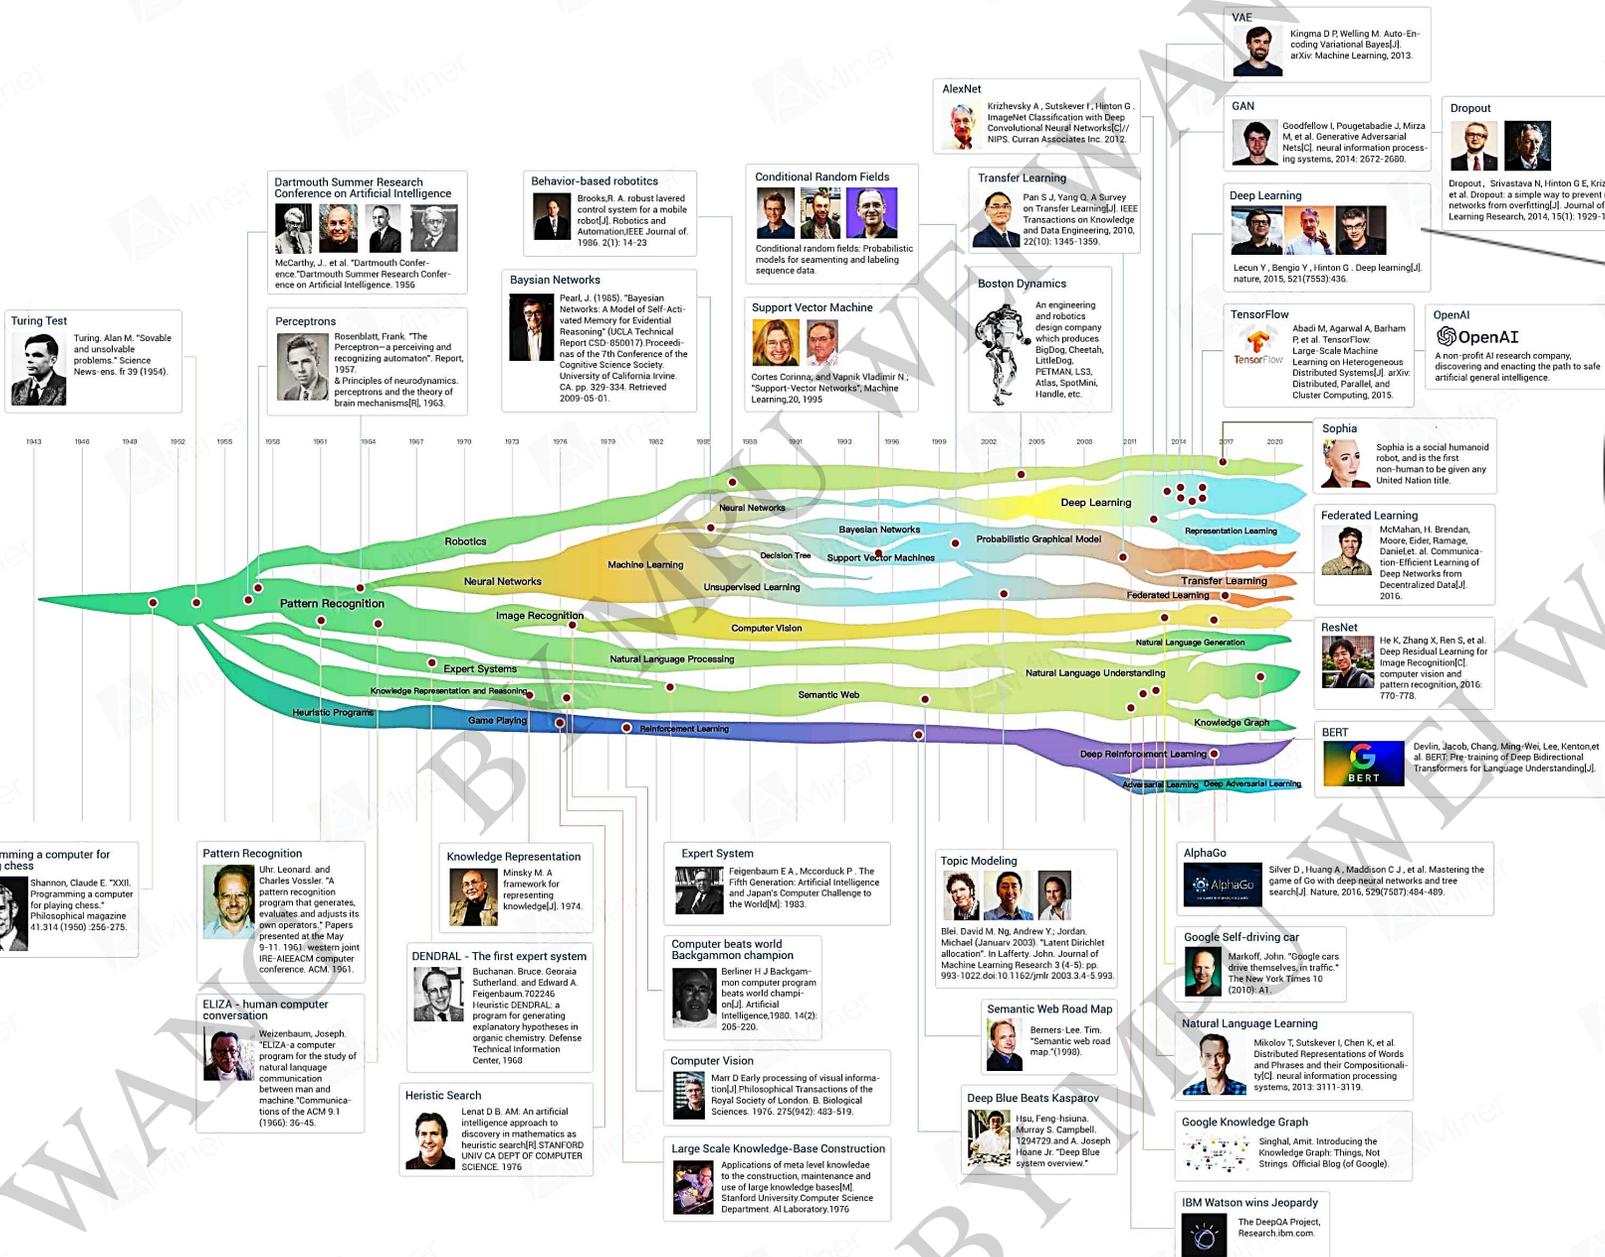
Pattern Recognition
McCarthy, J. et al. "Dartmouth Summer Research Conference on Artificial Intelligence. 1956"

Perceptrons
Rosenblatt, Frank. "The Perceptron—a perceiving and recognizing automaton". Report, 1957

Bayesian Networks
Pearl, J. (1985). "Bayesian Networks: A Model of Self-Activated Memory for Evidential Reasoning" (UCLA Technical Report CSD-850017). Proceedings of the 7th Conference of the Cognitive Science Society, University of California Irvine. CA. pp. 329-334. Retrieved 2009-05-01.

Behavior-based robotics
Brooks, R. A. robust layered control system for a mobile robot[J]. Robotics and Automation, IEEE Journal of, 1996, 2(1): 14-23.

Brief History of AI



Deep Learning

Lecun Y, Bengio Y, Hinton G. Deep learning[J]. nature, 2015, 521(7553):436.

Abadi M, Agarwal

TensorFlow

Abadi M, Agarwal A, Barham P, et al. TensorFlow: Large Scale Machine Learning on Heterogeneous Distributed Systems[J]. arXiv: Distributed, Parallel, and Cluster Computing, 2015.

OpenAI

OpenAI
A non-profit AI research company, discovering and enacting the path to safe artificial general intelligence.

Sophia

Sophia is a social humanoid robot, and is the first non-human to be given any United Nation title.

Federated Learning

McMahan, H. Brendan, Moore, Edgar, Ramage, Daniel, et al. Communication-Efficient Learning of Deep Networks from Decentralized Data[J]. 2016.

ResNet

He K, Zhang X, Ren S, et al. Deep Residual Learning for Image Recognition[C]. computer vision and pattern recognition, 2016: 770-778.

BERT

Devlin, Jacob, Chang, Ming-Wei, Lee, Kenton, et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding[J].

THE NOBEL PRIZE IN PHYSICS 2024

John J. Hopfield **Geoffrey E. Hinton**

"for foundational discoveries and inventions that enable machine learning with artificial neural networks"

THE ROYAL SWEDISH ACADEMY OF SCIENCES

Turing Test

Turing, Alan M. "Solvable and unsolvable problems." Science News-ns. fr 39 (1954).

