CS440/ECE448: Artificial Intelligence
Lecture 2: History and Themes

Slides by Svetlana Lazebnik, 9/2016
Modified by Mark Hasegawa-Johnson, 1/2019 and Julia Hockenmaier 1/2019

Last time: What is AI?

Thinking Humanly?

Examples: embodied cognition, trying to reconstruct a brain cell-by-cell

Acting Humanly?

Examples: Turing test, Winograd schema

Thinking Rationally?

Example: Aristotle, especially the *Analytics*

Example: the logicist approach to AI, symbolic reasoning, fuzzy logic

Acting Rationally?

Example: John Stuart Mill, *Utilitarianism*

Example: rational agent theory, Economics

Today's lecture: History of AI

- Nice overview video
  https://www.youtube.com/watch?v=BFWt5Bxfcjo

- AI Successes

Image source

Examples of AI Successes

IBM Watson and "cognitive computing"

- [2010 NY Times article](), [trivia demo]()
- February 2011: [IBM Watson wins on Jeopardy]()
- Since then: [Watson Analytics](), [social services](), [personal shopping](), [health care]()

Self-driving cars

If the age of self-driving cars is upon us, what's keeping them off the ...
The Guardian - 9 hours ago
Sitting in the passenger seat of Google's self driving car is a less bizarre experience than sitting in the driving seat, but it's still unsettling. In the ...

How will self-driving cars affect your insurance?
The Conversation UK - 1 hour ago
And software bugs in self-driving cars could create a new reason manufacturers might have to shoulder the cost of crashes. Yet if drivers ...

If Uber and Lyft switch to self-driving cars, what about the drivers?
Digital Trends - Aug 21, 2016
Uber and Lyft have tens of thousands of drivers in the U.S. alone. The two ridesharing companies have each been in the news recently for their ...
The summer that changed everything for Uber: China, self-driving ...
In-Depth - VentureBeat - 3 hours ago
View all

Uber is about to start giving rides in self-driving cars
Los Angeles Times - Aug 18, 2016
The robot cars aren't coming. The robot cars are here. A fleet of Fords and Volvos, capable of driving themselves, is fully equipped and ready to ...
How Pittsburgh Birthed the Age of the Self-Driving Car
In-Depth - WIRED - Aug 19, 2016
View all

Here's how Ford's self-driving cars will work
Business Insider - Aug 20, 2016
On Tuesday, the automaker announced that it aims to roll out an autonomous taxi fleet in at least one city in 2021. Ford said its driverless cars will have level 4 ...

Legislators rush to keep up with self-driving cars
East Valley Tribune - 1 hour ago
There apparently are no laws which would prohibit manufacturers from marketing self-driving cars today to consumers. And nothing keeps ...

From Microsoft to Self-Driving Cars, Invention Springs From Data
Adweek - 16 hours ago
The product data arms race is also at play in the nascent but exploding space of self-driving cars. While design will certainly be important, ...

Ford Promises Fleets of Driverless Cars Within Five Years
New York Times - Aug 16, 2016
In the race to develop driverless cars, several automakers and technology companies are already testing vehicles that pilot themselves on ...
Want a ride? Ford hopes to race ahead of self-driving pack sans ...
Opinion - Chicago Tribune - Aug 17, 2016
View all

Self-driving cars safe, say firms developing such vehicles here
The Straits Times - Aug 20, 2016
Despite recent accidents involving driverless cars in the United States, two companies involved in developing such vehicles here said they are ...

Uber Buys Self-Driving Truck Biz
Sci-Tech Today - Aug 21, 2016
Uber Buys Self-Driving Truck Biz and Tests Autonomous Cars .... Ford's approach to the autonomous car breaks from many other companies, ...

Google News snapshot as of August 22, 2016

Speech and natural language

Skype Translator

Break down the language barrier with your friends, family and colleagues.

Our online translator can help you communicate in 7 languages for voice calls, and in more than 50 languages while instant messaging.

Skype Translator uses machine learning. So the more you use it, the better it gets. Thanks for being patient as the technology graduates from Preview mode.

