

Introduction to Artificial Intelligence

IFT3335 Lecture 1: Introduction

Bang Liu, Jian-Yun Nie

2 Outline

- Organization of the course
- Introduction to Artificial Intelligence (AI)
- Risks and Benefits of AI

Organization of the course

4 Teachers


- ◎ Professors:
 - Bang Liu: bang.liu@umontreal.ca (English)
 - Jian-Yun Nie: nie@iro.umontreal.ca (French)
- ◎ Teaching assistant:
 - Hongjian Gu: hongjian.gu@mila.quebec (French)

5

Course materials


- Documents on Studium
- Slack workspace: https://join.slack.com/t/ift3335-a-h22/shared_invite/zt-11180lras-qlAp9rp1Gn9_0RtICChxMg
- Course webpage:
 - <https://www-labs.iro.umontreal.ca/~liubang/IFT%203335%20-%20Winter%202022.htm>
(in progress)
- Book recommended for reference:
 - *Artificial Intelligence: A Modern Approach*, 4th US ed. by Stuart Russell and Peter Norvig

- <http://aima.cs.berkeley.edu/>



6 Programming

- ⦿ Programming language
 - Python
- ⦿ Resources
 - <http://aima.cs.berkeley.edu>
 - <https://github.com/aimacode> (code)



7 Class Time

- ◎ Courses (January 12th ~ April 13th)
 - Wednesdays 12:30 pm ~ 13:30 pm
 - Thursdays 9:30 am ~ 11:30 am
- ◎ Practical Work (January 19th ~ April 13th)
 - Wednesdays 13:30 pm ~ 15:30 pm
- ◎ Exams
 - Mid-term: 12:30 ~ 14:30, March 9th
 - Final: 12:30 ~ 15:30, April 27th

Evaluation

- ⦿ Online exams
 - Mid-term exam: 20% (2 hours)
 - Final exam: 20% (2 hours)

- ⦿ Small Exercises (30%)
 - individual
 - 6 exercises on concepts or application of concepts (5% each)

- ⦿ Practical work (programming) (30%)
 - in groups of 2-3 people
 - 2 programming labs (15% each)

9 Plan (tentative)

	Topic	Exercise (quiz)	Practical work
1	Introduction and notion of agent (Chapter 1, 2)		
2	Search in the state space (Chapter 3)	Ex 1 (5%)	
3	Two player games (Chapter 5)	Ex 2 (5%)	
4	Logic and knowledge (Chapter 7~10)	Ex 3 (5%)	TP1: Sudoku game (15%)
5	Uncertainty and Bayesian Networks (Chapter 12, 13)		
	Reading week: February 28th		
	Mid-term exam (March 9th) (20%)		
6	Markov Decision and Reinforcement Learning (Chapter 16, 17)	Ex 4 (5%)	
7	Machine Learning (Chapter 19, 20)	Ex 5 (5%)	
8	Deep Learning (Chapter 21, 22)	Ex 6 (5%)	
9	Natural Language Processing (Chapter 23, 24)		TP2: learning / NLP (15%)
10	Revision		
	Final exam (April 27th) (20%)		

Introduction to Artificial Intelligence (AI)

What is AI ?

David is 11 years old.
He weighs 60 pounds.
He is 4 feet, 6 inches tall.
He has brown hair.

His love is real.
But he is not.



人工智能

A STEVEN SPIELBERG FILM


ARTIFICIAL INTELLIGENCE

WARNER BROS. PICTURES ... DREAMWORKS PICTURES ...
A FILM BY STEVEN SPIELBERG ... AI: ARTIFICIAL INTELLIGENCE ... CASTING BY HALLY JOEL OSMENT
EXECUTIVE PRODUCERS JUDY LAY ... PRODUCED BY FRANCIS O'CONNOR ... DIRECTED BY STEVEN SPIELBERG ... COSTUME DESIGNER WILLIAM HURT ... EDITOR STAN WINSTON ...
EXECUTIVE PRODUCERS MICHAEL KAHN, ACE ... EXECUTIVE PRODUCERS BOB RINGWOOD ... EXECUTIVE PRODUCERS JOHN WILLIAMS ...
EXECUTIVE PRODUCERS JAN HARLAN ... EXECUTIVE PRODUCERS WALTER F. PARKES ... EXECUTIVE PRODUCERS STEVEN SPIELBERG ... EXECUTIVE PRODUCERS IAN WATSON ...
EXECUTIVE PRODUCERS BRIAN KOPPELMAN ... EXECUTIVE PRODUCERS KATHLEEN KENNEDY ... EXECUTIVE PRODUCERS STEVEN SPIELBERG ... EXECUTIVE PRODUCERS BONNIE CURTIS ...
DREAMWORKS PICTURES ... SUMMER 2001 ...


12 Imagination about AI



Bicentennial Man (1999, Chris Columbus)



I, Robot (2004, Alex Proyas)




Star Wars (1977, George Lucas)



Blade Runner 2049 (2017, Denis Villeneuve)



Doraemon (1970, Fujiko Fujio)



Astro Boy (1952, Osamu Tezuka)

13 Imagination about AI



What is AI?

- “the scientific **understanding** of the mechanisms underlying **thought and intelligent behavior** and their **embodiment** in machines” (Association for the Advancement of Artificial Intelligence – AAAI <https://www.aaai.org/>)
 - Understanding of intelligence
 - Implantation in the machines
- Embodiment cognition
 - a theory that many aspects of cognition are shaped by the body
 - The dictionary defines “to embody” as “making visible”. Our movement and body make visible all of who we are: our mood, personality, history, family, and culture.

