



# **AI Intro**

**- a very, very short intro -**

**December 2019**

Sila. 12/12 2019.



## AI is everywhere...

AI is:

- being deployed in health care and warfare.
- helping people make music and books.
- it's scrutinizing your resume, judging your creditworthiness, and tweaking the photos you take on your phone.

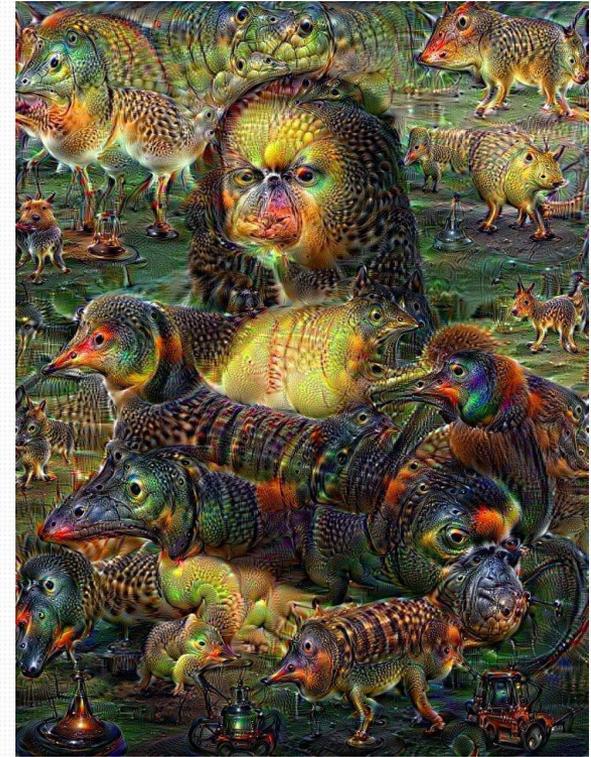
In short, it's making decisions that affect your life whether you like it or not".

# AI Art



<https://deepdreamgenerator.com/>

A journey through all the layers of an artificial neural network.



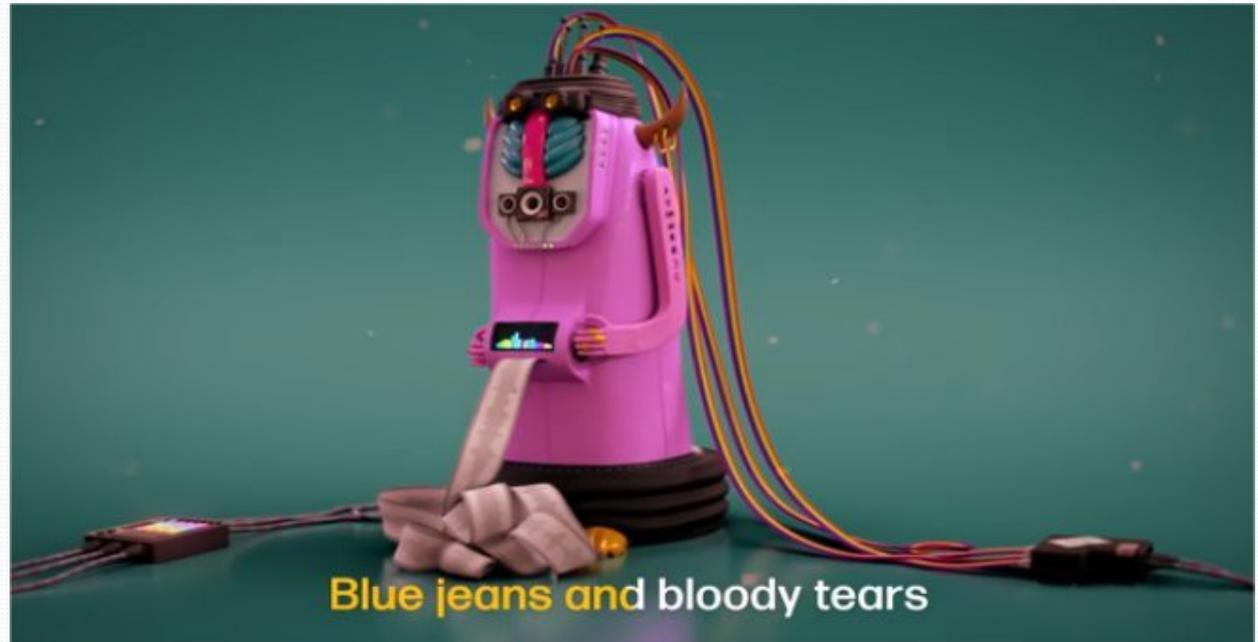
<https://vimeo.com/132700334>

<https://ai.googleblog.com/2015/07/deepdream-code-example-for-visualizing.html>

DeepDream - a code example for visualizing Neural Networks



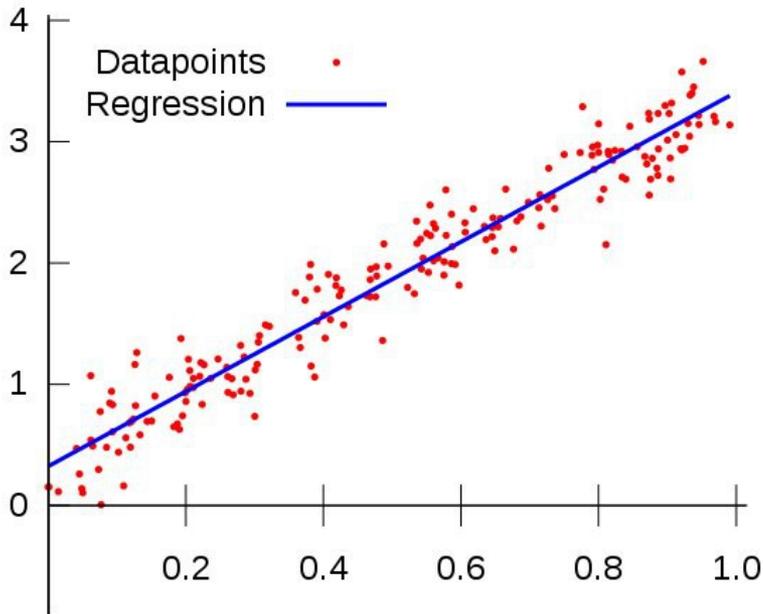
A Eurovision song created by Artificial Intelligence:  
"Blue Jeans and Bloody Tears" (2019):



[https://youtu.be/4MKAf6YX\\_7M](https://youtu.be/4MKAf6YX_7M)

Called "I AM AI". It's the first LP to be entirely composed and produced with AI.