Hallo, hola, olá to the new, more powerful Google Translate app

Posted: Wednesday, January 14, 2015

- Instant translation with Word Lens
- Have a conversation with Google Translate

Computer Vision (Face Recognition)

Games

- 1997: IBM's Deep Blue defeats the reigning world chess champion Garry Kasparov
  - **1996: Kasparov Beats Deep Blue**
    "I could feel – I could smell – a new kind of intelligence across the table."
  - **1997: Deep Blue Beats Kasparov**
    "Deep Blue hasn't proven anything."
- 2007: Checkers is solved
  - Though checkers programs had been beating the best human players for at least a decade before then
- 2014: Heads-up limit Texas Hold-em poker is solved
  - First game of imperfect information
- 2016: AlphaGo computer beats Go grandmaster Lee Sedol 4-1

Mathematics

- In 1996, a computer program written by researchers at Argonne National Laboratory proved a mathematical conjecture unsolved for decades
  - NY Times story: *"[The proof] would have been called creative if a human had thought of it"*
- Mathematical software:

$$\partial_r^2 u = -\left[E' - \frac{l(l+1)}{r^2} - r^2\right] u(r)$$

$$e^{-2s}\left(\partial_s^2 - \partial_s\right) u(s) = -\left[E' - l(l+1)e^{-2s} - e^{2s}\right] u(s)$$

$$e^{-2s}\left[e^{\frac{1}{2}s}\left(e^{-\frac{1}{2}s}u(s)\right)'' - \frac{1}{4}u\right] = -\left[E' - l(l+1)e^{-2s} - e^{2s}\right] u(s)$$

$$e^{-2s}\left[e^{\frac{1}{2}s}\left(e^{-\frac{1}{2}s}u(s)\right)''\right] = -\left[E' - \left(l + \frac{1}{2}\right)^2 e^{-2s} - e^{2s}\right] u(s)$$

$$v'' = -e^{2s}\left[E' - \left(l + \frac{1}{2}\right)^2 e^{-2s} - e^{2s}\right] v$$

Logistics, scheduling, planning

- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people

- NASA's Remote Agent software operated the Deep Space 1 spacecraft during two experiments in May 1999

- In 2004, NASA introduced the MAPGEN system to plan the daily operations for the Mars Exploration Rovers

Robotics

- Mars rovers
- Autonomous vehicles
  - [DARPA Grand Challenge](#)
  - Self-driving cars
- [Autonomous helicopters](#)
- Robot soccer
  - [RoboCup](#)
- Personal robotics
  - [Humanoid robots](#)
  - [Robotic pets](#)
  - Personal assistants?

History of AI

Origins of AI: Early excitement

1939   Hodgin & Huxley measure **action potentials** of squid giant axon

1940s First model of a **neuron** (W. S. McCulloch & W. Pitts)

       Hebbian learning rule

       Cybernetics

       Turing Test

1950s **Perceptrons** (F. Rosenblatt)

       Computer **chess** and **checkers** (C. Shannon, A. Samuel)

       Machine translation (Georgetown-IBM experiment)

       Theorem provers (A. Newell & H. Simon, H. Gelernter & N. Rochester)

1956   Dartmouth meeting:
the term "**Artificial Intelligence**" is adopted

Squid giant axon

From Wikipedia, the free encyclopedia

The **squid giant axon** is the very large (up to 1 mm in diameter; typically around 0.5 mm) axon that controls part of the water jet propulsion system in squid. It was first described by L. W. Williams[1] in 1909,[2][*page needed*] but this discovery was forgotten until English zoologist and neurophysiologist J. Z. Young demonstrated the axon's function in the 1930s while working in the Stazione Zoologica in Naples, the Marine Biological Association in Plymouth and the Marine Biological Laboratory in Woods Hole.[3][4] Squids use this system primarily for making brief but very fast movements through the water.

Between the tentacles of a squid is a siphon through which water can be rapidly expelled by the fast contractions of the body wall muscles of the animal. This contraction is initiated by action potentials in the giant axon. Action potentials travel faster in a larger axon than a smaller one,[5][*page needed*] and squid have evolved the giant axon to improve the speed of their escape response. The increased diameter of the squid axon decreases the internal resistance of the axon, as resistance is inversely proportional to the cross sectional area of the object. This increases the space constant ($\lambda = \sqrt{r_m/r_i}$), leading to faster local depolarization and a faster action potential ($E = E_o e^{-x/\lambda}$).[6][*page needed*]

In their Nobel Prize-winning work uncovering ionic mechanism of action potentials, Alan Hodgkin and Andrew Huxley performed experiments on the squid giant axon. The prize was shared with John Eccles. The large diameter of the axon provided a great experimental advantage for Hodgkin and Huxley as it allowed them to insert voltage clamp electrodes inside the lumen of the axon.