Perceptrons

Rosenblatt, Frank. "The Perceptron—a perceiving and recognizing automaton." Report, 1957 & Principles of neurodynamics, perceptrons and the theory of brain mechanisms[R], 1963.

Behavior-based robotics

Brooks, R. A. robust layered control system for a mobile robot[J]. Robotics and Automation, IEEE Journal of, 1996, 2(1): 14-23

Bayesian Networks

Pearl, J. (1985). "Bayesian Networks: A Model of Self-Activated Memory for Evidential Reasoning" (UCLA Technical Report CSD-850017). Proceedings of the 7th Conference of the Cognitive Science Society, University of California Irvine. CA. pp. 329-334. Retrieved 2009-05-01.

Conditional Random Fields

Conditional random fields: Probabilistic models for segmenting and labeling sequence data

Support Vector Machine

Cortes Corinna, and Vapnik Vladimir N. "Support-Vector Networks". Machine Learning, 20, 1995.

AlexNet

Krizhevsky A., Sutskever I., Hinton G. ImageNet Classification with Deep Convolutional Neural Networks[C]// NIPS. Curran Associates Inc. 2012.

Transfer Learning

Pan S. J. Yang G. A Survey on Transfer Learning[J]. IEEE Transactions on Knowledge and Data Engineering, 2010, 22(10): 1345-1359.

Boston Dynamics

An engineering and robotics design company which produces BigDog, CheetaH, LittleDog, PETMAN, LS3, Atlas, SpotMini, Handle, etc.

VAE

Kingma D P, Welling M. Auto-encoding Variational Bayes[J]. arXiv: Machine Learning, 2013.

GAN

Goodfellow I, Pougetabadi J, Mirza M, et al. Generative Adversarial Nets[C]. neural information processing systems, 2014: 2672-2680.

Dropout

Dropout., Srivastava N, Hinton G E, Krizhevsky A, et al. Dropout: a simple way to prevent neural networks from overfitting[J]. Journal of Machine Learning Research, 2014, 15(1): 1929-1958.

Deep Learning

Lecun Y, Bengio Y, Hinton G. Deep learning[J]. nature, 2015, 521(7553):436.

Programming a computer for playing chess

Shannon, Claude E. "XXII. Programming a computer for playing chess." Philosophical magazine 41.314 (1950): 256-275.

Pattern Recognition

Uhr, Leonard and Charles Vossler. "A pattern recognition program that generates, evaluates and adjusts its own operators." Papers presented at the May 9-11, 1961 western joint IRE-AIEE-ACM computer conference. ACM, 1961.

ELIZA - human computer conversation

Weizenbaum, Joseph. "ELIZA-a computer program for the study of natural language communication between man and machine." Communications of the ACM 9.1 (1966): 36-45.

Knowledge Representation

Minsky M. A framework for representing knowledge[J]. 1974.

DENDRAL - The first expert system

Buchanan Bruce, Goelaia Sutherland, and Edward A. Feigenbaum 702246 Heuristic DENDRAL: a program for generating explanatory hypotheses in organic chemistry. Defense Technical Information Center, 1968

Heristic Search

Lenat D B. AM: An artificial intelligence approach to discovery in mathematics as heuristic search[R]. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE, 1976

Expert System

Feigenbaum E. A., McCorduck P. The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World[M]. 1983.

Computer beats world Backgammon champion

Berliner H. J. Backgammon computer program beats world champion[J]. Artificial Intelligence, 1980, 14(2): 205-220.

Computer Vision

Marr D. Early processing of visual information[J]. Philosophical Transactions of the Royal Society of London. B, Biological Sciences. 1976, 275(942): 483-519.

Large Scale Knowledge-Base Construction

Applications of meta-level knowledge to the construction, maintenance and use of large knowledge bases[M]. Stanford University Computer Science Department. AI Laboratory. 1976

Topic Modeling

Blei David M, Ng Andrew Y, Jordan Michael (January 2003). "Latent Dirichlet allocation". In Lafferty, John. Journal of Machine Learning Research 3 (4-5). pp. 993-1022. doi:10.1162/mlr.2003.3.4-5.993.

Semantic Web Road Map

Berners-Lee, Tim. "Semantic web road map." (1998).

Deep Blue Beats Kasparov

Hsu, Feng Hsiang, Murray S, Campbell, 1294729 and A. Joseph Hoane Jr. "Deep Blue system overview."

AlphaGo

Silver D, Huang A, Maddison C J, et al. Mastering the game of Go with deep neural networks and tree search[J]. Nature, 2016, 529(7527): 484-489.

Google Self-driving car

Markoff, John. "Google cars drive themselves, in traffic." The New York Times 10 (2010): A1.

Natural Language Learning

Mikolov T, Sutskever I, Chen K, et al. Distributed Representations of Words and Phrases and their Compositionality[C]. neural information processing systems, 2013: 3111-3119.