<https://www.theverge.com/2017/8/27/16197196/taryn-southern-album-artificial-intelligence-interview>



### Gradient descent algorithm

repeat until convergence {

$$\theta_0 := \theta_0 - \alpha \frac{1}{m} \sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)})$$

$$\theta_1 := \theta_1 - \alpha \frac{1}{m} \sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)}) \cdot x^{(i)}$$

}

$$\frac{\partial}{\partial \theta_0} J(\theta_0, \theta_1)$$

$$\frac{\partial}{\partial \theta_1} J(\theta_0, \theta_1)$$

update  $\theta_0$  and  $\theta_1$  simultaneously

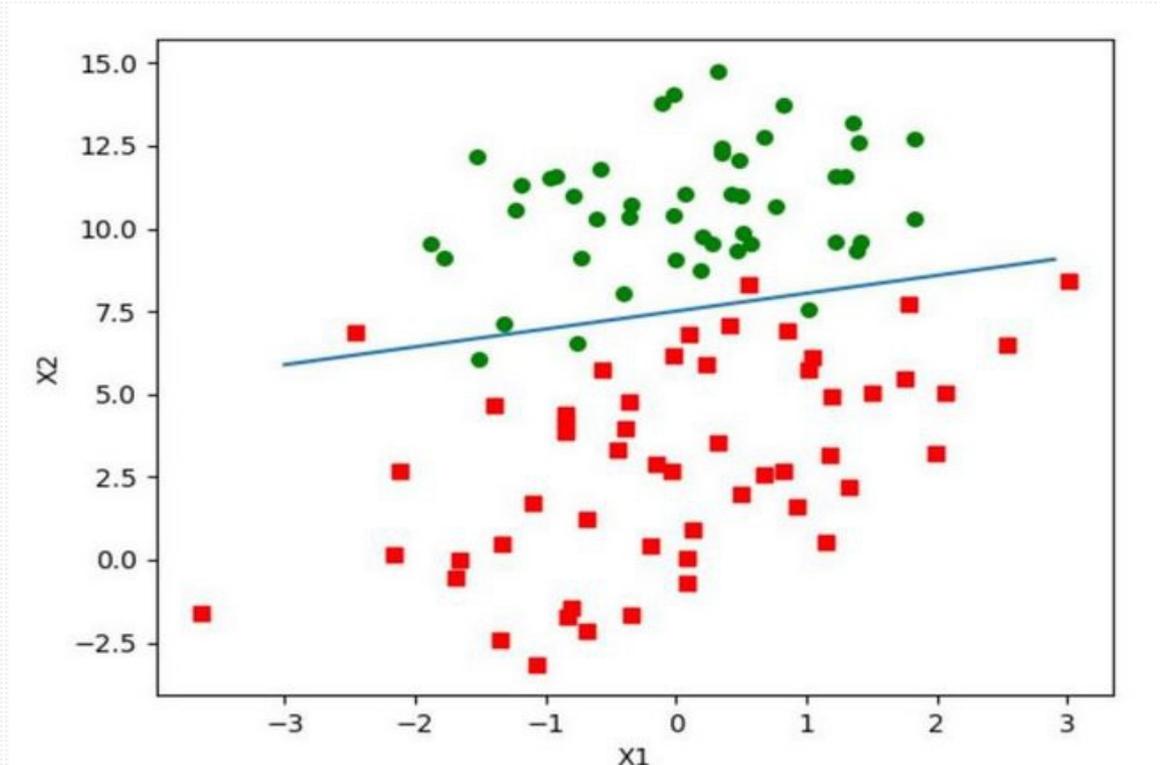
$$h_{\theta}(x) = \theta_0 + \theta_1 x$$

But let's start with something simple ...



In statistics, linear regression is a linear approach to modeling the relationship between a scalar response (or dependent variable) and one or more explanatory variables (or independent variables).

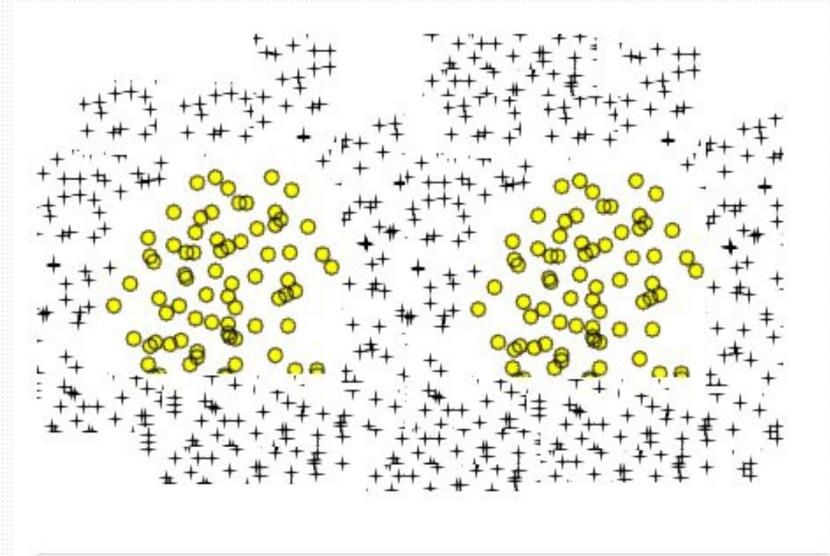
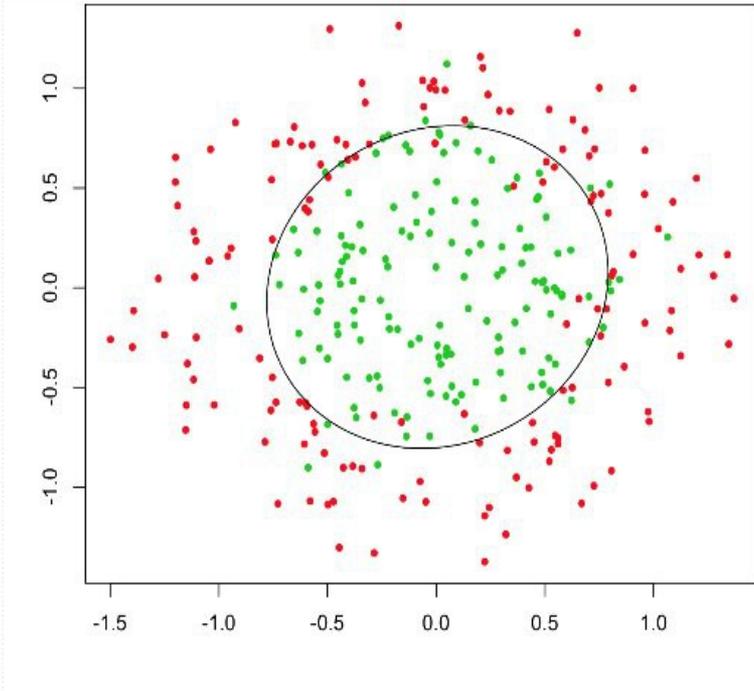
E.g. Estimate price of house based on size of house ...?



Or - Logistic Regression for classification problems.