While the squid axon is very large in diameter it is unmyelinated which decreases the conduction velocity substantially. The conduction velocity of a typical 0.5 mm squid axon is about 25 m/s. During a typical action potential in the cuttlefish *Sepia* giant axon, an influx of 3.7 pmol/cm^2 (picomoles per centimeter^2) of sodium is offset by a subsequent efflux of 4.3 pmol/cm^2 of potassium.[7]

See also [edit]

- Lateral giant neuron

1943: McCulloch-Pitts Neuron Model

$x_0 = +1$

$w_{0k} = b_k$

x_1

w_{1k}

x_2

w_{2k}

w_{3k}

$v_k = \sum_i w_{ik} x_i$

$y_k = \varphi(v_k)$

x_3

x_m

w_{mk}

An Electronic Model of Neural Learning: the Perceptron, 1957, Rosenblatt

Attribution: Cornell University Library

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 7 (UPI) —The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's $2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.,

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of $100,000.

Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human beings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

Without Human Controls

The Navy said the perceptron would be the first non-living mechanism "capable of receiving, recognizing and identifying its surroundings without any human training or control."

The "brain" is designed to remember images and information it has perceived itself. Ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

In today's demonstration, the "704" was fed two cards, one with squares marked on the left side and the other with squares on the right side.

Learns by Doing

In the first fifty trials, the machine made no distinction between them. It then started registering a "Q" for the left squares and "O" for the right squares.

Dr. Rosenblatt said he could explain why the machine learned only in highly technical terms. But he said the computer had undergone a "self-induced change in the wiring diagram."

The first Perceptron will have about 1,000 electronic "association cells" receiving electrical impulses from an eye-like scanning device with 400 photo-cells. The human brain has 10,000,000,000 responsive cells, including 100,000,000 connections with the eyes.

Herbert Simon, 1957

- *"It is not my aim to surprise or shock you – but ... there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – in a visible future – the range of problems they can handle will be coextensive with the range to which human mind has been applied.* **More precisely: within 10 years a computer would be chess champion, and an important new mathematical theorem would be proved by a computer.***"*