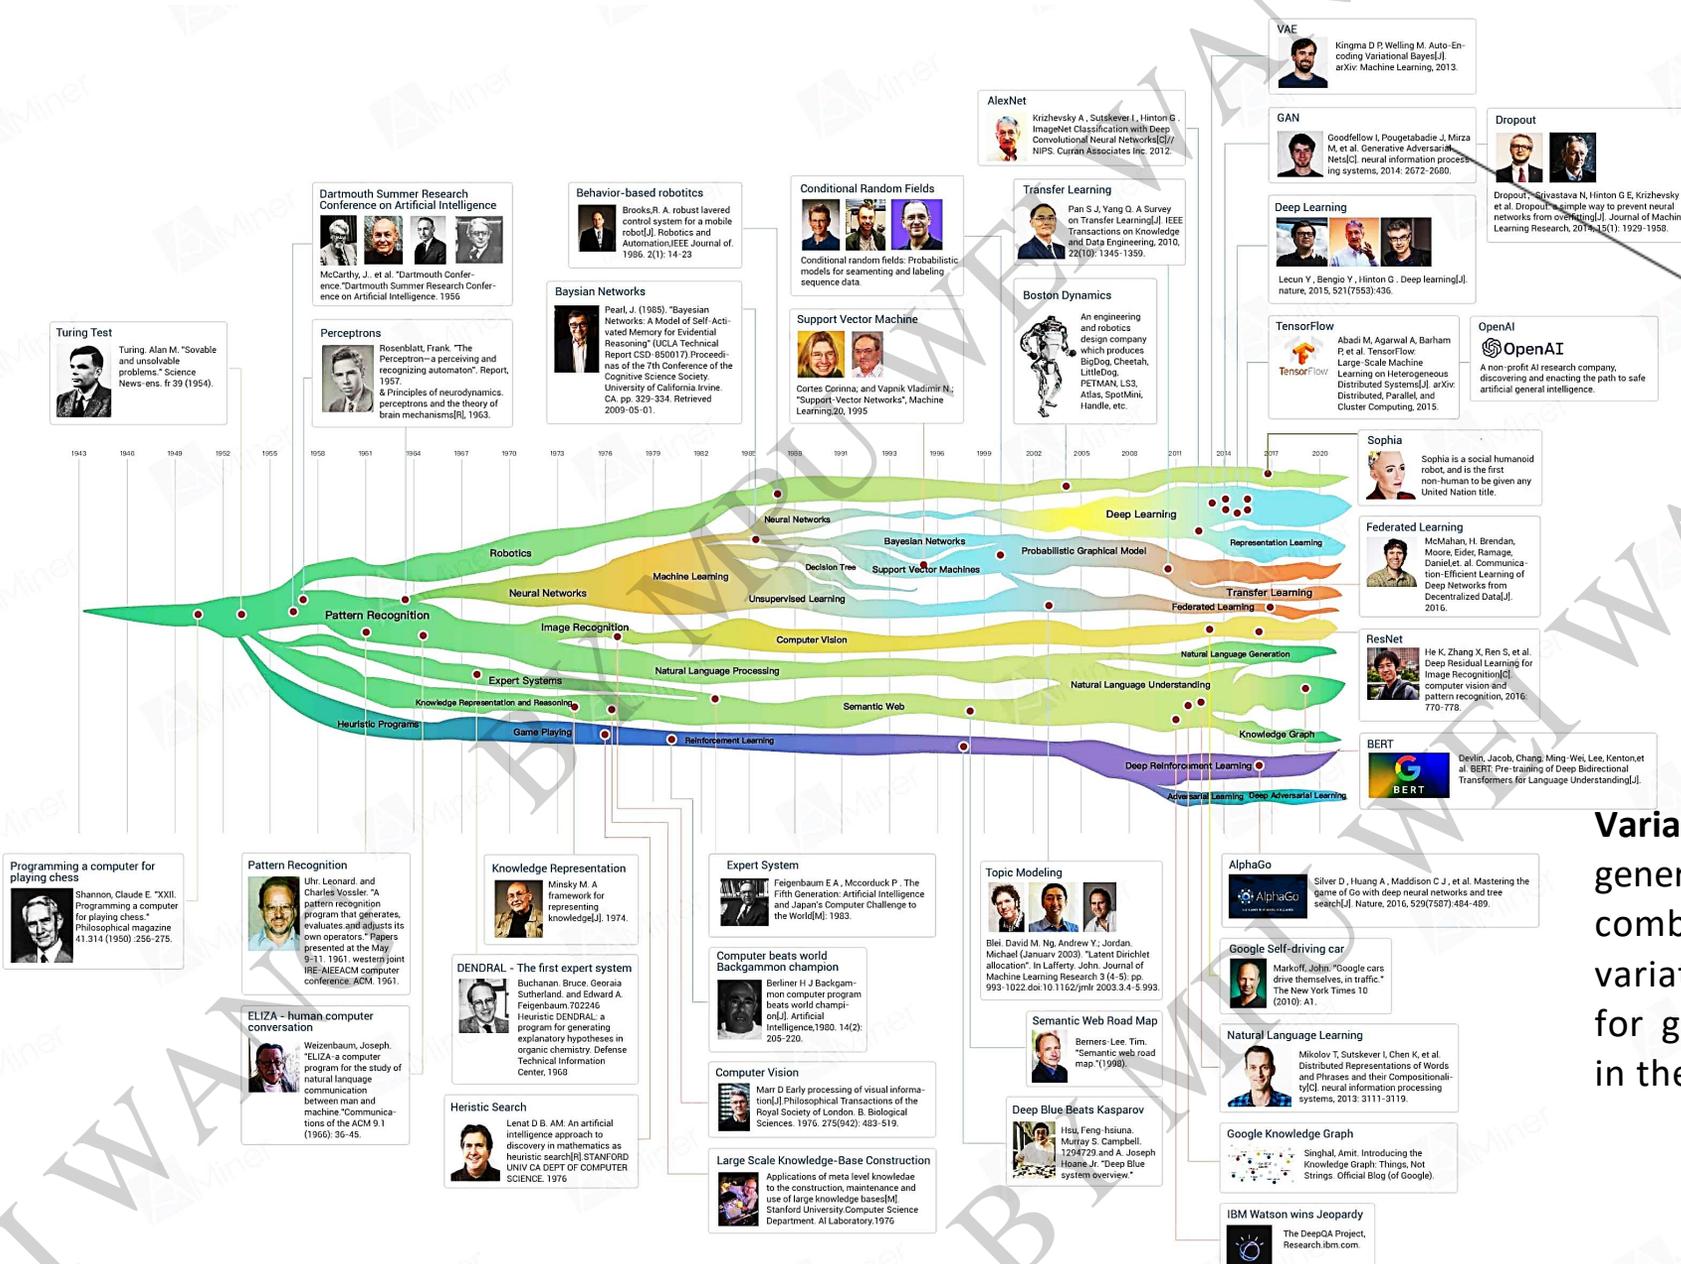
Google Knowledge Graph

Singhal, Amit. Introducing the Knowledge Graph. Things, Not