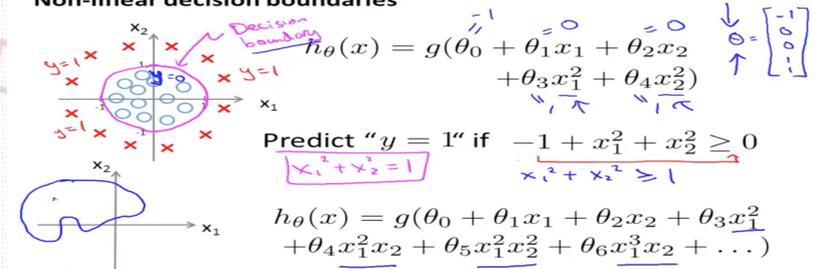
## More advanced classifications tasks:



ML : Use polynomial extensions to logistic regression in case of arbitrary classification



### Non-linear decision boundaries





Goals of AI:

- ***Technological:***

”Using computers to get useful things done.  
(Sometimes employing methods unlike those used by minds)”.

- ***Scientific:***

”Using AI concepts and models to help answer questions about human beings and other living things”.

See page 2 in

”Artificial Intelligence. A Very Short Introduction”

Margaret A. Boden.

## Arthur Samuel working on IBM 701



On February 24, 1956, Arthur Samuel's Checkers program, which was developed for play on the IBM 701, was demonstrated to the public on television.

<https://www.ibm.com/ibm/history/ibm100/us/en/icons/ibm700series/impacts/>

Samuels believed teaching computers to play games was very fruitful for developing tactics appropriate to general problems, and he chose checkers as it is relatively simple though has a depth of strategy.

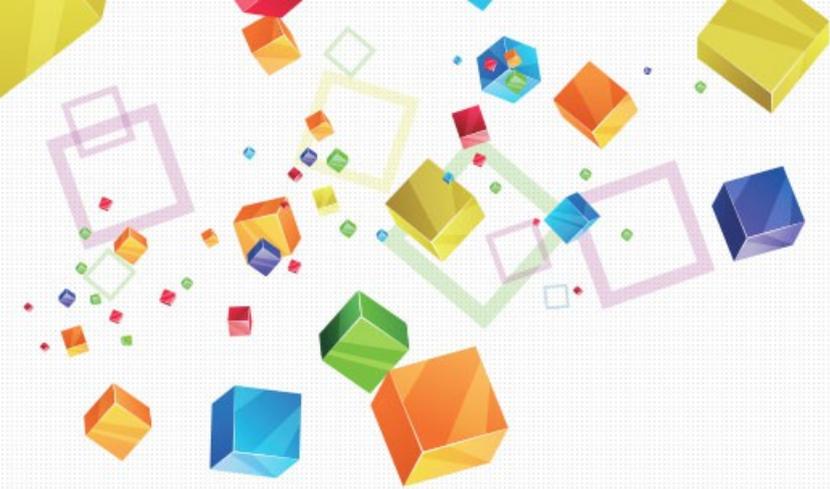
The main driver of the machine was a search tree of the board positions reachable from the current state. Since he had only a very limited amount of available computer memory, Samuel implemented what is now called alpha-beta pruning.

[https://en.wikipedia.org/wiki/Arthur\\_Samuel](https://en.wikipedia.org/wiki/Arthur_Samuel)



The computer is about to win this game and type out: 'Sorry, you lose.' John Pfeiffer, *The Thinking Machine*, 1962.





## Basic structuring of a problem:

### define the goals:

e.g., take my son to nursery school

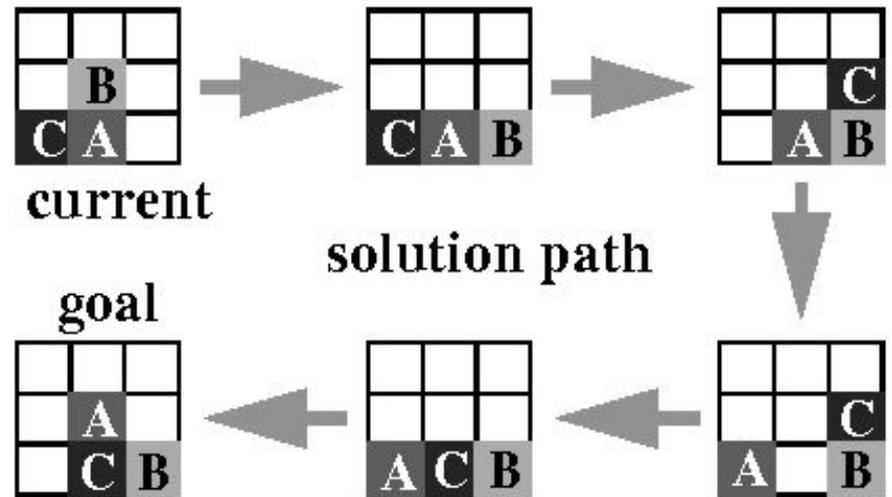
### define preconditions for the goals

e.g., the precondition for dropping my son off at nursery school is that my son is at home and the car works

### define the means (“operators”) for turning one set of conditions into another

e.g., to ensure the car repair shop has my money, we can define an operator “give shop money” that changes the world so that “shop has money” is true and “I have money” is false.

# GENERAL PROBLEM SOLVER



The solution to the problem in a global sense is a sequence of configurations between the current configuration and the goal state.

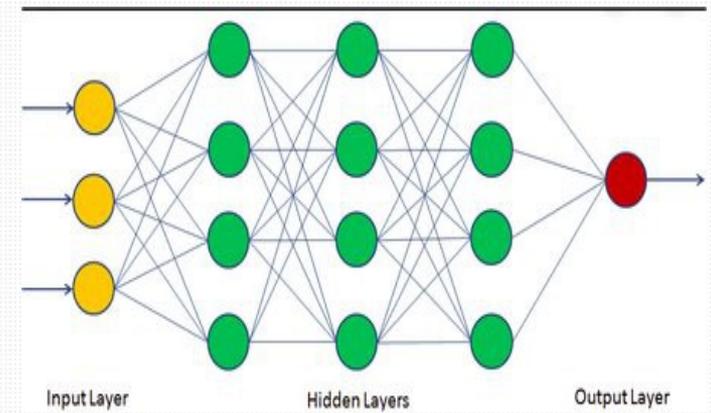
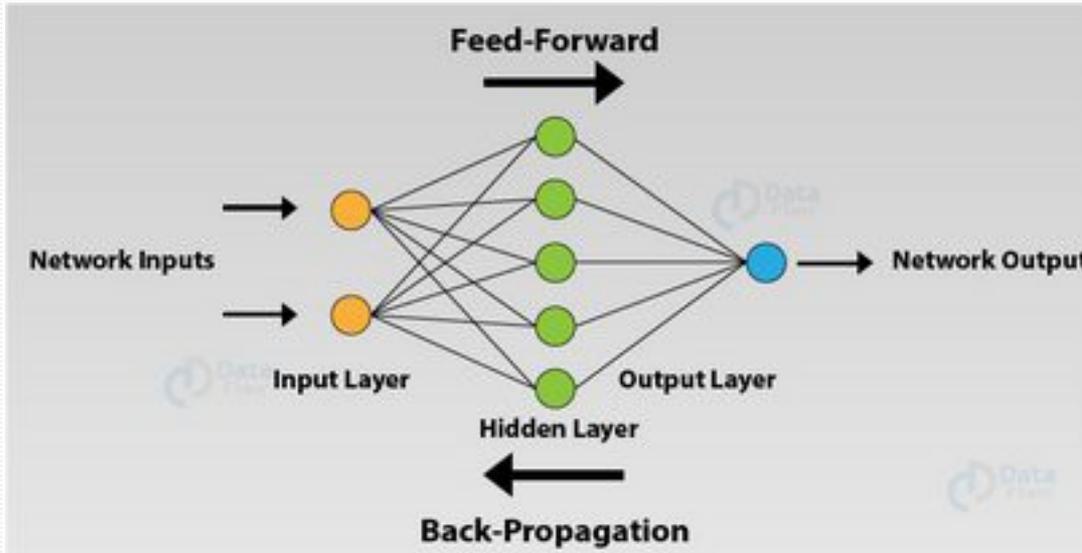
In this program, a problem is represented as a table of connections showing the distances between all pairs of states (initial, intervening, and final), and problem solving is modelled as a search through the problem space using permissible operators (actions), the task being to find a path of operators from an initial state to a goal state.