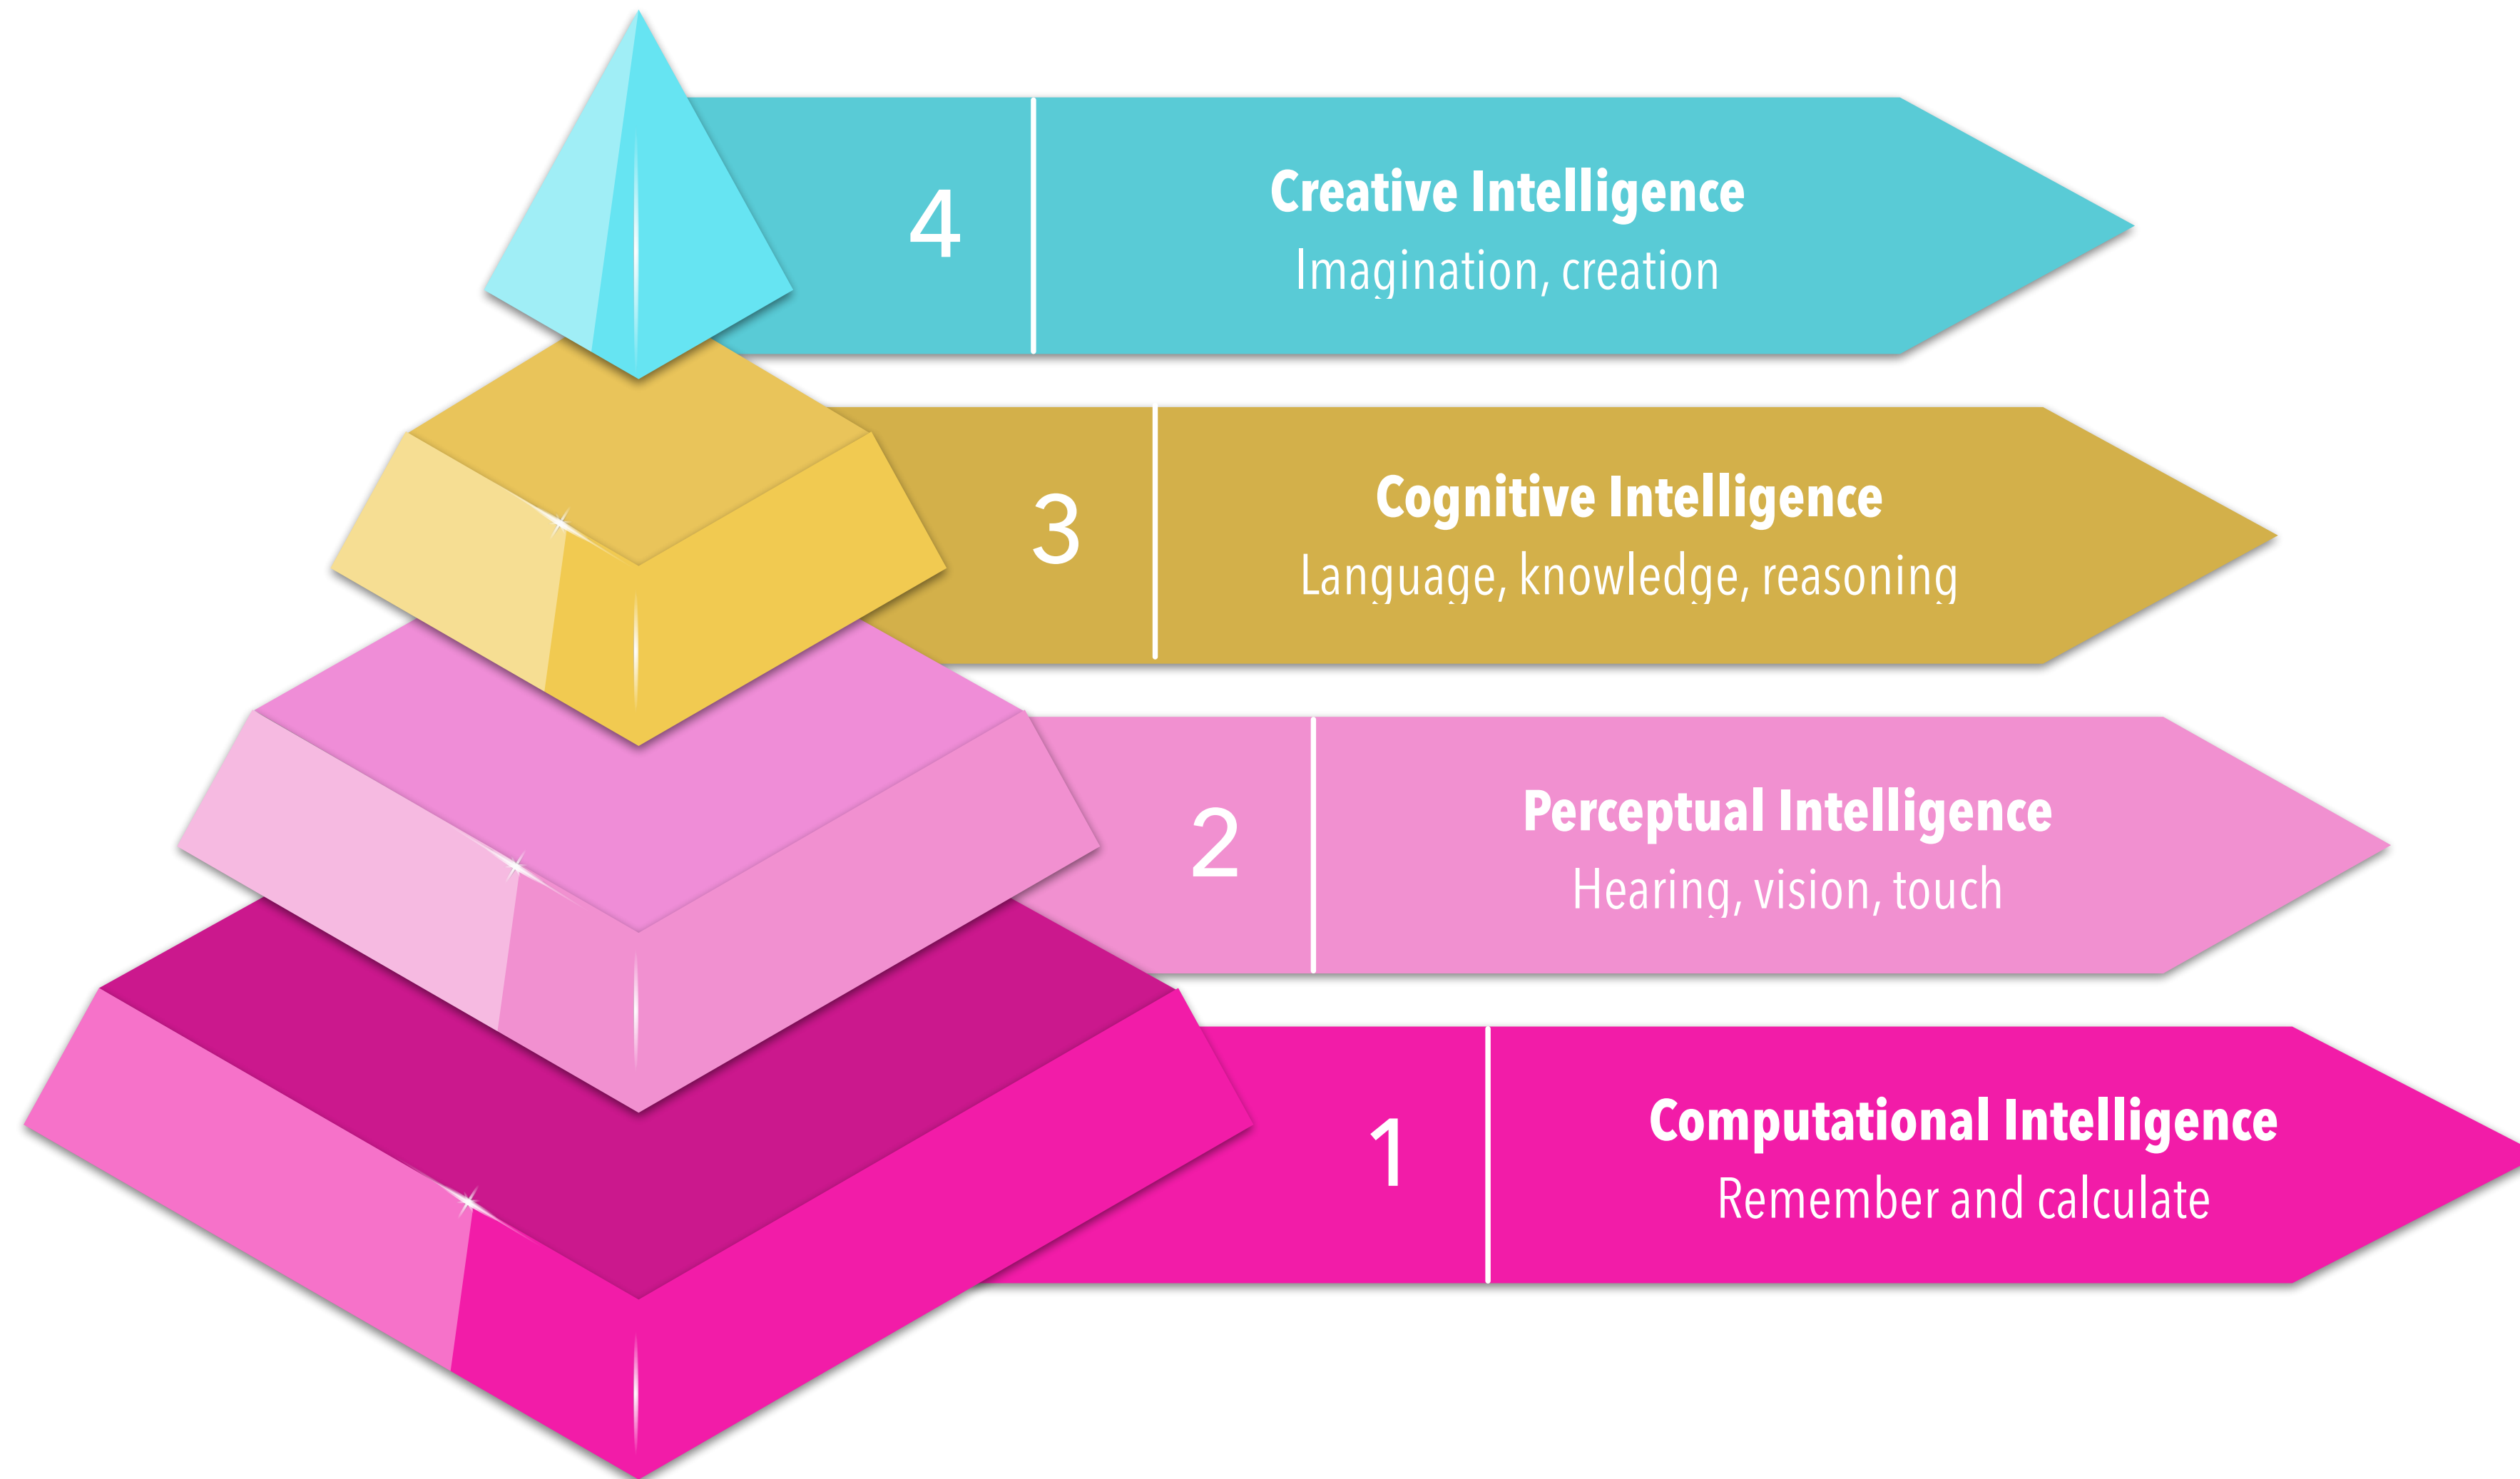
■ WELCOME TO THE ASSOCIATION FOR THE ADVANCEMENT OF ARTIFICIAL INTELLIGENCE!