Strings. Official Blog (of Google)

IBM Watson wins Jeopardy

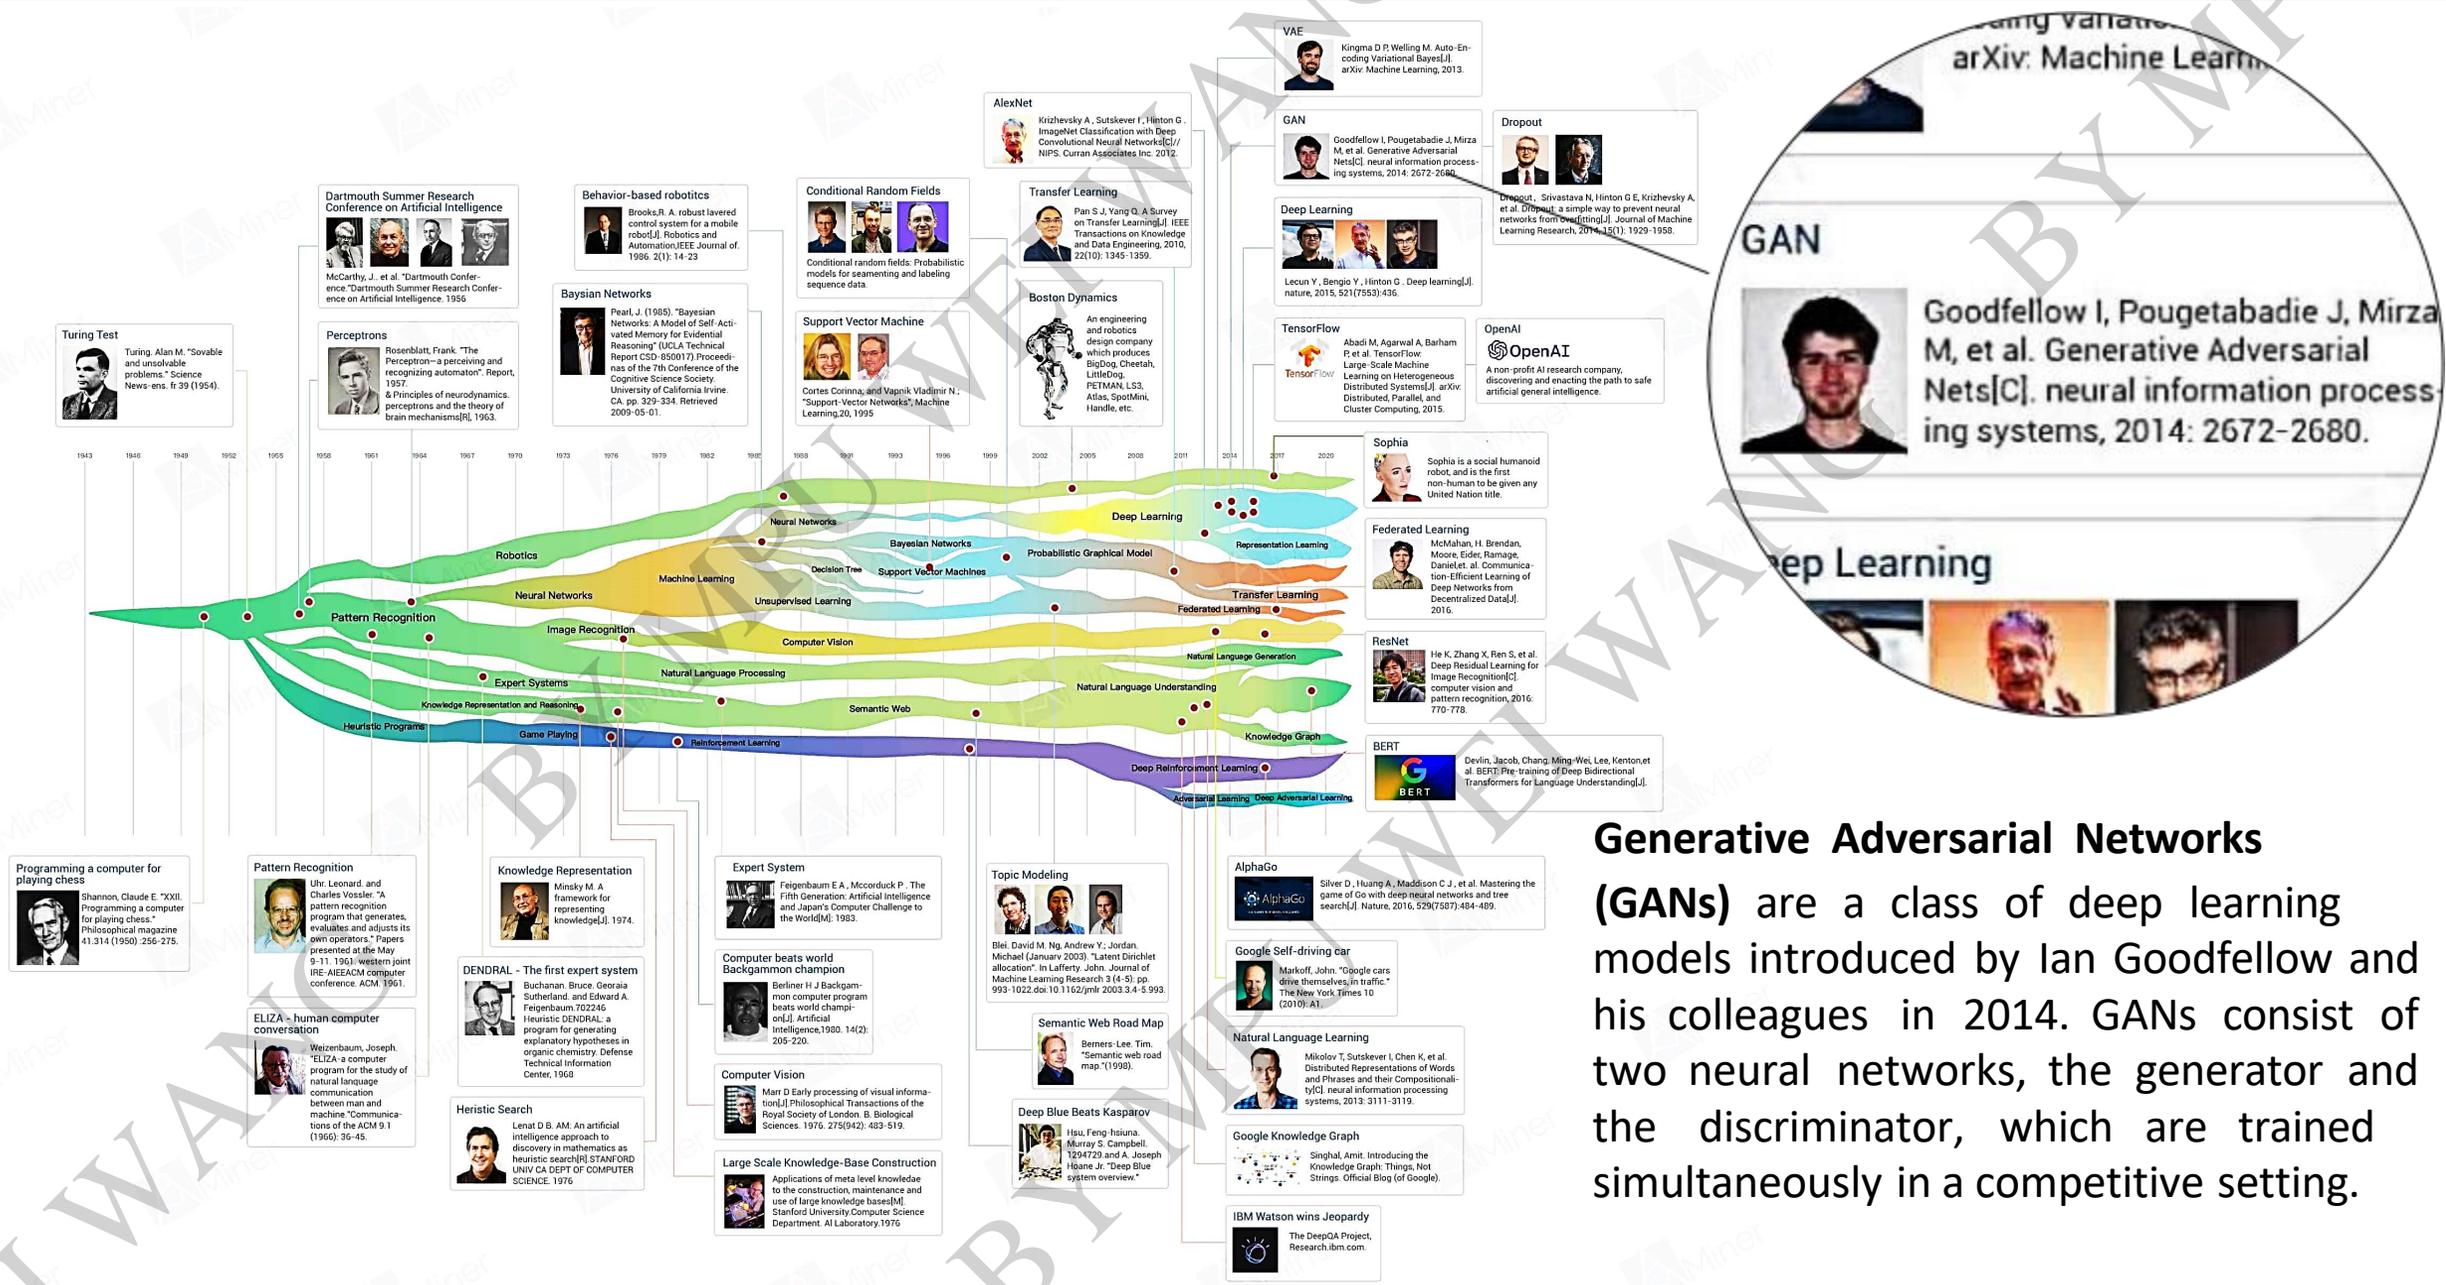
The DeepQA Project, Research.ibm.com.