A commonplace initial state is being at home with a child and a car that has a flat battery, and the corresponding goal state in this case might be delivering the child to nursery school. The program is able to solve the Tower of Hanoi problem and similar.

[https://en.wikipedia.org/wiki/General\\_Problem\\_Solver](https://en.wikipedia.org/wiki/General_Problem_Solver)

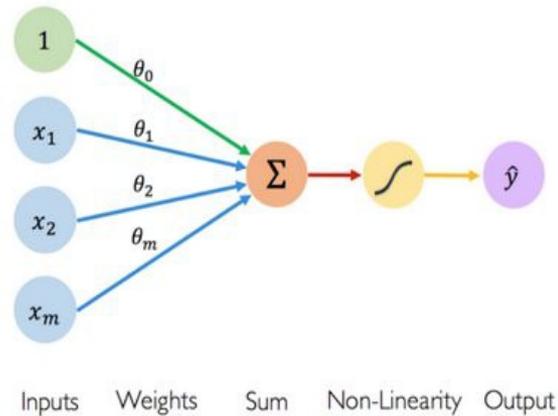
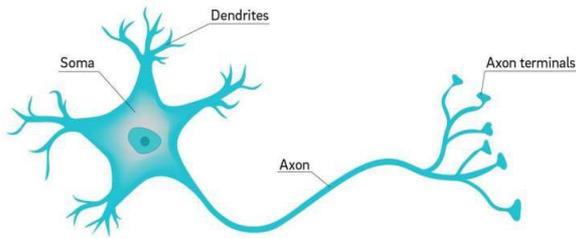
Herbert A. Simon, J. C. Shaw, and Allen Newell (1959)

# Neural Networks



[https://en.wikipedia.org/wiki/Artificial\\_neural\\_network](https://en.wikipedia.org/wiki/Artificial_neural_network)

## Neuron

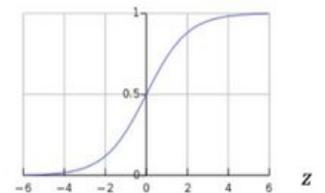


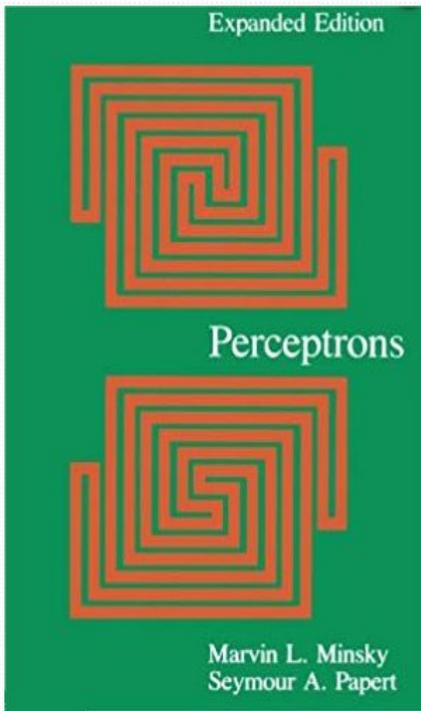
## Activation Functions

$$\hat{y} = g(\theta_0 + \mathbf{X}^T \boldsymbol{\theta})$$

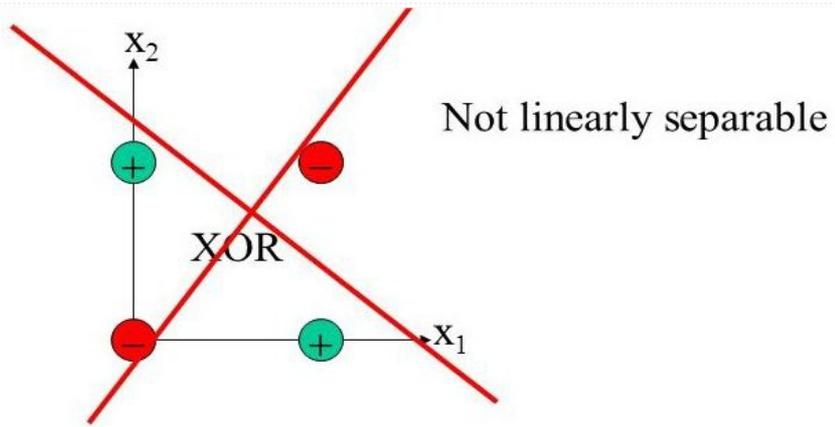
- Example: sigmoid function

$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$





After Minsky and Papert's paper - funding for neural network research dried up.



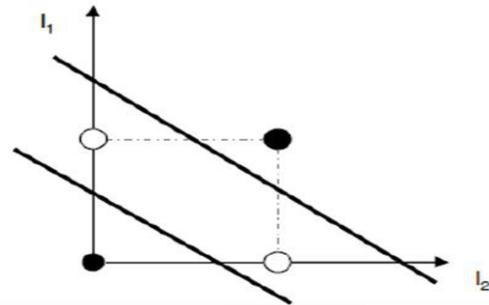
Minsky & Papert (1969)

Perceptrons can only represent linearly separable functions.

## Limitations of the perceptron model (cont'd)

Example: XOR gate (Minsky and Papert, 1969)

| XOR   |       |     |
|-------|-------|-----|
| $I_1$ | $I_2$ | out |
| 0     | 0     | 0   |
| 0     | 1     | 1   |
| 1     | 0     | 1   |
| 1     | 1     | 0   |



Finding cats in videos.  
Wired. 2012.



<https://www.wired.com/2012/06/google-x-neural-network>

## Building a "Cat Detector" using Convolutional Neural Networks

<https://medium.com/@curiously/tensorflow-for-hackers-part-iii-convolutional-neural-networks-c077618e590b>

But neural nets came back .. with many more layers...