- Prediction came true – but 40 years later instead of 10

Harder than originally thought

- 1966: Eliza chatbot (Weizenbaum)
  - **" … mother …" → "Tell me more about your family"**
  - **"I wanted to adopt a puppy, but it's too young to be separated from its mother."**
- 1954: Georgetown-IBM experiment
  - Completely automatic translation of more than sixty Russian sentences into English
  - Only six grammar rules, 250 vocabulary words, restricted to organic chemistry
  - Promised that machine translation would be solved in three to five years (press release)
  - Automatic Language Processing Advisory Committee (ALPAC) report (1966): machine translation has failed
    - **"The spirit is willing but the flesh is weak." → "The vodka is strong but the meat is rotten."**
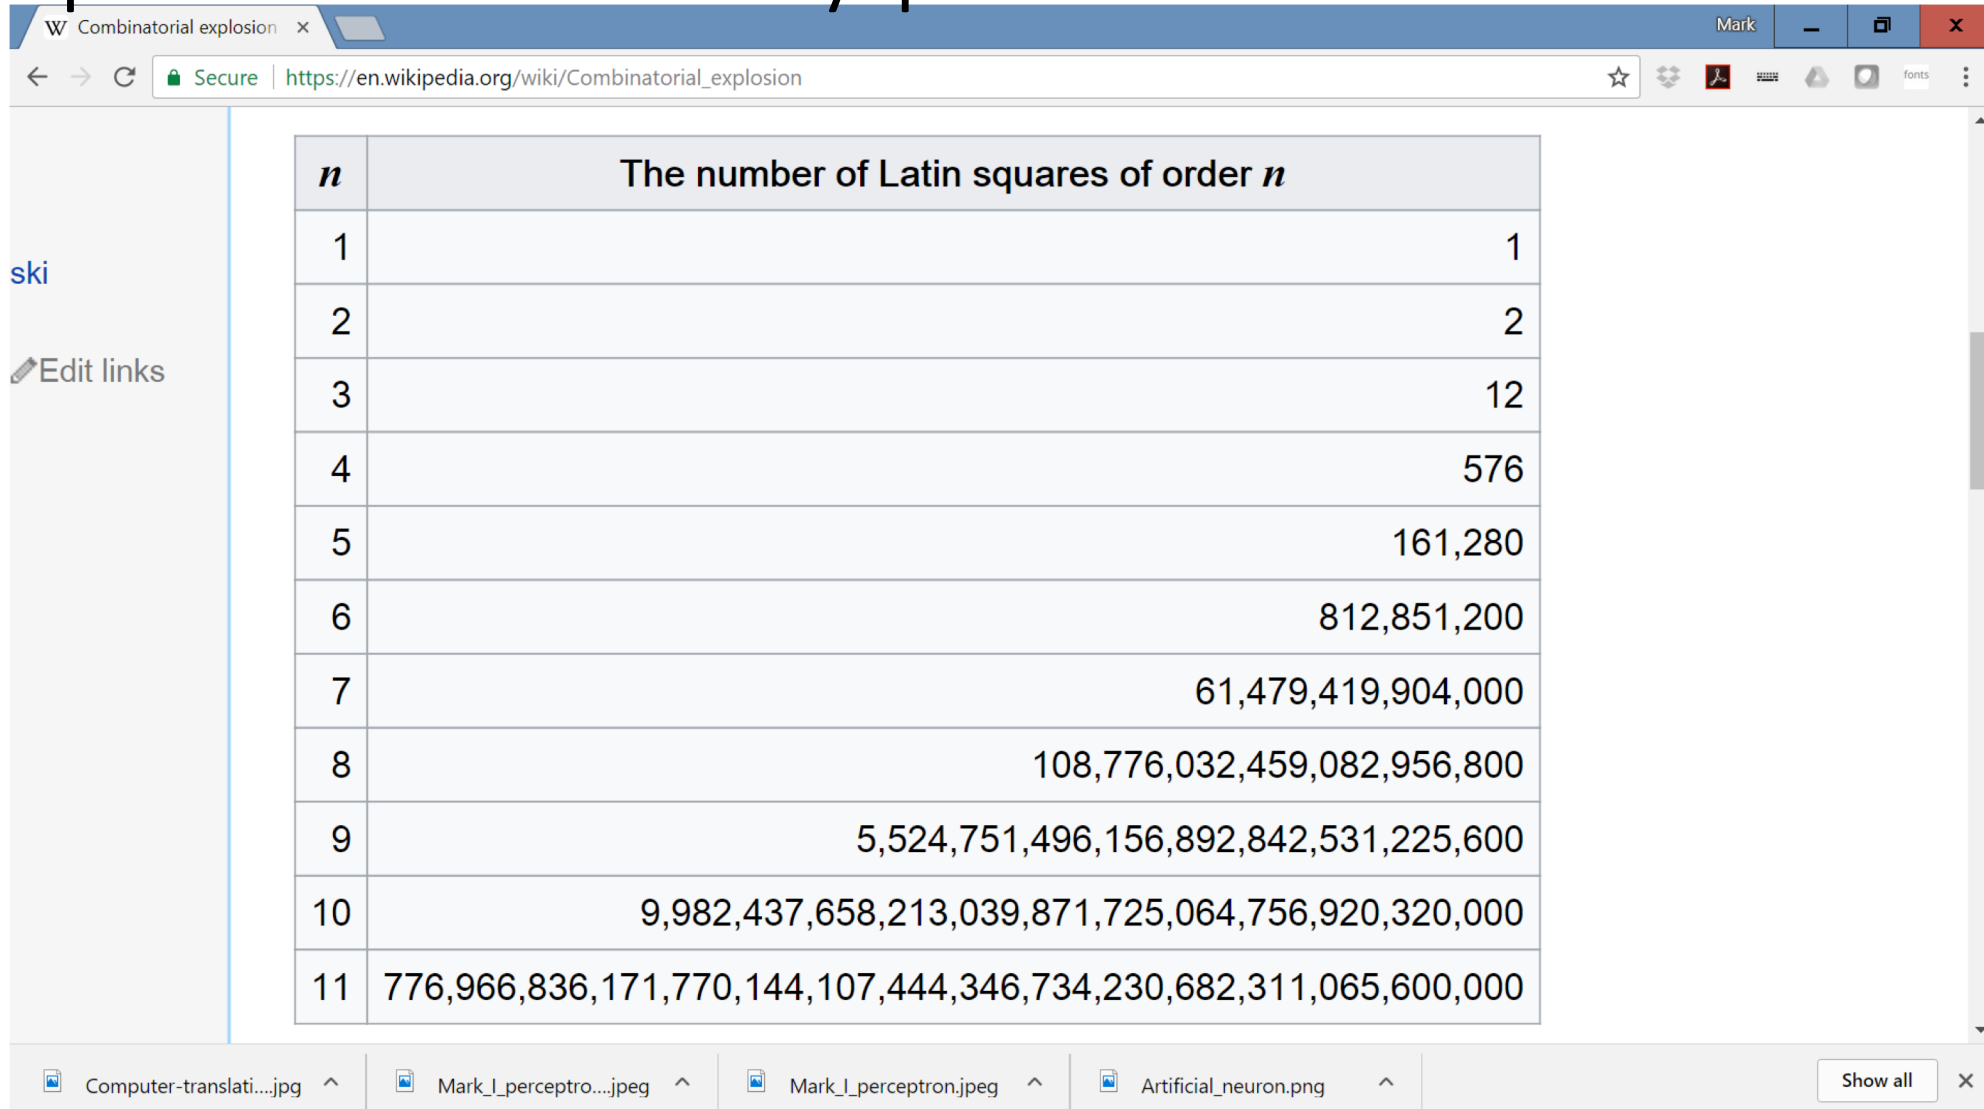
The ALPAC Report of 1966