Brief History of AI



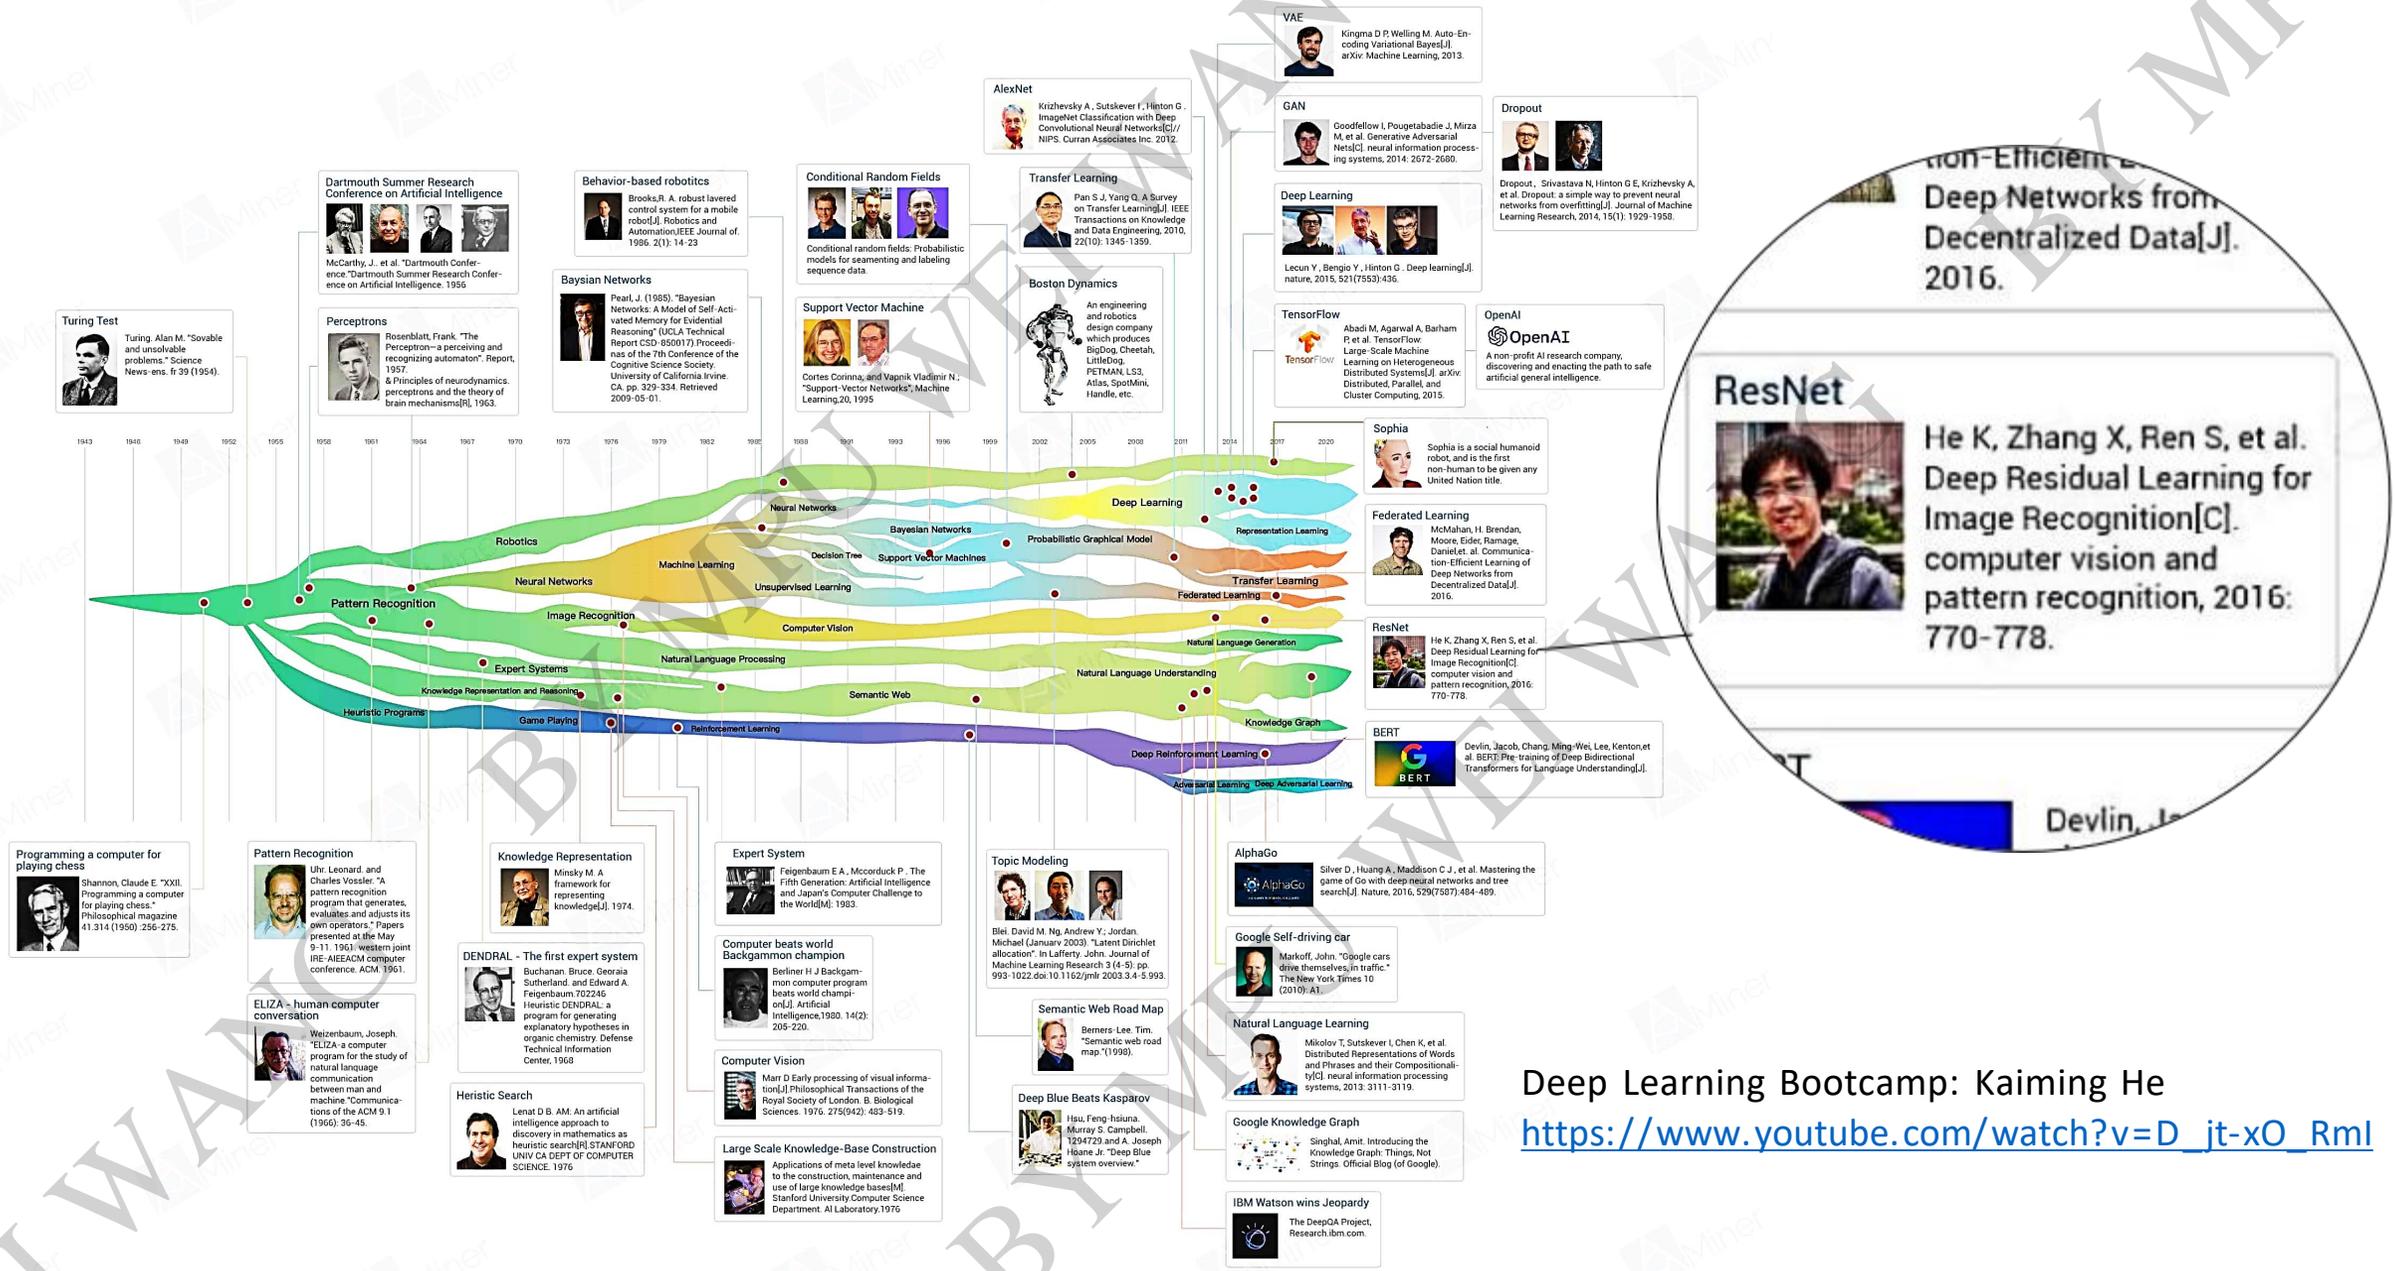
Variational Autoencoder (VAE) is a type of generative model in machine learning that combines elements of deep learning and variational inference. VAEs are commonly used for generating new data samples, particularly in the field of unsupervised learning.

Brief History of AI



Generative Adversarial Networks (GANs) are a class of deep learning models introduced by Ian Goodfellow and his colleagues in 2014. GANs consist of two neural networks, the generator and the discriminator, which are trained simultaneously in a competitive setting.

Brief History of AI



Deep Networks from Decentralized Data [J]. 2016.