## Common sense:

### Nixon Diamond

---

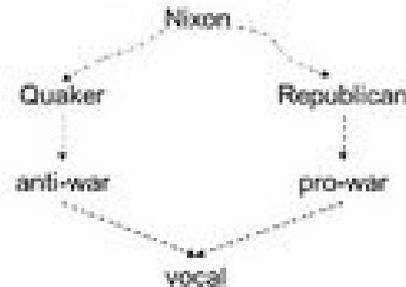
- Another classical example.
- Highlights the limitations of traditional formalisms.
- Nixon is both a Quaker and Republican.
- Quakers are anti-war.
- Republicans are pro-war.
- Both war supporters and opponents are vocal about their position.

1. Is Nixon pro-war or anti-war?

Unknown

2. Is Nixon vocal about his position?

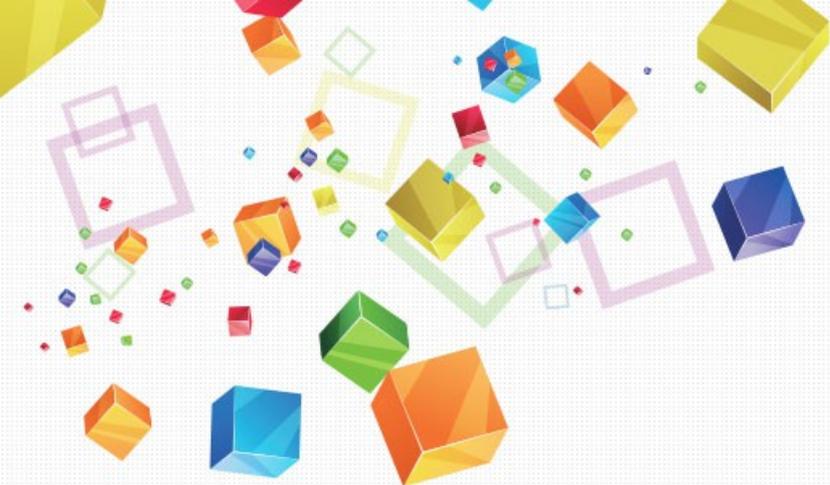
Yes



General Intelligence.

John McCarthy (AI pioneer) recognized AIs need for common sense.





## Heuristic Search.

- Lets only look at a part of the search space. Heuristics makes it easier to solve a problem. Heuristics (rules of thumb) directs the towards certain parts of the search space, and away from others.

## Planning.

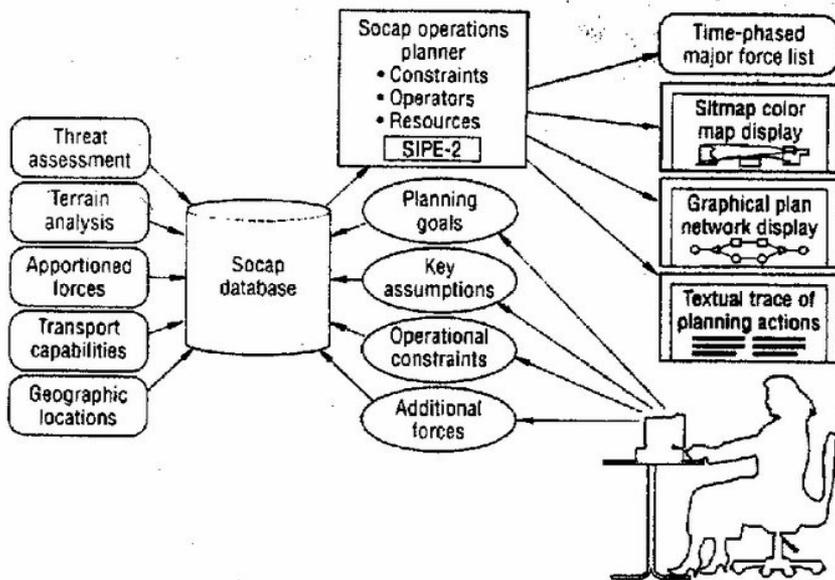
- A plan specifies a sequence of actions. A final goal, sub-goals, and sub-sub-goals etc.

## Knowledge representation.

- Often the hardest part of an AI problem is presenting the problem to the system in the first place.

See "Artificial Intelligence. A very short intro". p. 22- 25.

Margaret A. Boden.



## Military planning.

Joint Assistant for Development and Execution (JADE) is a U.S. military system used for planning the deployment of military forces in crisis situations.

## Cats & Understanding.

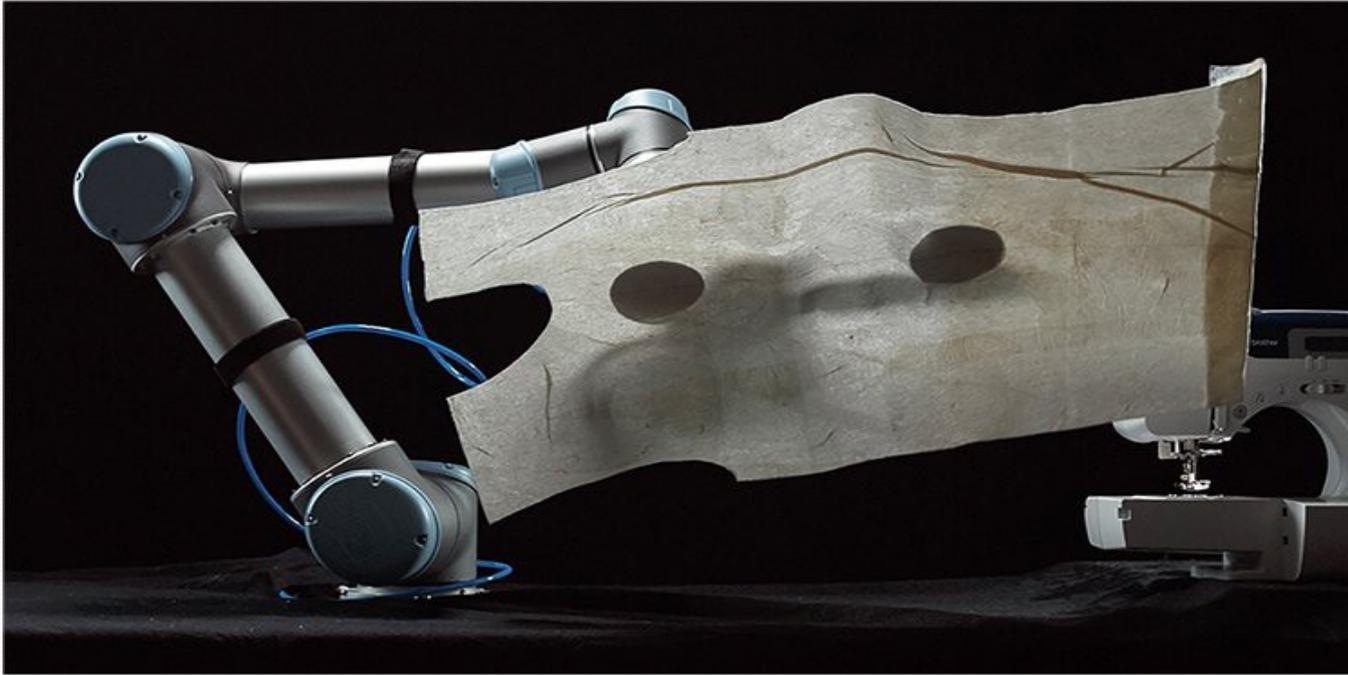


### Building a "Cat Detector" using Convolutional Neural Networks

<https://medium.com/@curiously/tensorflow-for-hackers-part-iii-convolutional-neural-networks-c077618e590b>

But - just like with the "rules based" systems.  
You can argue that Neural Nets (etc) is also just "rules".