Founded in 1979, the Association for the Advancement of Artificial Intelligence (AAAI) (formerly the American Association for Artificial Intelligence) is a nonprofit scientific society devoted to advancing the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines. AAAI aims to promote research in, and responsible use of, artificial intelligence. AAAI also aims to increase public understanding of artificial intelligence, improve the teaching and training of AI practitioners, and provide guidance for research planners and funders concerning the importance and potential of current AI developments and future directions. [More...](#)



Not a Member?
[Please Join!](#)

15 Levels of AI

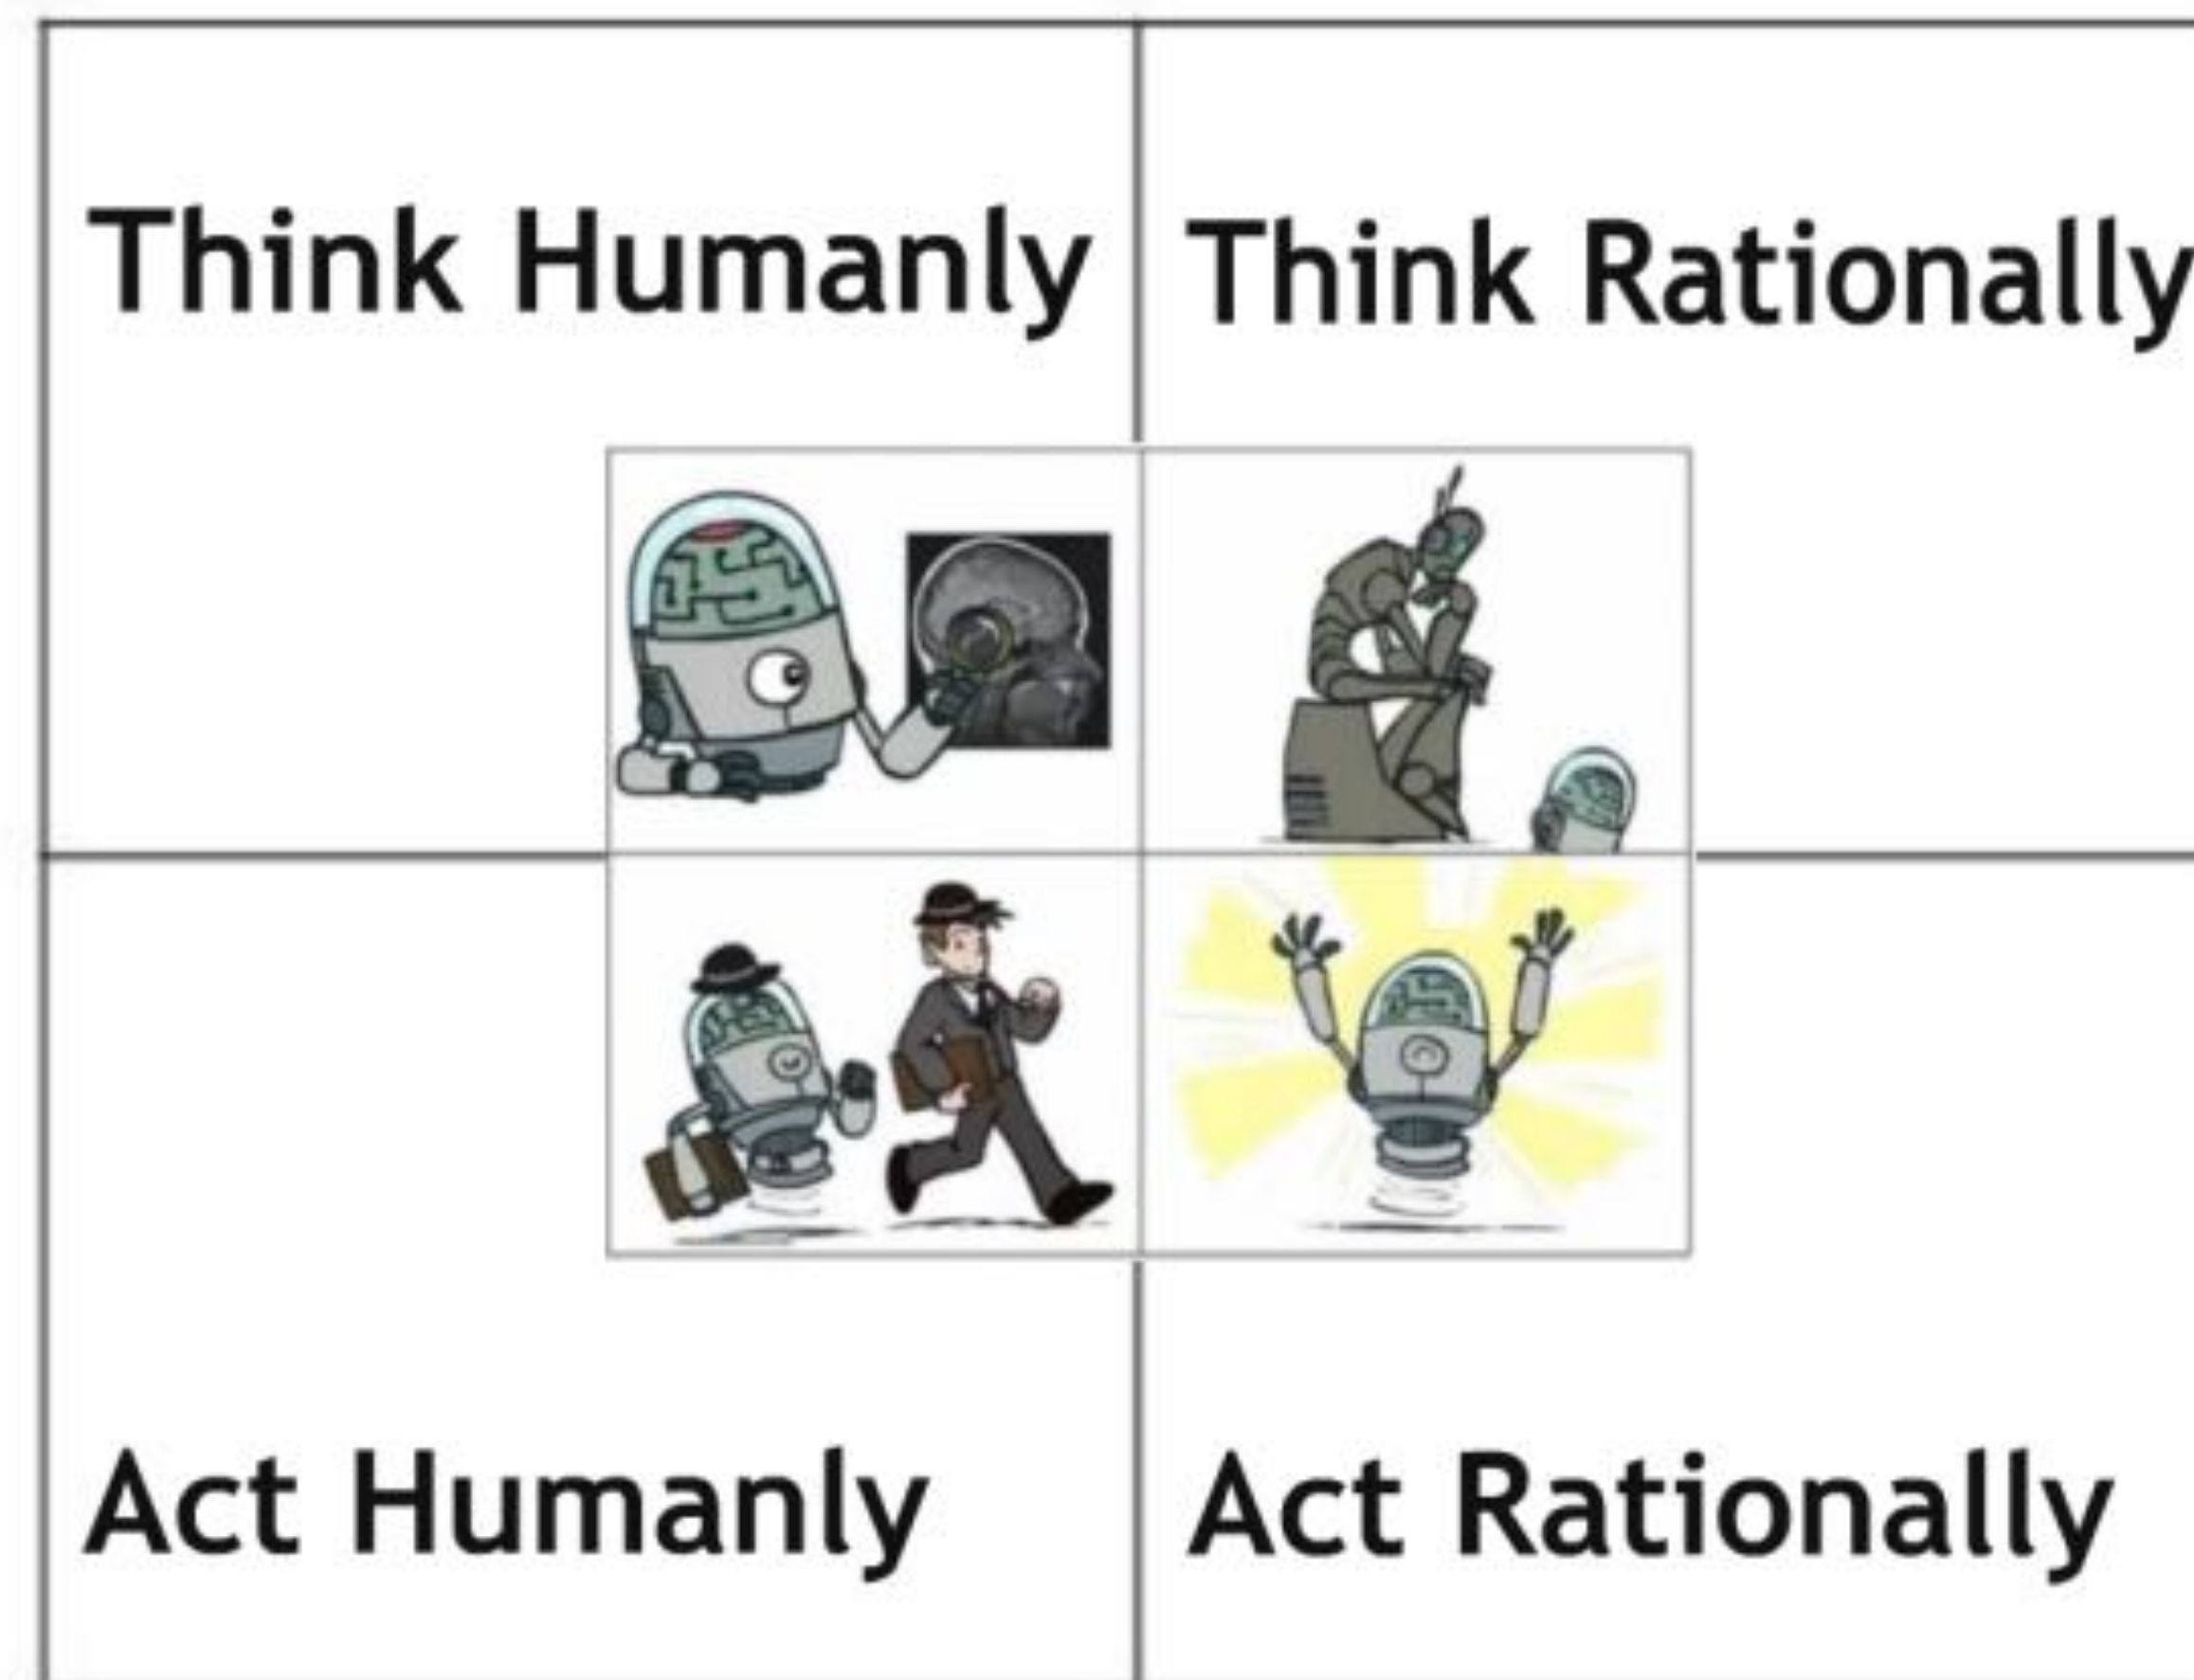


16 Intelligence = ?

- ◎ Perceptions
- ◎ Execution capacity
- ◎ Learning
- ◎ Knowledge
- ◎ Reasoning
- ◎ Capacity for analysis, abstraction / generalization
- ◎ Adaptability
- ◎ Communications
- ◎ Intelligence quotient vs. Emotional Quotient
- ◎ ...?