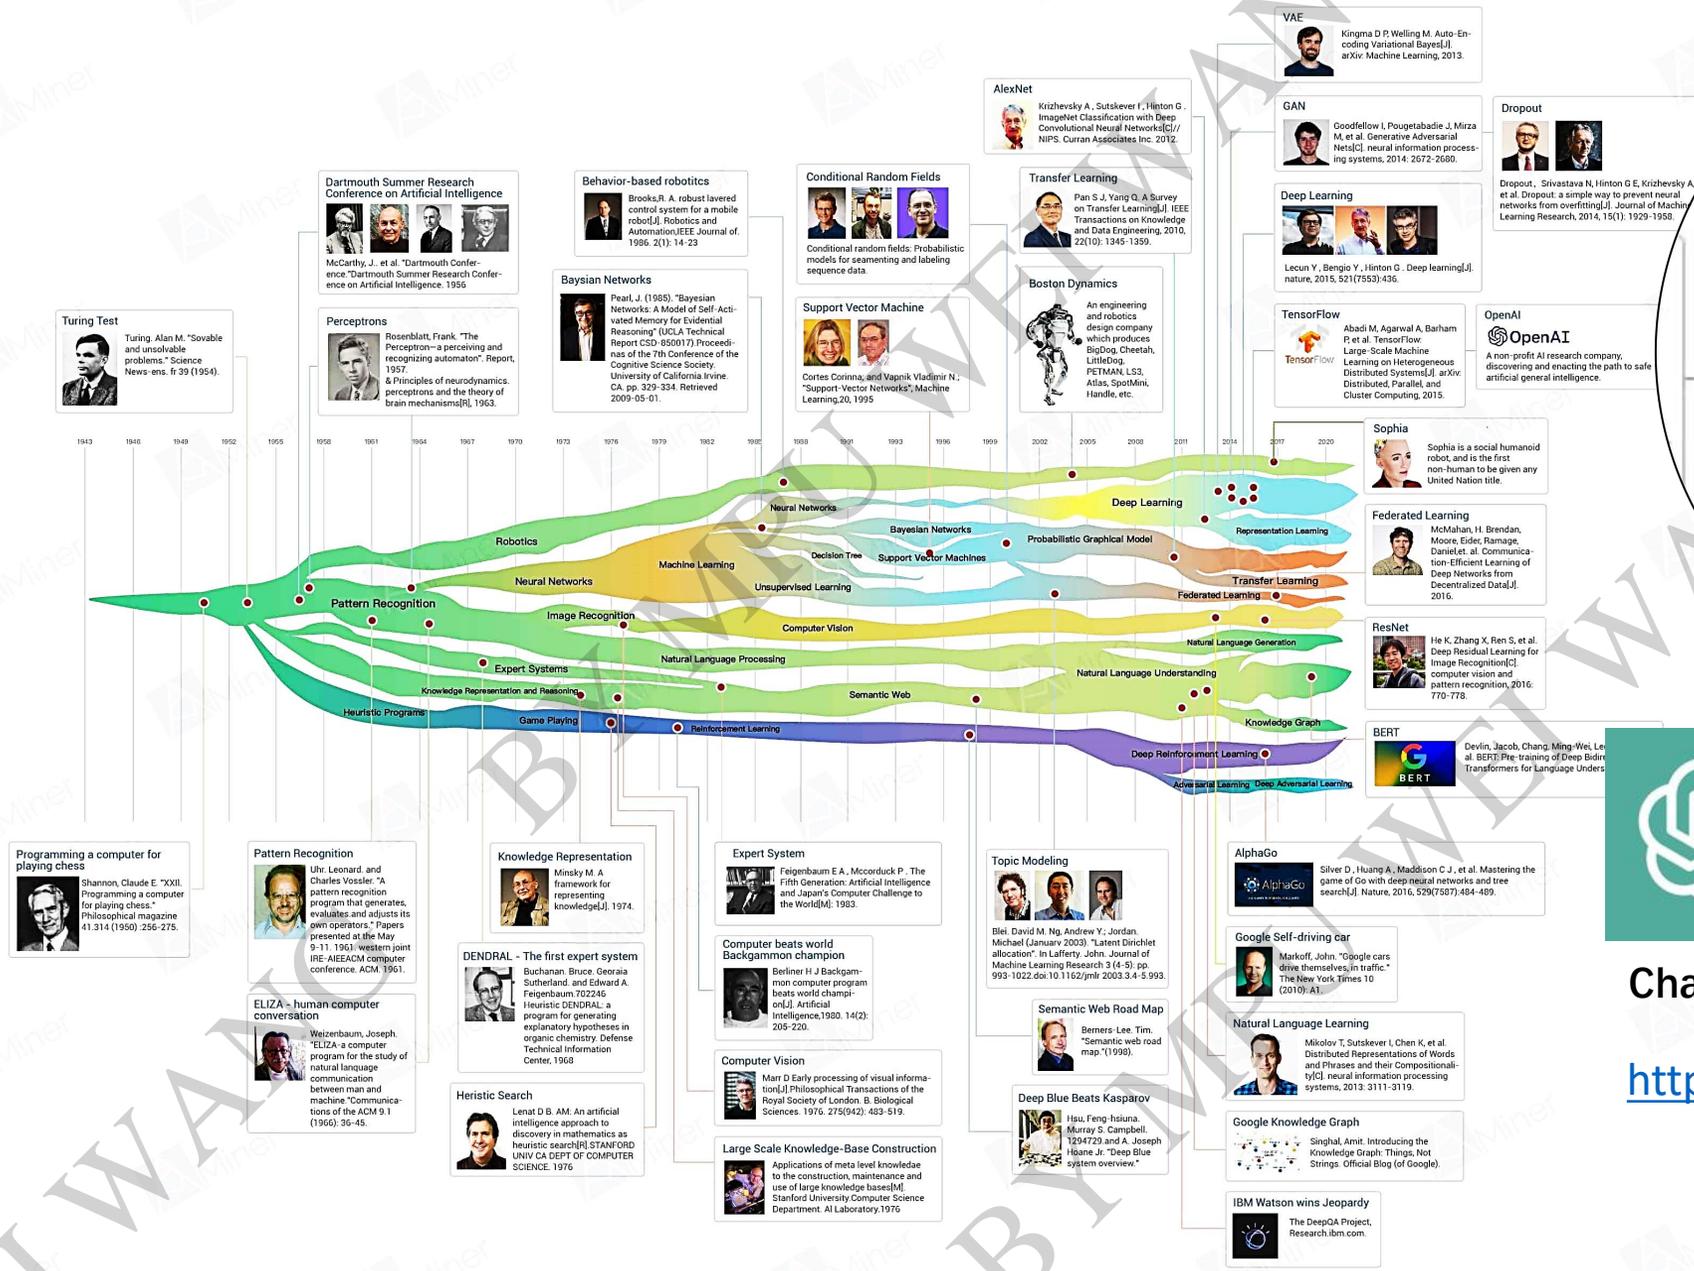
ResNet

He K, Zhang X, Ren S, et al. Deep Residual Learning for Image Recognition [C]. computer vision and pattern recognition, 2016: 770-778.

Devlin, J.

Deep Learning Bootcamp: Kaiming He
https://www.youtube.com/watch?v=D_jt-xO_Rml

Brief History of AI



OpenAI
OpenAI
 A non-profit AI research company, discovering and enacting the path to safe artificial general intelligence.



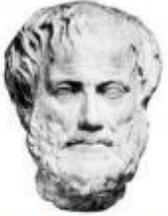
ChatGPT

Midjourney

Stable Diffusion

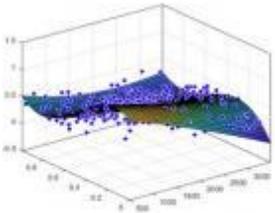
<https://poe.com/>

Three Intellectual Traditions



symbolic AI

Symbolic AI, also known as rule-based or logic-based AI, focuses on using logical rules and symbolic representations to perform reasoning and problem-solving tasks. It involves the use of knowledge representation, expert systems, and logical inference to manipulate symbols and make logical deductions. Symbolic AI emphasizes explicit knowledge representation and rule-based decision-making.



statistical AI

Statistical AI, also referred to as machine learning, involves the use of statistical models and algorithms to enable computers to learn from data and make predictions or decisions. It relies on analyzing patterns and relationships in large datasets to extract meaningful insights and generalize from examples. Statistical AI encompasses techniques such as regression, classification, clustering, and probabilistic models.



neural AI

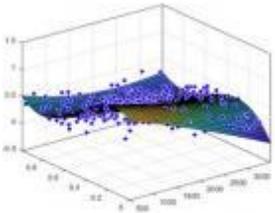
Neural AI, also known as deep learning, is a subset of machine learning that focuses on artificial neural networks inspired by the structure and function of the human brain. It involves training deep neural networks with multiple layers of interconnected nodes (neurons) to learn hierarchical representations of data. Neural AI excels in tasks such as image and speech recognition, natural language processing, and pattern recognition. It has achieved remarkable success in recent years due to advancements in computational power and the availability of large-scale datasets.

Three Intellectual Traditions



symbolic AI

Symbolic AI, also known as **rule-based** or **logic-based AI**, focuses on using logical rules and symbolic representations to perform reasoning and problem-solving tasks. It involves the use of knowledge representation, expert systems, and logical inference to manipulate symbols and make logical deductions. Symbolic AI emphasizes explicit **knowledge representation and rule-based decision-making**.



statistical AI

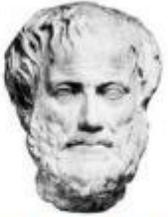
Statistical AI, also referred to as **machine learning**, involves the use of statistical models and algorithms to enable computers to learn from data and make predictions or decisions. It relies on analyzing patterns and relationships in large datasets to extract meaningful insights and generalize from examples. Statistical AI encompasses techniques such as **regression, classification, clustering, and probabilistic models**.



neural AI

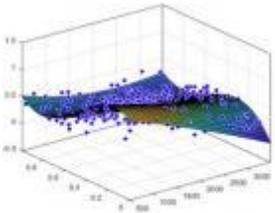
Neural AI, also known as **deep learning**, is a subset of machine learning that focuses on artificial neural networks inspired by the structure and function of the human brain. It involves training deep neural networks with multiple layers of interconnected nodes (neurons) to learn hierarchical representations of data. Neural AI excels in tasks such as image and speech recognition, natural language processing, and pattern recognition. It has achieved remarkable success in recent years due to advancements in computational power and the availability of large-scale datasets.