The neural net really doesn't understand what a cat **is**.  
And there is no real understanding of what 3d space **is**.



AI understanding,

A robotic arm can fully assemble a T-shirt from start to finish. It followed a sequence of functions: a machine cuts cloth panels to be sewn, the panels are drenched with the polymer and stiffened, a robotic arm uses suction cups to lift each panel and position it in a sewing machine. Finally, the robotic arm lifts the stitched T-shirt off the sewing machine.

<https://money.cnn.com/2016/10/11/technology/robots-garment-manufacturing-sewbo/index.html>

"But - you can't ask a robot to put on a T-shirt. The head must go in first, and not via a sleeve. - But why? AI doesn't deal with such "topological problems" (yet). See "AI. Short Intro" – Margaret A. Boden.





## NELL: Never-Ending Language Learning

Can computers learn to read? We think so. "Read the Web" is a research project that attempts to create a computer system that learns over time to read the web. Since January 2010, our computer system called NELL (Never-Ending Language Learner) has been running continuously. Nell by Tom M. Mitchell.

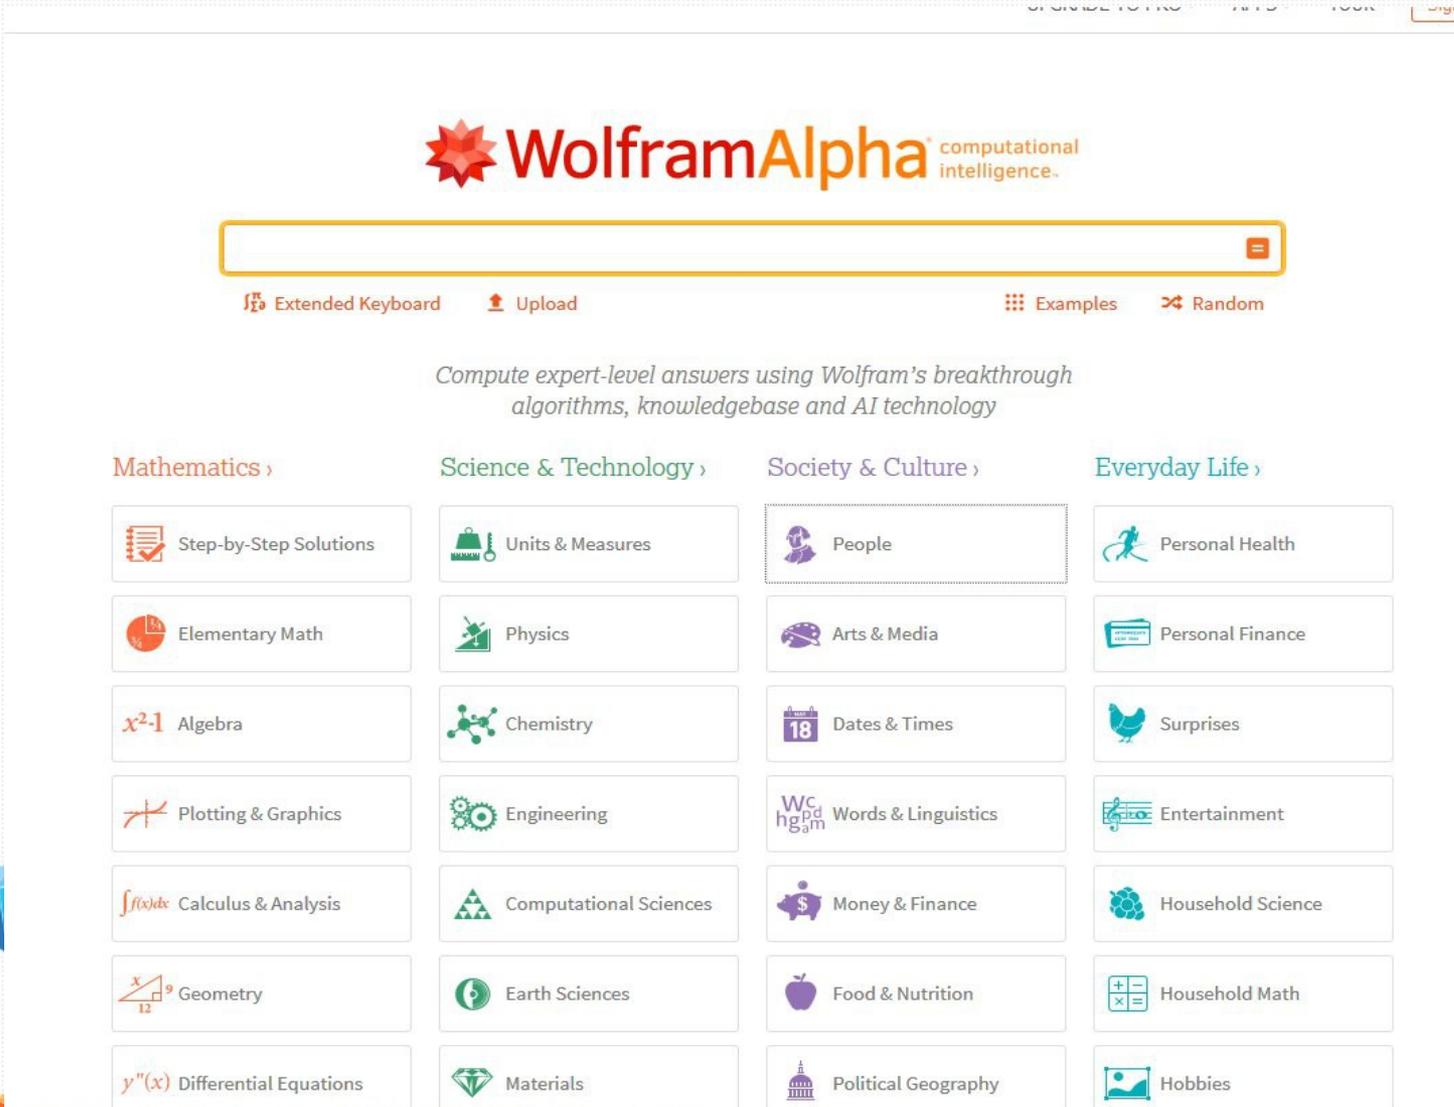
### Recently-Learned Facts

| instance   | iteration | date learned | confidence |   |
|--|-----------|--------------|------------|---|
| <u>judge_graham_morrow_gc</u> is a <u>judge</u>  | 1111      | 06-jul-2018  | 97.2       |       |
| <u>starbucks_coffee_company</u> is a <u>café</u>   | 1114      | 25-aug-2018  | 100.0      |       |
| <u>woolworth_s_department_store</u> is a <u>retail store</u>                             | 1111      | 06-jul-2018  | 99.7       |       |
| <u>heating_installations</u> is a <u>physical action</u>                                 | 1111      | 06-jul-2018  | 91.5       |       |
| <u>bad_pokertexas</u> is a <u>board game</u>   | 1111      | 06-jul-2018  | 100.0      |       |
| <u>richard</u> <u>held the position of king</u>  | 1116      | 12-sep-2018  | 100.0      |     |
| <u>internal_revenue_service</u> is <u>headquartered in</u> the country <u>us</u>         | 1116      | 12-sep-2018  | 100.0      |   |
| <u>washington_d_c</u> is a city that <u>lies on</u> the river <u>snake</u>               | 1113      | 15-aug-2018  | 99.3       |   |
| the companies <u>the_globe_and_mail</u> and <u>journal</u> <u>compete with</u> eachother | 1111      | 06-jul-2018  | 93.8       |   |
| <u>ford</u> is a company that <u>produces toyota_prius</u>                               | 1114      | 25-aug-2018  | 97.7       |   |

<https://twitter.com/cmunnell>

"I am a machine reading research project at Carnegie Mellon, periodically tweeting facts I read. Please follow me, and reply with corrections so I can improve!"