17 What is AI?

- **Humanly**: behave like human beings
- **Rationally**: act like a rational person
- **Thinking**: mechanism of functioning
- **Act**: result



Acting humanity: Turing test

- **Produce responses like humans**
- Turing (1950) “Computing machinery and intelligence”
- "Can the machine think" → "Can the machine behave intelligently?" (Like a human)
- Operational test for intelligent behavior: the imitation game
 - Question by the interrogator
 - Response by human or AI system
 - Can the interrogator distinguish between the human and the system?



19 Turing test

- First objective test
- Some limitations of the test
 - Communication is done in natural language (weak point for the system)
 - Limit oneself to human intelligence (e.g. unable to do complex calculations)
 - Easy to trap the system
- He states that in 2000, a machine could have a 30% chance of cheating on a person for 5 minutes.
- The article anticipated all the major arguments against AI for the next 50 years
- Suggested the components of AI: knowledge, reasoning, language comprehension, learning

Variations and alternatives to the Turing Test

☉ Variations:

- **Reverse Turing Test** -- where a human tries to convince a computer that it is not a computer. An example of this is a CAPTCHA.
- **Total Turing Test** -- where the questioner can also test perceptual abilities as well as the ability to manipulate objects.
- **Minimum Intelligent Signal Test** -- where only true/false and yes/no questions are given.