Three Intellectual Traditions



symbolic AI

Symbolic AI, also known as **rule-based** or **logic-based AI**, focuses on using logical rules and symbolic representations to perform reasoning and problem-solving tasks. It involves the use of knowledge representation, expert systems, and logical inference to manipulate symbols and make logical deductions. Symbolic AI emphasizes explicit **knowledge representation** and **rule-based decision-making**.



statistical AI

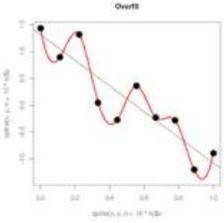
Statistical AI, also referred to as **machine learning**, involves the use of statistical models and algorithms to enable computers to learn from data and make predictions or decisions. It relies on analyzing patterns and relationships in large datasets to extract meaningful insights and generalize from examples. Statistical AI encompasses techniques such as **regression, classification, clustering, and probabilistic models**.



neural AI

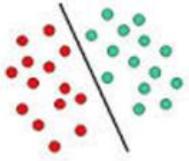
Neural AI, also known as **deep learning**, is a subset of machine learning that focuses on artificial neural networks inspired by the structure and function of the human brain. It involves training deep neural networks with multiple layers of interconnected nodes (neurons) to learn hierarchical representations of data. Neural AI excels in tasks such as image and speech recognition, natural language processing, and pattern recognition. It has achieved remarkable success in recent years due to advancements in computational power and the availability of large-scale datasets.

Early Ideas from Outside AI



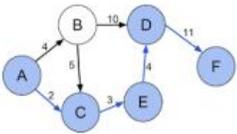
1801: linear regression (Gauss, Legendre)

Linear regression is a statistical modeling technique used to understand the relationship between a dependent variable and one or more independent variables. It assumes a linear relationship between the variables, meaning that the relationship can be represented by a straight line: $y = \beta_0 + \beta_1x + \epsilon$.



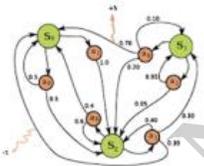
1936: linear classification (Fisher)

Linear classification is a machine learning technique used to classify data into different categories or classes based on a linear decision boundary. It assumes that the classes can be separated by a linear hyperplane in the feature space.



1956: Uniform cost search for shortest paths (Dijkstra)

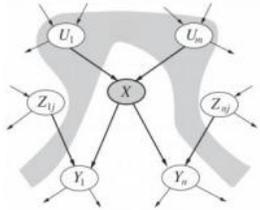
Uniform Cost Search (UCS) is a graph traversal algorithm used to find the shortest path in a weighted graph. It explores the graph by considering the cost of reaching each node from the starting node and selects the node with the lowest cost as the next node to visit.



1957: Markov decision processes (Bellman)

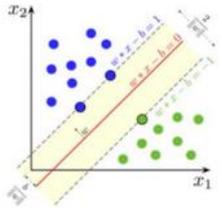
Markov Decision Processes (MDPs) are mathematical frameworks used to model decision-making in situations where outcomes are uncertain and influenced by both the current state and the actions taken.

Statistical Machine Learning



1985: Bayesian networks (Pearl)

Bayesian networks are a type of probabilistic graphical model that uses Bayesian inference for probability computations. Bayesian networks aim to model conditional dependence, and therefore causation, by representing conditional dependence by edges in a directed graph. Through these relationships, one can efficiently conduct inference on the random variables in the graph through the use of factors.



1995: Support vector machines (Cortes/Vapnik)

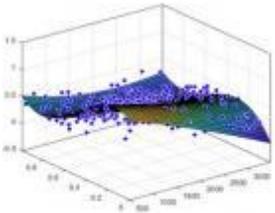
Support Vector Machine (SVM) is a powerful machine learning algorithm used for linear or nonlinear classification, regression, and even outlier detection tasks. SVMs can be used for a variety of tasks, such as text classification, image classification, spam detection, handwriting identification, gene expression analysis, face detection, and anomaly detection. SVMs are adaptable and efficient in a variety of applications because they can manage high-dimensional data and nonlinear relationships.

Three Intellectual Traditions



symbolic AI

Symbolic AI, also known as rule-based or logic-based AI, focuses on using logical rules and symbolic representations to perform reasoning and problem-solving tasks. It involves the use of knowledge representation, expert systems, and logical inference to manipulate symbols and make logical deductions. Symbolic AI emphasizes explicit knowledge representation and rule-based decision-making.



statistical AI

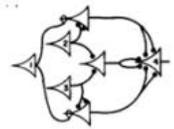
Statistical AI, also referred to as machine learning, involves the use of statistical models and algorithms to enable computers to learn from data and make predictions or decisions. It relies on analyzing patterns and relationships in large datasets to extract meaningful insights and generalize from examples. Statistical AI encompasses techniques such as regression, classification, clustering, and probabilistic models.



neural AI

Neural AI, also known as deep learning, is a subset of machine learning that focuses on **artificial neural networks** inspired by the structure and function of the human brain. It involves training deep neural networks with multiple layers of interconnected nodes (neurons) to learn hierarchical representations of data. Neural AI excels in tasks such as **image and speech recognition, natural language processing, and pattern recognition**. It has achieved remarkable success in recent years due to advancements in computational power and the availability of large-scale datasets.

Artificial Neural Networks



1943: artificial neural networks, relate neural circuitry and mathematical logic (McCulloch/Pitts)



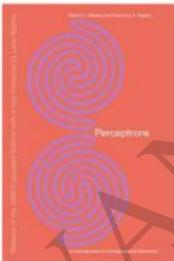
1949: "cells that fire together wire together" learning rule (Hebb)



1958: Perceptron algorithm for linear classifiers (Rosenblatt)



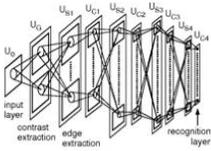
1959: ADALINE device for linear regression (Widrow/Hoff)



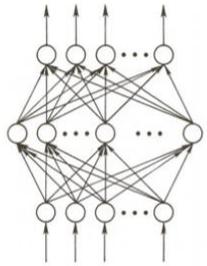
1969: Perceptrons book showed that linear models could not solve XOR, killed neural nets research (Minsky/Papert)

The XOR (Exclusive OR) function is a logical operation used in digital circuits and computing.

Revival of Connectionism



1980: Neocognitron, a.k.a. convolutional neural networks for images (Fukushima)



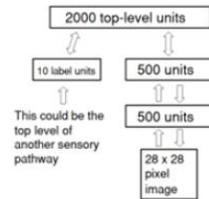
1986: popularization of backpropagation for training multi-layer networks (Rumelhardt, Hinton, Williams)



1989: applied convolutional neural networks to recognizing handwritten digits for USPS (LeCun)

USPS: United States Postal Service

Deep Learning



2006: unsupervised layerwise pre-training of deep networks (Hinton et al.)



2012: AlexNet obtains huge gains in object recognition; transformed computer vision community overnight



2016: AlphaGo uses deep reinforcement learning, defeat world champion Lee Sedol in Go

Generative AI

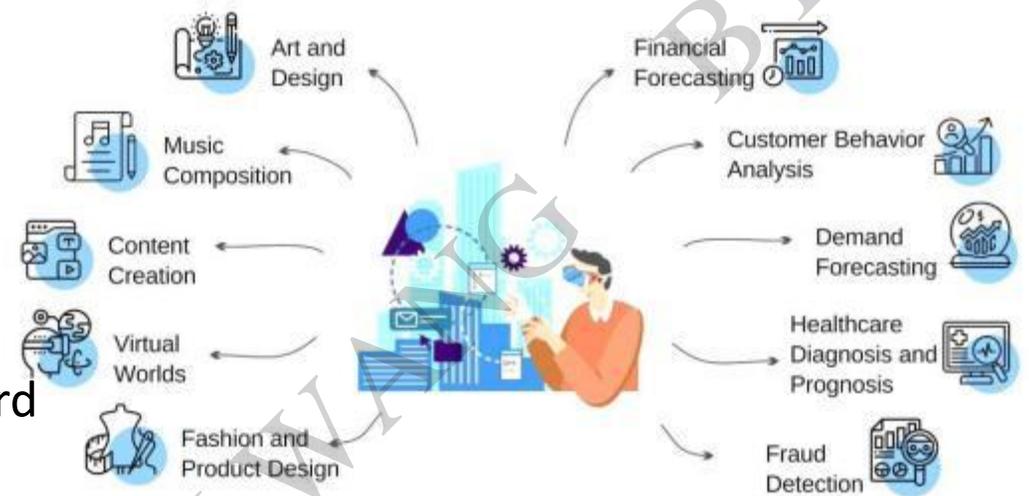
Generative Artificial Intelligence (AI) - AIGC (Artificial Intelligence Generated Content), also known as Generative AI or Generative Adversarial Networks (GANs), refers to the use of AI technologies such as GANs and large-scale pre-trained models to generate relevant content with appropriate generalization capabilities based on learning and recognition from existing data.