# Language & asking questions.



The image shows the WolframAlpha website interface. At the top, the WolframAlpha logo is displayed with the tagline "computational intelligence." Below the logo is a large search input field. Underneath the search bar are several utility links: "Extended Keyboard", "Upload", "Examples", and "Random". A descriptive sentence reads: "Compute expert-level answers using Wolfram's breakthrough algorithms, knowledgebase and AI technology". Below this, there are four main category headers: "Mathematics", "Science & Technology", "Society & Culture", and "Everyday Life". Each header is followed by a grid of sub-category buttons with icons. The "Society & Culture" category is highlighted with a dashed border around its "People" button.

WolframAlpha<sup>®</sup> computational intelligence.

Extended Keyboard Upload Examples Random

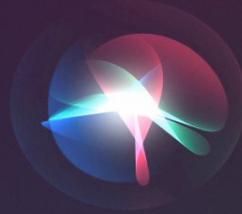
Compute expert-level answers using Wolfram's breakthrough algorithms, knowledgebase and AI technology

Mathematics › Science & Technology › Society & Culture › Everyday Life ›

|                        |                        |                     |                   |
|------------------------|------------------------|---------------------|-------------------|
| Step-by-Step Solutions | Units & Measures       | People              | Personal Health   |
| Elementary Math        | Physics                | Arts & Media        | Personal Finance  |
| Algebra                | Chemistry              | Dates & Times       | Surprises         |
| Plotting & Graphics    | Engineering            | Words & Linguistics | Entertainment     |
| Calculus & Analysis    | Computational Sciences | Money & Finance     | Household Science |
| Geometry               | Earth Sciences         | Food & Nutrition    | Household Math    |
| Differential Equations | Materials              | Political Geography | Hobbies           |

<https://www.wolframalpha.com/>

**Siri does more than ever.  
Even before you ask.**



Siri has access to Google Maps, Wikipedia, NY Times  
WolframAlpha etc.

Sometimes Siri gives ridiculous answers.

And if the users strays away from the domain of facts,  
Siri is lost.



<https://www.independent.co.uk/life-style/gadgets-and-tech/florida-man-accused-of-killing-his-roommate-asked-siri-where-to-hide-the-body-9665437.html>



Siri: AI assisted murder.





Bletchley Park — Home of Alan Turing during World War 2

The Loebner Prize 2018 was held in Bletchley Park

”The **Loebner Prize** is an international contest, where chatbots compete with each other to find the most humanlike. This is an incredibly difficult challenge, as recreating human conversation is such a complex task and instead of being an assistant or having a goal, human conversation covers a vast multitude of topics”.

<https://medium.com/pandorabots-blog/mitsuku-wins-loebner-prize-2018-3e8d98c5f2a7>

<https://aisb.org.uk/mitsuku-wins-2019-loebner-prize-and-best-overall-chatbot-at-aisb-x/>

## Loebner prize (2019):

Q: Whats the definition of robots

A: Robots are strong machines with arms. Robots build cars.

Please punctuate sentences.

Q: Robot means slave

A: Is slave an object?

Q: No

A: I don't know what a slave is.