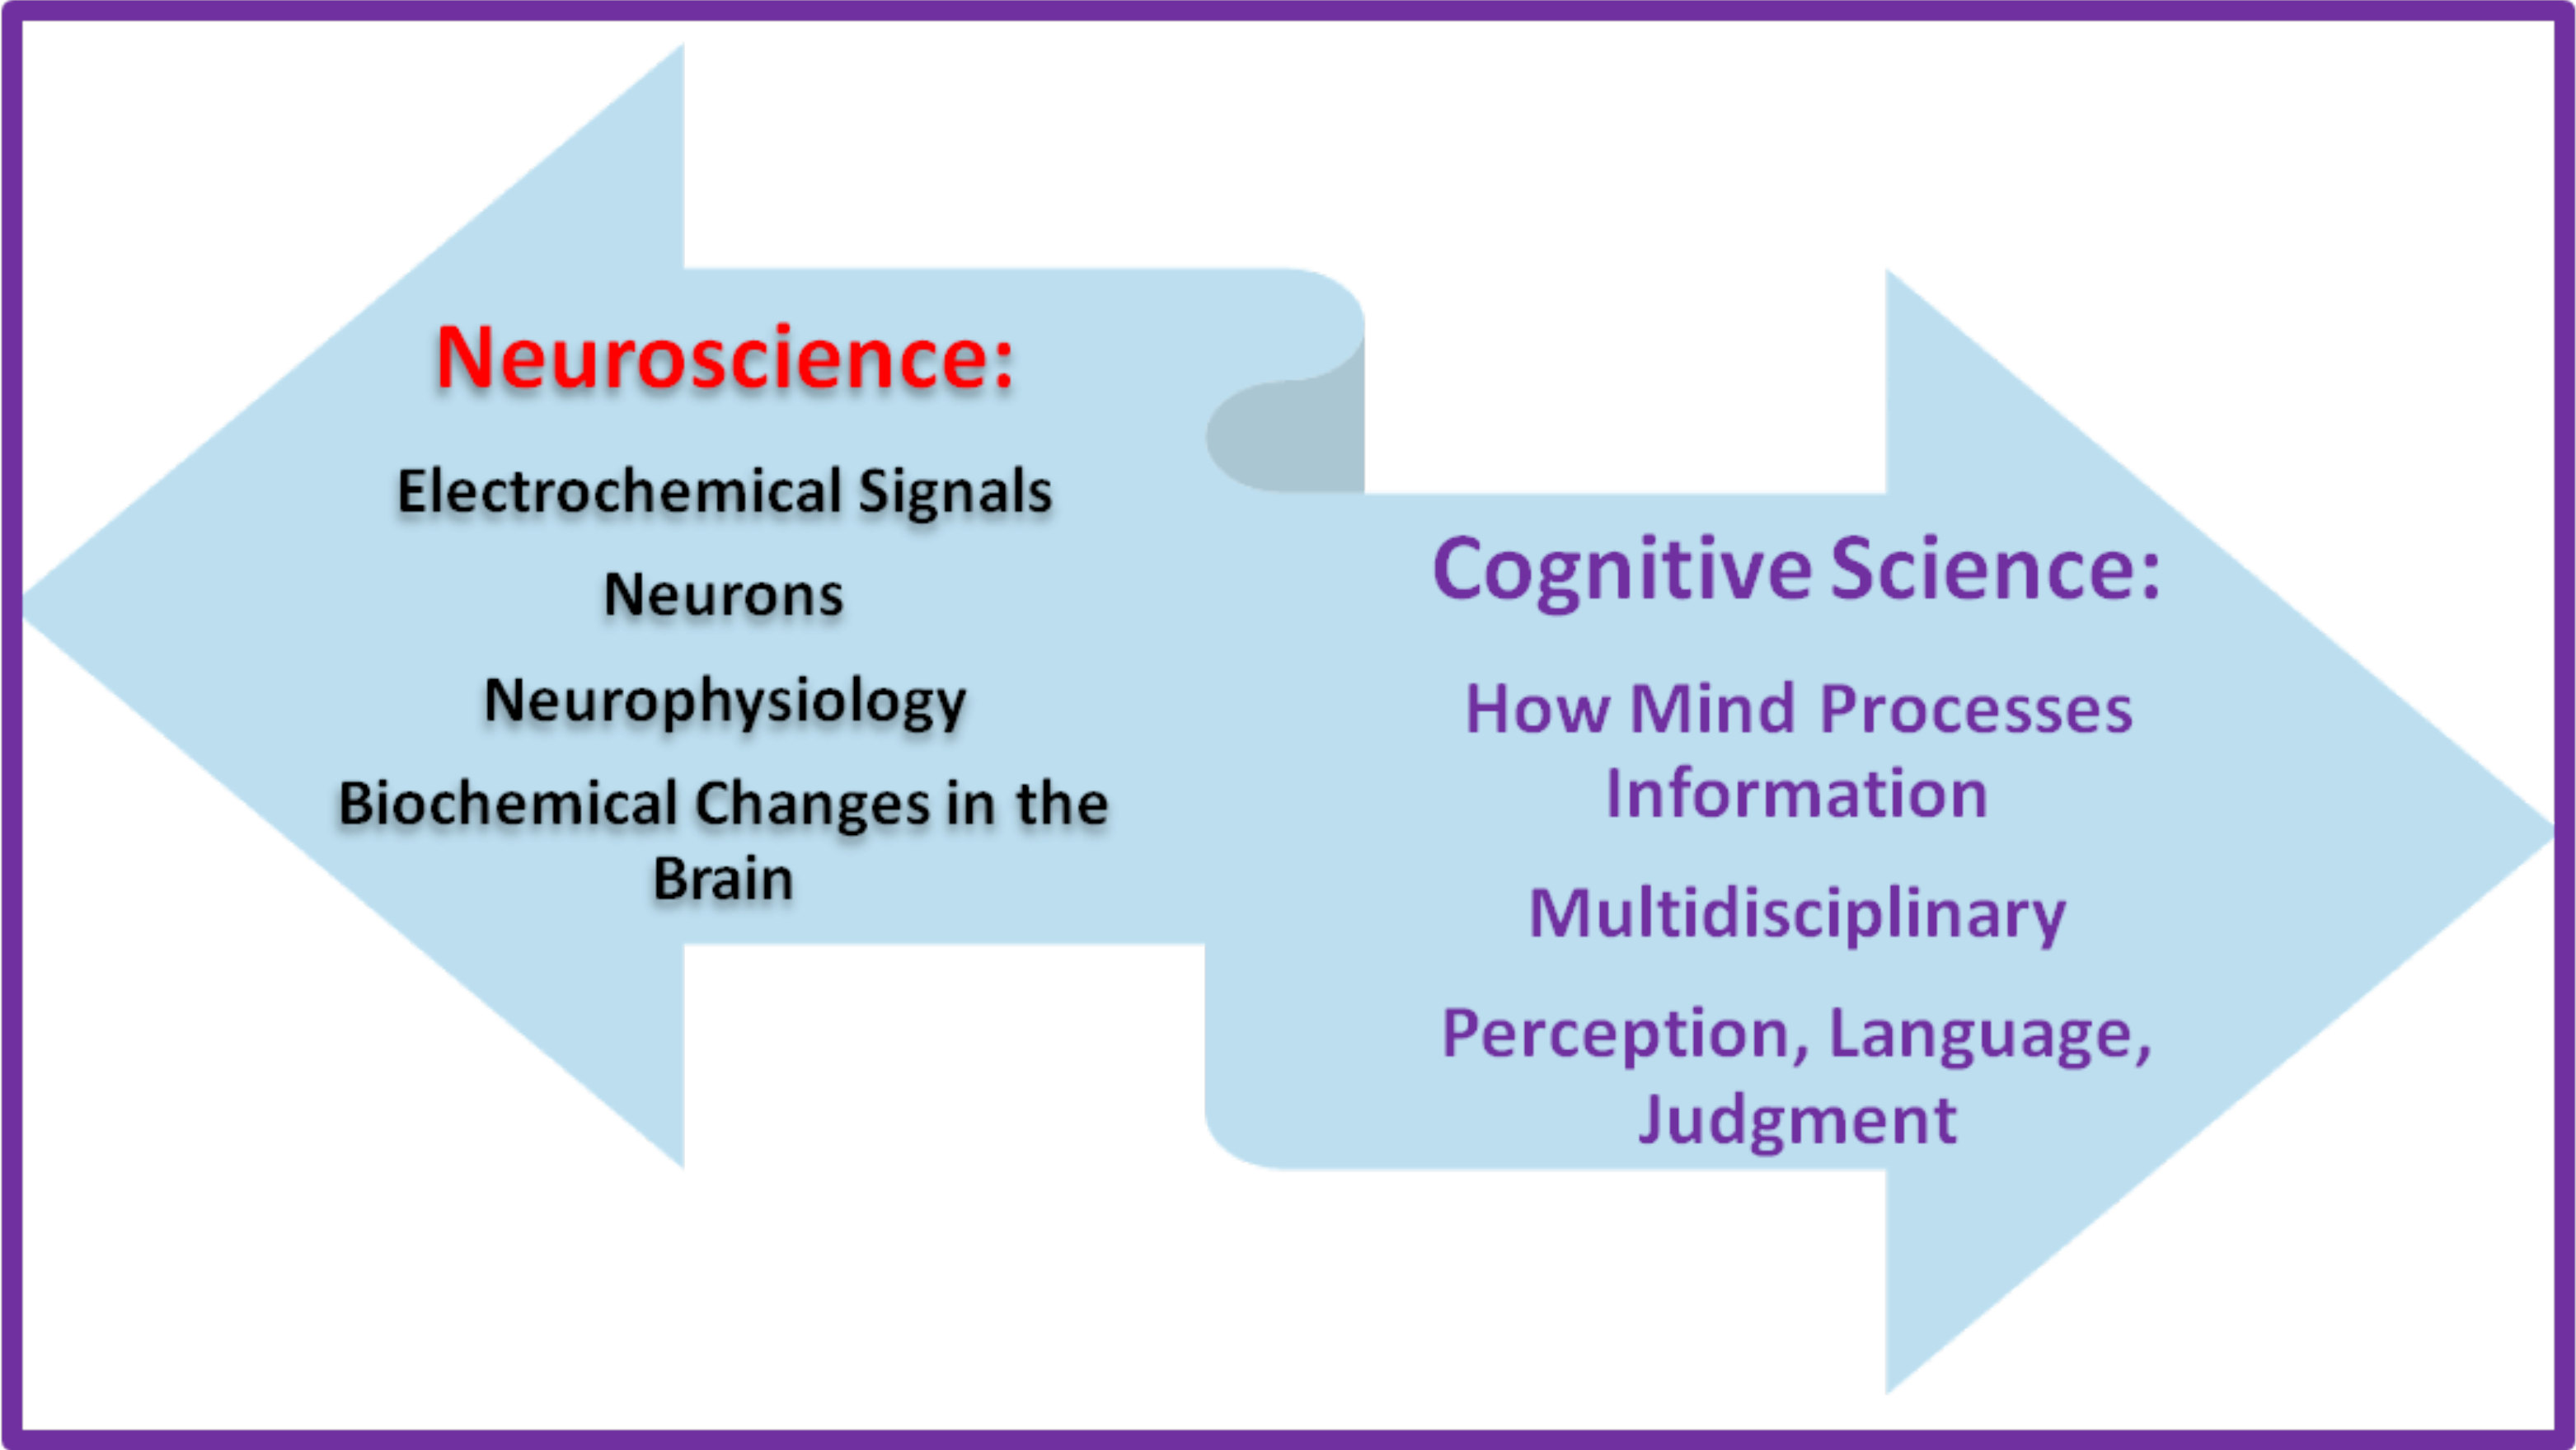
☉ Alternatives:

- **The Marcus Test** -- in which a program that can 'watch' a television show is tested by being asked meaningful questions about the show's content.
- **The Lovelace Test 2.0** -- which is a test made to detect AI through examining its ability to create art.
- **Winograd Schema Challenge** -- which is a test that asks multiple-choice questions in a specific format.