The core idea behind AIGC technology is to utilize AI algorithms to generate content with **a certain level of creativity and quality**. By training models on large amounts of data, AIGC can generate content that is related to specific inputs or guidance. For example, by inputting keywords, descriptions, or samples, AIGC can generate corresponding articles, images, audio, and more.

Generative AI

AIGC Chronology

- 2012: Preliminary Image Learning and Generation
- 2014: GAN Adversarial Models
- 2021: OPENAI Releases DALL·E
- 2022.04: Disco Diffusion
- 2022.06: Midjourney
- 2022.08: Midjourney Defeats Humans and Wins the Award
- 2022.08: Stability AI Open Sources Stable Diffusion
- 2022.11: MJV4
- 2022.11: ChatGPT
- 2022.11: Stable Diffusion 2.0
- 2023.02: ControlNet
- 2023.03: MJV5 GPT4
- Now : Multimodal, interactive, real-time generation



ChatGPT



Midjourney



Stable Diffusion

<https://poe.com/>

Artificial Intelligence

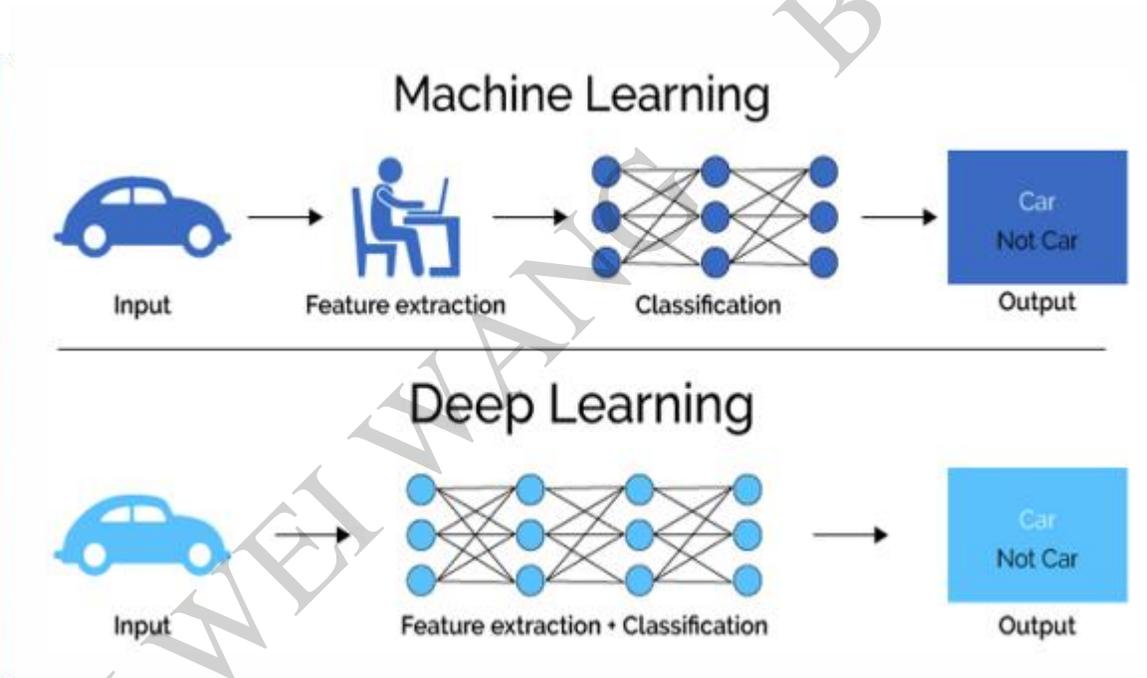
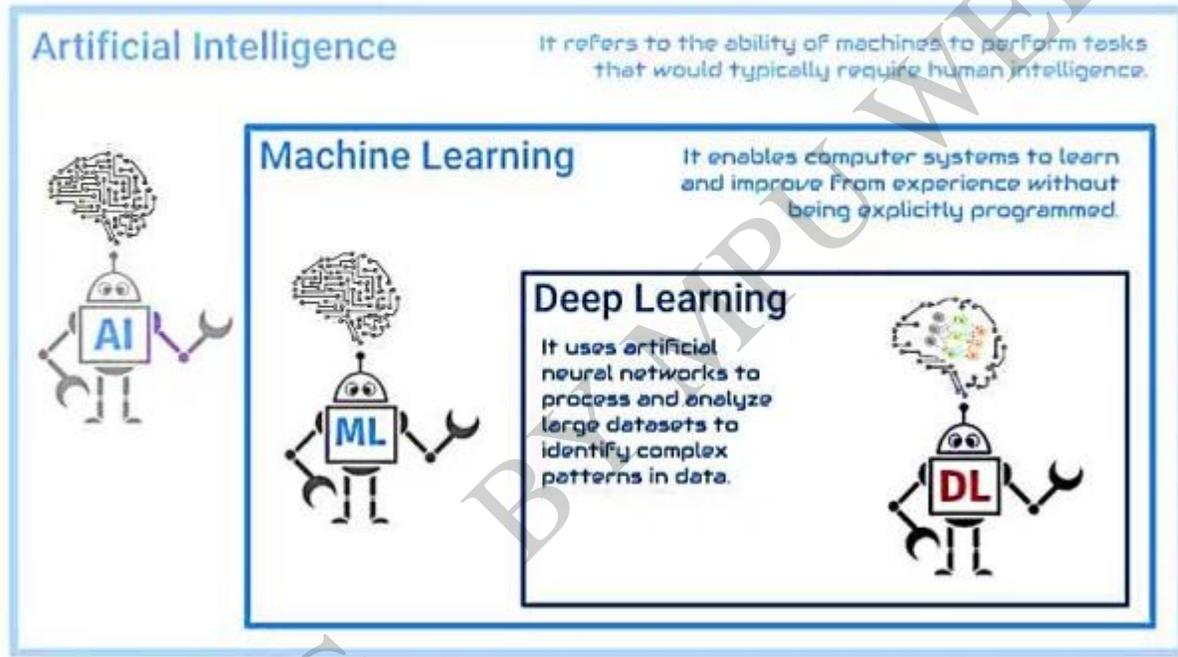
Computer Science (CS), Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) are interconnected fields within the broader domain of information technology. Here's an overview of their relationships:

Computer Science (CS): Computer Science is a discipline that encompasses the study of computation, algorithms, data structures, programming languages, software development, and various other aspects related to computing. It provides the **foundation and theoretical underpinnings** for many areas of technology, including AI, ML, and DL.

Artificial Intelligence (AI): Artificial Intelligence focuses on developing **intelligent systems** that can perform tasks that typically require human intelligence. AI encompasses a wide range of approaches, techniques, and algorithms aimed at enabling machines to perceive, reason, learn, and make decisions. AI involves the study of both symbolic and statistical methods to model and replicate intelligence in machines.

Machine Learning (ML): Machine Learning is a **subset of AI** that focuses on designing algorithms and statistical models that allow computer systems to learn from data and improve their performance on specific tasks without being explicitly programmed. ML algorithms enable computers to automatically analyze and extract patterns from large datasets and make predictions or decisions based on those patterns.

Deep Learning (DL): Deep Learning is a **subfield of Machine Learning** that specifically focuses on the design and training of deep neural networks. Deep neural networks are artificial neural networks with multiple layers of interconnected nodes (neurons). DL algorithms leverage these deep neural networks to learn hierarchical representations of data, enabling them to model complex patterns and relationships. DL has been particularly successful in tasks such as image and speech recognition, natural language processing, and other areas requiring high-level abstractions.



AI Versus CS/ML/DL

	AI	Machine learning	Deep learning
Optimal data volumes	Varying data volumes	Thousands of data points	Big data: millions of data points
Outputs	Anything from predictions to recommendations to decision-making	Numerical value, like a classification or score	Anything from numerical values to free-form elements, like free text and sound
How it works	Machines are programmed to mimic human activity with human-like accuracy	Uses various types of automated algorithms that learn to model functions and predict future actions from data	Uses neural networks that pass data through many processing layers to interpret data features and relationships
How it's managed	Algorithms require human oversight in order to function properly	Algorithms are directed by data analysts to examine specific variables in data sets	Algorithms are largely self-directed on data analysis once they're put into production

Special Thanks



Special thanks to Dr. **Zhang Jie** from the Faculty of Applied Science at the Macau Polytechnic University for preparing most of the course materials. My slides drew heavily on his content.



Thank you!

Innovating into the Future