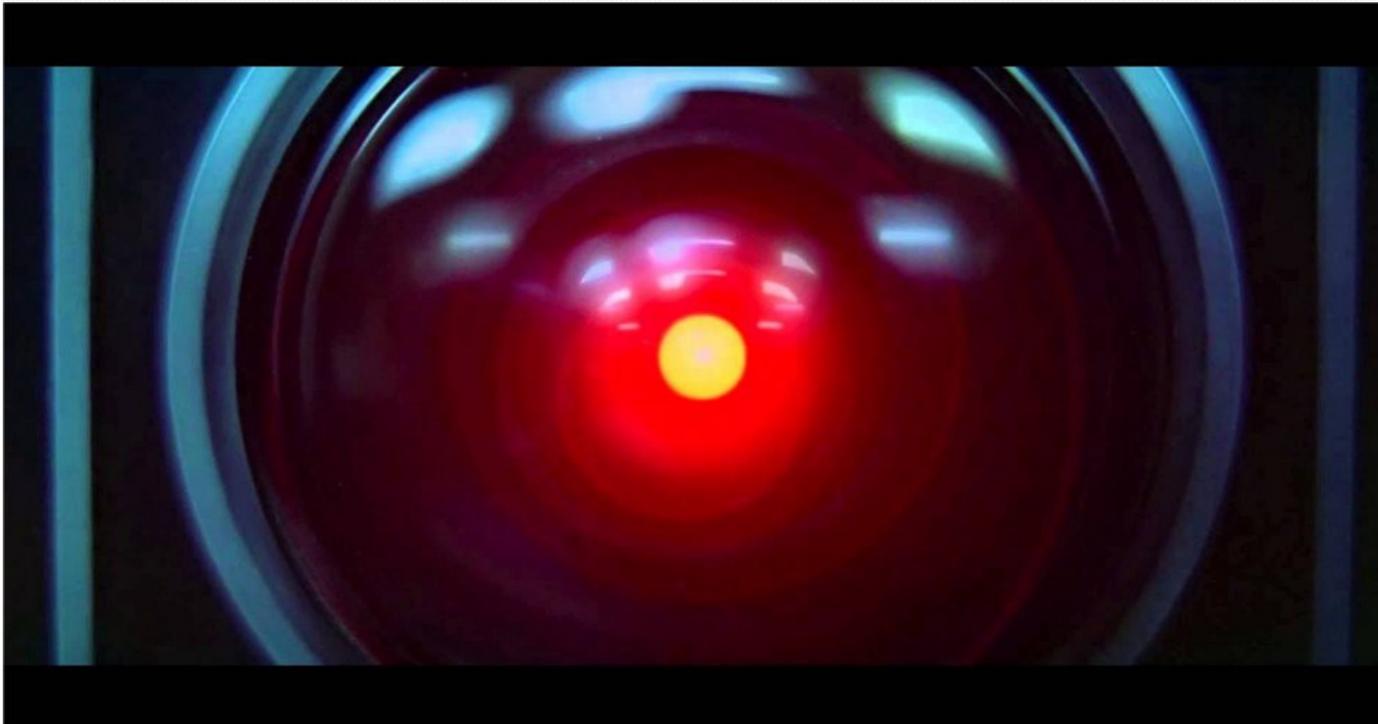
Q: A slave is a robot

A: Oh. A slave is not a robot though, because slaves can think.

Q: Robots can think

A: Slaves can think too, but robots can not think, because robots are hardware and metal.





AI in popular culture.

April 2018 marks the 50th anniversary of Stanley Kubrick's seminal, groundbreaking film 2001: A Space Odyssey. The main protagonist in the plot is a seemingly sentient artificial intelligence named HAL 9000, who appears to possess human-like intelligence and emotional capabilities.

<https://medium.com/stradigiai/50-years-after-2001-a-space-odyssey-can-we-build-a-hal-9000-64ed1b2e65d1>





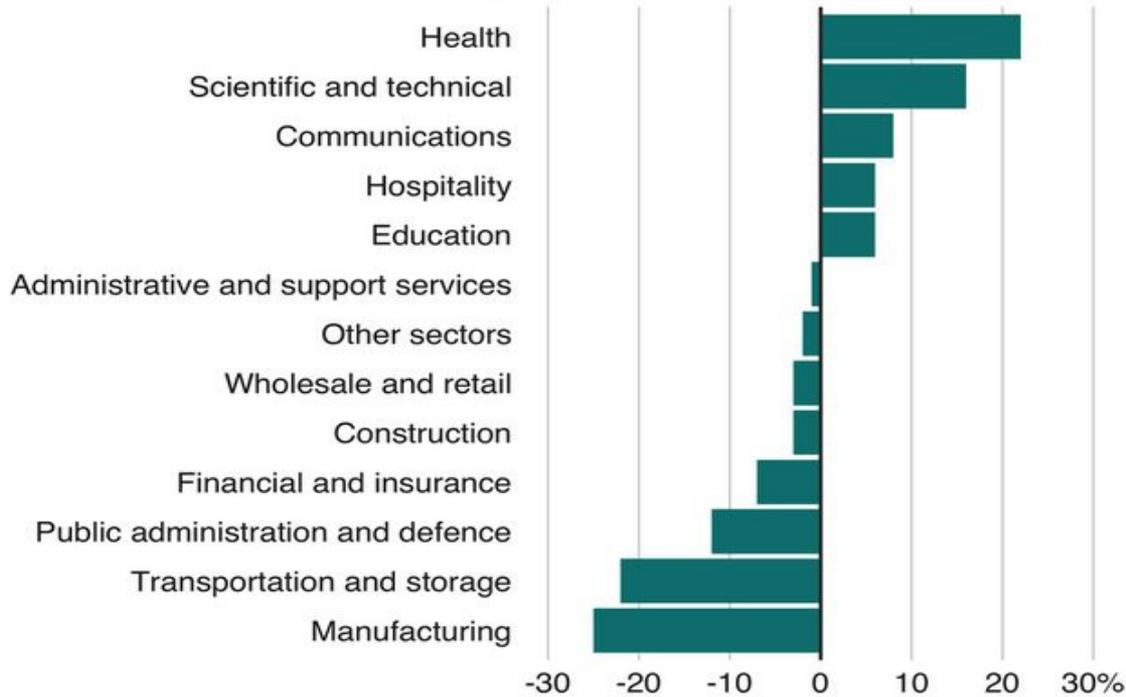
“I’m sorry Dave, I’m afraid I can’t do that”

<https://medium.com/datadriveninvestor/123-e32aa90ba472>



# How AI could change the job market

Estimated net job creation by industry sector, 2017-2037



Source: PwC

BBC

<https://www.bbc.com/news/business-44849492>

AI will create as many jobs as it displaces – report (2018).



**As Kevin Kelly  
pointed out the next  
10,000 startups will  
be AI + something**



What AI systems can't do (in 2019):

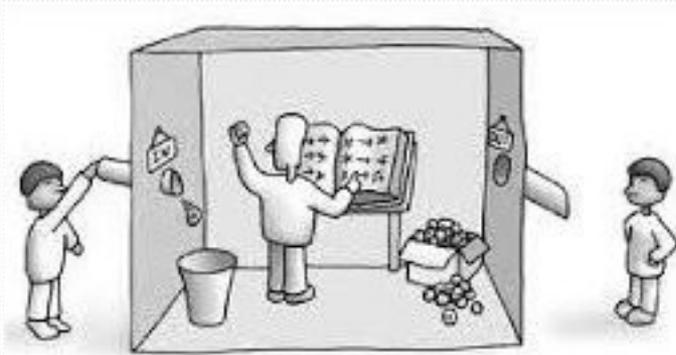
- "Lacks a human sense of relevance (within a area of expertise) ”.
- "AI focus on intellectual rationality, while ignoring social/emotional intelligence”.
- " Lack good psychological/computational theories about human minds work (which makes the prospects for AGI look dim)”.

Etc.

See: "Artificial Intelligence. A very short Introduction". p. 136.

Margaret A. Boden





Chinese Room

**The Chinese room** argument holds that a digital computer executing a program cannot be shown to have a "mind", "understanding" or "consciousness", regardless of how intelligently or human-like the program may make the computer behave.

The argument was first presented by philosopher John Searle in his paper, "Minds, Brains, and Programs".

<https://plato.stanford.edu/entries/chinese-room/>





In the philosophy of mind, the China brain thought experiment (also known as the Chinese Nation or Chinese Gym) considers what would happen if each member of the Chinese nation were asked to simulate the action of one neuron in the brain, using telephones or walkie-talkies to simulate the axons and dendrites that connect neurons. Would this arrangement have a mind or consciousness in the same way that brains do?

By Lawrence Davis (1974) & Ned Block (1978).





”Sure, you might enjoy discussing your personal memories with an AI?

But is it really a discussion?

Isn't it more like being able to talk about yourself, and remembering episodes from your life?

- What most people really want is acknowledgement of courage, suffering. From someone who shares the same human condition”.

?

See: "Artificial Intelligence. A very short Introduction". p. 143.  
Margaret A. Boden