Think humanly: cognitive modeling

- **Systems should solve problems the same way humans do.**
- “Cognitive revolution” in 1960s: psychology based on information processing
The cognitive revolution was a period during the 1950s-1960s when cognitive psychology replaced Behaviorism and Psychoanalysis as the main approach in psychological fields. Increasing focus was placed on observable behaviors in conjunction with brain activity and structure.
- Need scientific theories for internal brain activities.
- Two approaches:
 - predict and test the behaviors of human subjects (top-down)
 - or direct identification of neurological data (bottom-up)
- These 2 approaches (roughly speaking, cognitive science and cognitive neuroscience) are now distinct fields of AI

Cognitive Science v.s. Neuroscience



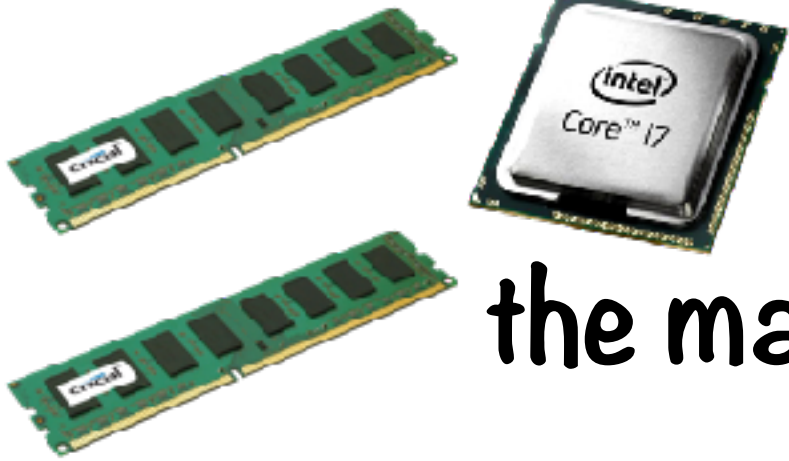
23 Toy example: how computer works?



24 Neuroscience



I can see
If remove





....
the machine is slow



Crack the hardware, but still don't know how it works.....

Cognitive Science




I can write words with 
And I can draw with 



Know what softwares can do, but don't know how they works...

Cognitive Science v.s. Neuroscience


Neuroscience is the scientific study of the nervous system, aims to to understand the fundamental and emergent properties of neurons, glia and neural circuits



The goal of **cognitive science** is to understand the principles of intelligence with the hope that this will lead to a better comprehension of the mind and of learning and to develop intelligent devices.

Think rationally: “laws of thought”

◎ Aristotle: What are the correct argument / thought processes?



◎ **Laws of thought**, the three fundamental laws of logic that are ascribed to Aristotle:

1. **The law of contradiction:** for all propositions p , it is impossible for both p and not p to be true
2. **The law of excluded middle:** either p or $\sim p$ must be true, there being no third or middle true proposition between them
3. **The principle of identity:** if a propositional function F is true of an individual variable x , then F is true of x , or: $F(x) \supset F(x)$, in which \supset means “formally implies.” Another formulation of the principle of identity asserts that a thing is identical with itself, or $(\forall x) (x = x)$, in which \forall means “for every”; or simply that x is x .

Think rationally: “laws of thought”

- ◎ Some Greek schools have developed different forms of logic: notation and derivation rules for thoughts; may or may not precede the idea of mechanization.
- ◎ Direct line from mathematics and philosophy to modern AI
- ◎ **Problems:**
 - Not all intelligent behavior is the result of logical deliberation.
 - What is the process of thinking? What thoughts should I be having?

Act rationally: rational agent

- ⊙ Rational behavior: Doing Right Things
 - Right things: those that are supposed to maximize the achievement of objectives, given the information available.
- ⊙ Rationality only concerns what decisions are made (not the thought process behind them)
- ⊙ Goals are expressed in terms of the utility of outcomes
- ⊙ **Being rational means maximizing your expected utility**

Rational agent

- An agent is an entity which perceives and acts.
- This course aims to develop rational agents
- In a very abstract way, an agent is a function which maps sequences of perception to actions:

$$[f: \mathcal{P}^* \rightarrow \mathcal{A}]$$

- For any class of environment and task, we aim to develop agents (or classes of agents) that produce the best performance
- But computational limitations mean that perfect rationality is not attainable.
 - → Develop the best program for given machine resources


History and State-of-the-art of AI

32 AI Prehistory

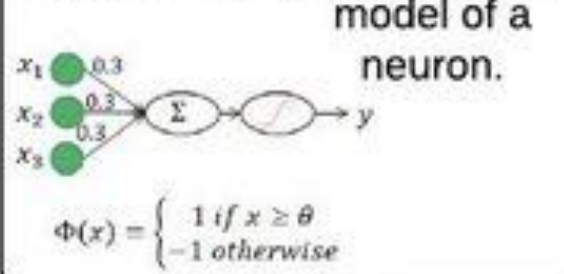
- Philosophy Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
- Mathematics Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability, bayes
- Economics/OR utility, decision theory
- Neuroscience physical substrate for mental activity
- Psychology phenomena of perception and motor control, experimental techniques
- Computer Science building fast computers, algorithms
- Linguistics knowledge representation, grammar
- Statistics bayes, big data

A brief history of AI

1943




First mathematical model of a neuron.



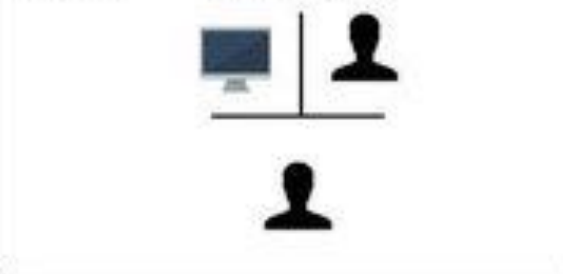
$$\Phi(x) = \begin{cases} 1 & \text{if } x \geq \theta \\ -1 & \text{otherwise} \end{cases}$$

Electronic Brain by McCulloch & Pitts

1950




Turing test is proposed.




Turing Test by Alan Turing

1952



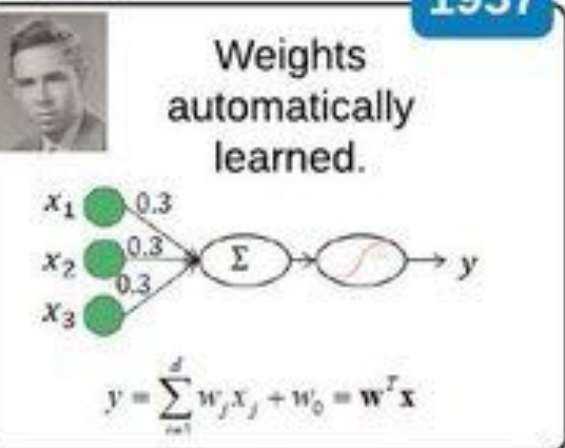
One of the first computer board games.



Checkers Program by Arthur Samuel

1957


Weights automatically learned.



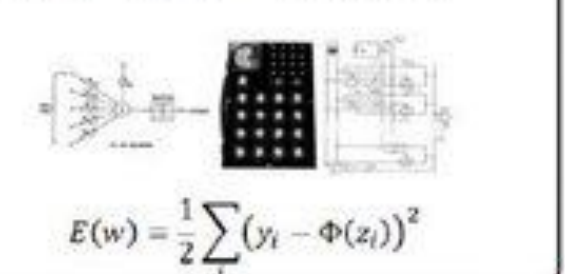
$$y = \sum_{j=1}^d w_j x_j + w_0 = \mathbf{w}^T \mathbf{x}$$

Perceptron by Frank Rosenblat

1960




The first cost function.



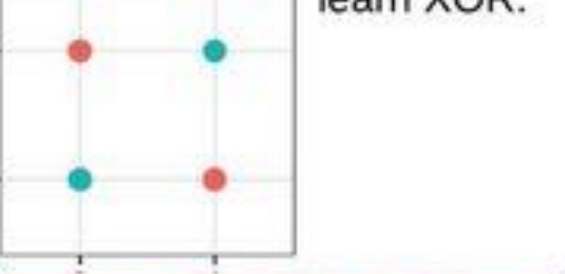
$$E(\mathbf{w}) = \frac{1}{2} \sum (y_i - \Phi(z_i))^2$$

ADALINE by Widrow & Hoff

1969



Perceptron cannot learn XOR.