"They concluded, in a famous 1966 report, that machine translation was more expensive, less accurate and slower than human translation."

Photo: Eldon Lyttle, https://commons.wikimedia. org/wiki/File:Computer-translation_Briefing_for_Ger ald_Ford.jpg

The Lighthill Report, 1973: combinatorial explosion is the key problem

n	The number of Latin squares of order n
1	1
2	2
3	12
4	576
5	161,280
6	812,851,200
7	61,479,419,904,000
8	108,776,032,459,082,956,800
9	5,524,751,496,156,892,842,531,225,600
10	9,982,437,658,213,039,871,725,064,756,920,320,000
11	776,966,836,171,770,144,107,444,346,734,230,682,311,065,600,000

Blocks world (1960s – 1970s)

(a) Original picture.

(b) Differentiated picture.

(c) Line drawing.

(d) Rotated view.
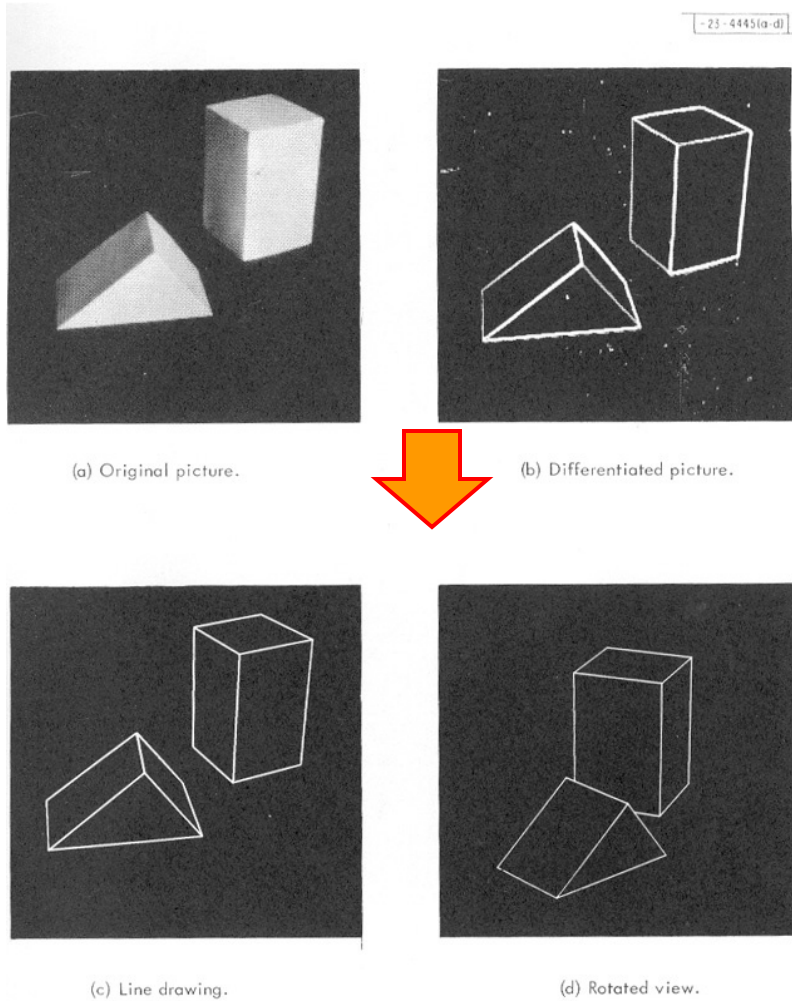
Larry Roberts, MIT, 1963

???

WIKIPEDIA
The Free Encyclopedia

AI winter

From Wikipedia, the free encyclopedia

> ⚠️ **This article has multiple issues.** Please help improve it or discuss these issues on the talk page. (Learn how and when to remove these template messages)   [hide]
>
> - The **lead section of this article** may need to be rewritten. (January 2016)
> - This article's **tone or style** may not reflect the encyclopedic tone used on Wikipedia. (July 2016)

In the history of artificial intelligence, an **AI winter** is a period of reduced funding and interest in artificial intelligence research.[1] The term was coined by analogy to the idea of a nuclear winter.[2] The field has experienced several hype cycles, followed by disappointment and criticism, followed by funding cuts, followed by renewed interest years or decades later.

The term first appeared in 1984 as the topic of a public debate at the annual meeting of AAAI (then called the "American Association of Artificial Intelligence"). It is a chain reaction that begins with pessimism in the AI community, followed by pessimism in the press, followed by a severe cutback in funding, followed by the end of serious research.[2] At the meeting, Roger Schank and Marvin Minsky—two leading AI researchers who had survived the "winter" of the 1970s—warned the business community that enthusiasm for AI had spiraled out of control in the '80s and that disappointment would certainly follow. Three years later, the billion-dollar AI industry began to collapse.[2]

Hypes are common in many emerging technologies, such as the railway mania or the dot-com bubble. The AI winter is primarily a collapse in the *perception* of AI by government bureaucrats and venture capitalists. Despite the rise and fall of AI's reputation, it has continued to develop new and successful technologies. AI researcher Rodney Brooks would complain in 2002 that "there's this stupid myth out there that AI has failed, but AI is around you every second of the day."[3] In 2005, Ray Kurzweil agreed: "Many observers still think that the AI winter was the end of the story and that nothing since has come of the AI field. Yet today many thousands of AI applications are deeply embedded in the infrastructure of every industry."[4]

Enthusiasm and optimism about AI has gradually increased since its low point in 1990, and by the 2010s artificial intelligence (and especially the sub-field of machine learning) became widely used, well-funded and many in the technology predict that it will soon

Artificial intelligence

Major goals

Knowledge reasoning

Planning

Machine learning

Natural language processing

Computer vision

Robotics

Artificial general intelligence

Approaches

Symbolic

Deep learning

Recurrent neural networks

Bayesian networks

Evolutionary algorithms

Philosophy

History of AI to the present day

1975-1985: **Expert systems boom**

1985-1995: **Expert system bust; the second "AI winter"**

Expert system, brief comic explanation:
https://www.youtube.com/watch?v=sg6hLmuyQ54

1995-2009: **Probabilistic reasoning**/ Bayesian logic boom

2009-now: **Deep learning boom**

Neural nets solve expert system problems:
https://www.youtube.com/watch?v=n-YbJi4EPxc

History of AI on Wikipedia

Building Smarter Machines: NY Times Timeline

What accounts for recent successes in AI?

- Faster computers
  - The IBM 704 vacuum tube machine that played chess in 1958 could do about **50,000 calculations per second**
  - Deep Blue could do **50 billion calculations per second** – a million times faster!
- Dominance of statistical approaches, machine learning
- Big data
- Crowdsourcing (to cheaply obtain large amounts of labeled data)

Machine Intelligence LANDSCAPE

CORE TECHNOLOGIES

ARTIFICIAL INTELLIGENCE
IBM WATSON · MetaMind · Numenta · ai-one · Cycorp · Microsoft Research · nara · Reactor · SI SCALED INFERENCE

DEEP LEARNING
vicarious · Vision Factory · facebook · Bai IDL · Google · ersatz · SKYMIND · SignalSense

MACHINE LEARNING
rapidminer · context relevant · 0xdata H2O · DATARPM · Liftigniter · SPARKBEYOND · Azure ML · yhat · wise.io · Sense · GraphLab · Alpine · Nutonian

NLP PLATFORMS
cortical.io · idibon · LUMINOSO · wit.ai · Maluuba

PREDICTIVE APIS
AlchemyAPI · MINDOPS · Google · big ml · indico · ALGORITHMIA · Expect Labs · PredictionIO

IMAGE RECOGNITION
clarifai · MADBITS · DNNresearch · DEXTRO · ViSENZE · lookflow

SPEECH RECOGNITION
GRIDSPACE · popUP archive · NUANCE

RETHINKING ENTERPRISE

SALES
Preact · AVISO · RelateIQ · NGDATA · CLARABRIDGE · FRAMED · infer · ATTENSITY · causata

SECURITY / AUTHENTICATION
CROSSMATCH · conjur · EYEVERIFY · BITSIGHT · CYLANCE · AREA 1 · bionym

FRAUD DETECTION
sift science · SOCURE · ThreatMetrix · feedzai · Brighterion · VERAFIN

HR / RECRUITING
TalentBin · entelo · predikt · Connectifier · gild · hiQ · CONCEPTNODE

MARKETING
brightfunnel · bloomreach · CommandIQ · AIRPR · RADIUS · Telkpart · people pattern · Freshplum

PERSONAL ASSISTANT
Siri · x. · Google now · Cortana · cleversense · tempo · Robinlabs · KASISTO · fuse machines · VIV · CLARA LABS

INTELLIGENCE TOOLS
ADATAO · Palantir · Quid · Digital Reasoning · FirstRain

RETHINKING INDUSTRIES

ADTECH
METAMARKETS · dstillery · rocketfuel · YieldMo · ADBRAIN

AGRICULTURE
BLUE RIVER · TerrAvion · ceresimaging · HONEYCOMB · THE CLIMATE CORPORATION · tule · mavrx

EDUCATION
declara · coursera · KNEWTON · kidaptive

FINANCE
Bloomberg · FinGenius · alphasense · KENSHO · Dataminr · minettabrook · BINATIX

LEGAL
Lex Machina · brightleaf · COUNSELYTICS · RAVEL · JUDICATA · Brevia · DiligenceEngine

MANUFACTURING
SIGHT MACHINE · MICROSCAN · BOULDER IMAGING · IVISYS

MEDICAL
Parzival · transcriptic · Genescient · ZEPHYR HEALTH · grand round table · bina · TUTE GENOMICS

OIL AND GAS
kaggle · AYASDI · TACHYUS · biota · Flutura

MEDIA / CONTENT
Outbrain · newsle · ARRIA · SAILTHRU · wavii · Owlin · NarrativeScience · YSEOP · Summly · Prismatic · ai AUTOMATED INSIGHTS

CONSUMER FINANCE
Affirm · inVenture · zestfinance · BILLGUARD · LendUp · LendingClub · Kabbage

PHILANTHROPIES
DataKind · thorn · THE DATA GUILD

AUTOMOTIVE
Google · Continental · TESLA · MOBILEYE · CRUISE

DIAGNOSTICS
enlitic · 3SCAN · lumiata · ENTOPSIS

RETAIL
BAY SENSORS · PRISM SKYLABS · celect · euclid

RETHINKING HUMANS / HCI

AUGMENTED REALITY
wearable intelligence · APX · blippar · META · layar

GESTURAL COMPUTING
THALMICLABS · omek · SoftBank · FLUTTER · LEAP MOTION · eyeSight · 3Gear SYSTEMS · GestureTek · nod

ROBOTICS
intel · LIQUID ROBOTICS · iRobot · Boston Dynamics · jibo · evolution robotics · ANKI

EMOTIONAL RECOGNITION
affectiva · BEYOND VERBAL · emotient · cogito

SUPPORTING TECHNOLOGIES

HARDWARE
NVIDIA · XILINX · QUALCOMM · NERVANA SYSTEMS · TERADEEP · Artificial Learning · rigetti

DATA PREP
TRIFACTA · Paxata · tamr · Alation

DATA COLLECTION
diffbot · kimono · CrowdFlower · Connotate · WorkFusion · import io

Technology

Last updated: January 4, 2015 4:46 pm

Investor rush to artificial intelligence is real deal

Richard Waters in San Francisco

©Bloomberg

IBM's Dario Gil unveils its question-answering system Watson. AI can help predict customer behaviour, says venture capitalist Matt McIlwain

Silicon Valley loves a new fad. To judge by the spate of fundraising by start-ups in recent weeks, it has found one in an idea that is more than half a century old: artificial intelligence.

Historical themes

- Boom and bust cycles
  - Periods of (unjustified) optimism followed by periods of disillusionment and reduced funding

- Silver bulletism ([Levesque, 2013](#)):
  - *"The tendency to believe in a silver bullet for AI, coupled with the belief that previous beliefs about silver bullets were hopelessly naïve"*

- Image problems
  - ["AI effect"](#)/Moravec's paradox: As soon as a machine gets good at performing some task, the task is no longer considered to require much intelligence
  - AI as a threat to safety?
  - AI as a threat to jobs?

AI Effect/Moravec's paradox

- [Moravec's paradox](#)

  "It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility"
  [Hans Moravec, 1988]

- Why is this?
  - Early AI researchers concentrated on the tasks that they themselves found the most challenging, abilities of animals and two-year-olds were overlooked
  - We are least conscious of what our brain does best
  - Sensorimotor skills took millions of years to evolve, whereas abstract thinking is a relatively recent development

AI weapons are a threat to humanity, warn Hawking, Musk and Wozniak

by Jason Murdock   28 Jul 2015

The rush to develop autonomous weapons will cause a global arms race, according to an open letter signed by over 1,000 artificial intelligence (AI) researchers, academics and computer scientists.

The letter has been signed by high-profile figures including physicist Stephen Hawking, Tesla chief executive Elon Musk and Apple co-founder Steve Wosniak, and argues AI has reached a point where deployment of robotic weapons is feasible within years.

Autonomous weapons are described in the letter as those that "select and engage targets without human intervention".

This includes, for example, armed quadcopters searching and eliminating targets that meet pre-defined criteria, but not remote controlled missiles or piloted drones that still have human involvement.

The letter was presented at this year's International Joint Conferences on AI in Buenos Aires, and argues that there are advantages to replacing human soldiers with machines but that doing so would "lower the threshold" for warfare.

News | Sport

Stephen Hawking warns artificial intelligence could end mankind

💬 COMMENTS (1027)

By Rory Cellan-Jones
Technology correspondent

Safari Power Saver   Click to Start Flash Plug-in

▶

Stephen Hawking: "Humans, who are limited by slow biological evolution, couldn't compete and would be superseded"

Prof Stephen Hawking, one of Britain's pre-eminent scientists, has said that efforts to create thinking machines pose a threat to our very existence.

He told the BBC:"The development of full artificial intelligence could spell the end of the human race."

His warning came in response to a question about a revamp of the technology he uses to communicate, which involves a basic form of AI.

WIRED

Elon Musk Donates $10M to Keep AI From Turning Evil

BY DAVEY ALBA 01.15.15 | 2:24 PM | PERMALINK

Elon Musk, Tesla Chairman, Product Architect and CEO, speaks at the Automotive News World Congress in Detroit, Tuesday, Jan. 13, 2015. 📷 Paul Sancya/AP

Elon Musk is worried that artificial intelligence research could go wrong—very wrong.

This may seem unexpected, coming as it does from the architect of the conceptual high-speed transportation system Hyperloop, and the CEO of such moonshot-seeking companies as SpaceX and Tesla Motors. But Musk is so committed to this point of view that on Thursday, he announced a donation of $10 million to the Future of Life Institute (FLI), which will run a global research program aimed at keeping AI "beneficial to humanity." In other words, Musk wants to keep AI from running loose and growing into something that's a real danger to humans, a fear he's expressed before.

http://www.bbc.com/news/technology-30290540 http://www.wired.com/2015/01/elon-musk-ai-safety/

# Will robots take our jobs? Experts can't decide

theguardian

A new report from Pew Research brings together almost 2,000 experts to comprehensively assess the effect of robots on the workplace

Alex Hern

🐦 @alexhern

Wednesday 6 August 2014 10.00 EDT

http://www.theguardian.com/technology/2014/aug/06/robots-jobs-artificial-intelligence-pew

In this class

- Part 1: sequential reasoning (MP1, MP2)

- Part 2: pattern recognition and learning (MP3, MP4)

Philosophy of this class

- Goal: use machines to solve hard problems that are traditionally thought to require human intelligence
- We will try to follow a sound scientific/engineering methodology
  - Consider relatively limited application domains
  - Use well-defined input/output specifications
  - Define operational criteria amenable to objective validation
  - Zero in on essential problem features
  - Focus on principles and basic building blocks