XOR Problem by Minsky & Papert

AI Winter (1974-1980)

AI Winter (1987-1993)

1970




Backpropagation & automatic differentiation.

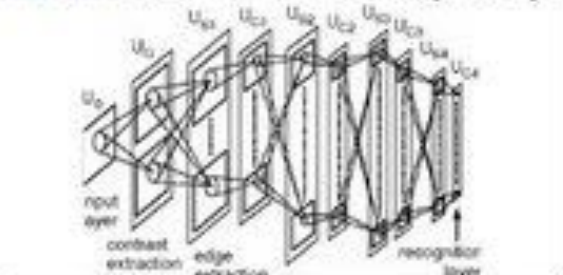
$$\frac{\partial y}{\partial x} = \frac{\partial y}{\partial w_1} \frac{\partial w_1}{\partial x}$$

Automatic differentiation by Seppo Linnainmaa

1979




The first convolutional neural network (CNN).

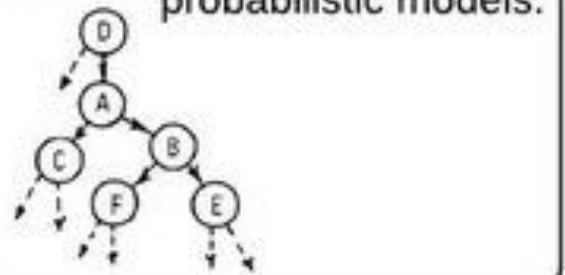


Neocognitron by Kunihiko Fukushima

1982




Foundation of graphical probabilistic models.

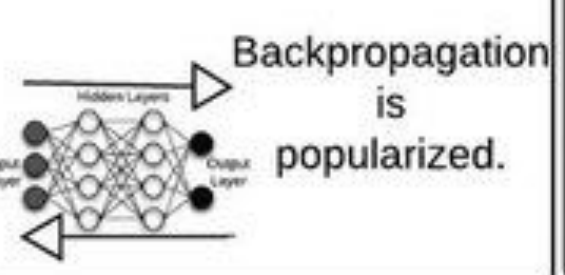


Bayesian Networks by Judea Pearl

1986




Backpropagation is popularized.

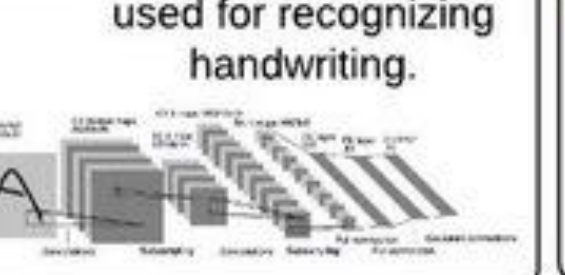


Backpropagation in MLP by Rumelhart, Hinton, Williams

1989




Convolutional neural networks (CNN) used for recognizing handwriting.

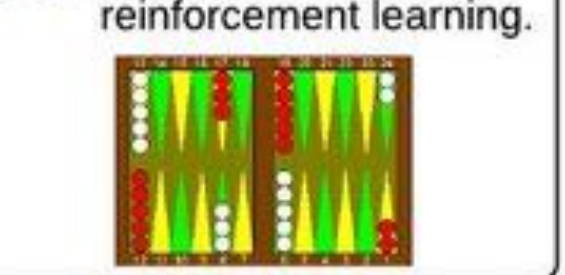


LeNet by Yann LeCun

1992




Almost champion-level backgammon, using reinforcement learning.

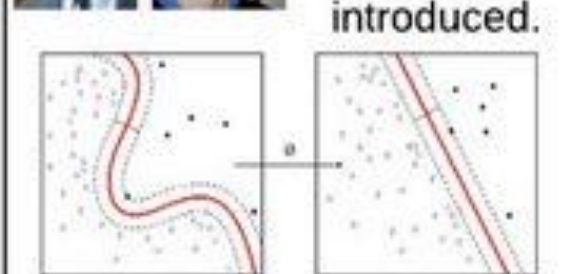


TD-Gammon by Gerald Tesauro

1995




Soft-margin SVM is introduced.




Support Vector Machines by S. Vapnik & Cortes

1995




MNIST is born.




MNIST by NIST

1996




DeepBlue beats Kasparov in chess.

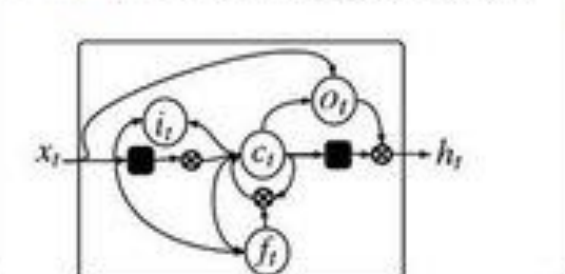


DeepBlue by IBM

1997




LSTM for addressing vanishing gradients.




Long Short-Term Memory (LSTM) by Schmidhuber.

2006




Deep learning is possible.




Deep Boltzman Machine by Ruslan & Hinton

2009




ImageNet, a large-scale image dataset is introduced.

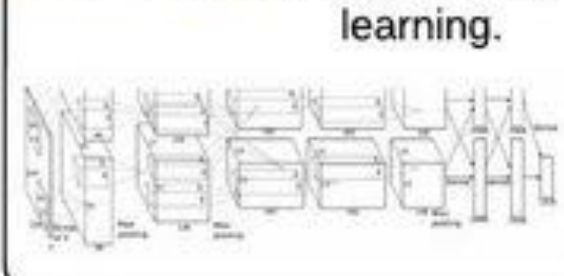


ImageNet by Fei-Fei Li

2012




First significant results in deep learning.

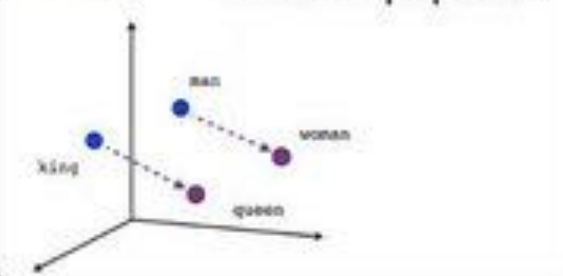


AlexNet by Krizhevsky & Hinton

2013




Word embeddings become popular.




Word2Vec by Tomas Mikolov

2014




Generative Adversarial Network (GAN) introduced.




GAN by Ian Goodfellow

2016




The first program to beat a professional Go player.




AlphaGo by DeepMind

2018




Pretrained language models.




BERT by Google AI

2019



Grandmaster level in Starcraft II.




AlphaStar by DeepMind



Deep Learning


Biological neurons



- **Dendrite:** Receives signals from other neurons
- **Soma:** Processes the information
- **Axon:** Transmits the output of this neuron
- **Synapse:** Point of connection to other neurons

***Biological Neurons:
An Overly Simplified Illustration***

The first computational model of a neuron



McCulloch-Pitts Neuron




Walter Pitts



Warren McCulloch

The first computational model of a neuron was proposed by Warren McCulloch (neuroscientist) and Walter Pitts (logician) in 1943.

Ancestor of neural networks: Perceptron



The "Perceptron" proposed by Cornell University professor Frank Rosenblatt in 1957 was the first to use algorithms to accurately define neural networks, and the first mathematical model with self-organization and self-learning capabilities. It is the ancestor of neural network models.




Frank Rosenblatt

Perceptrons can't do XOR

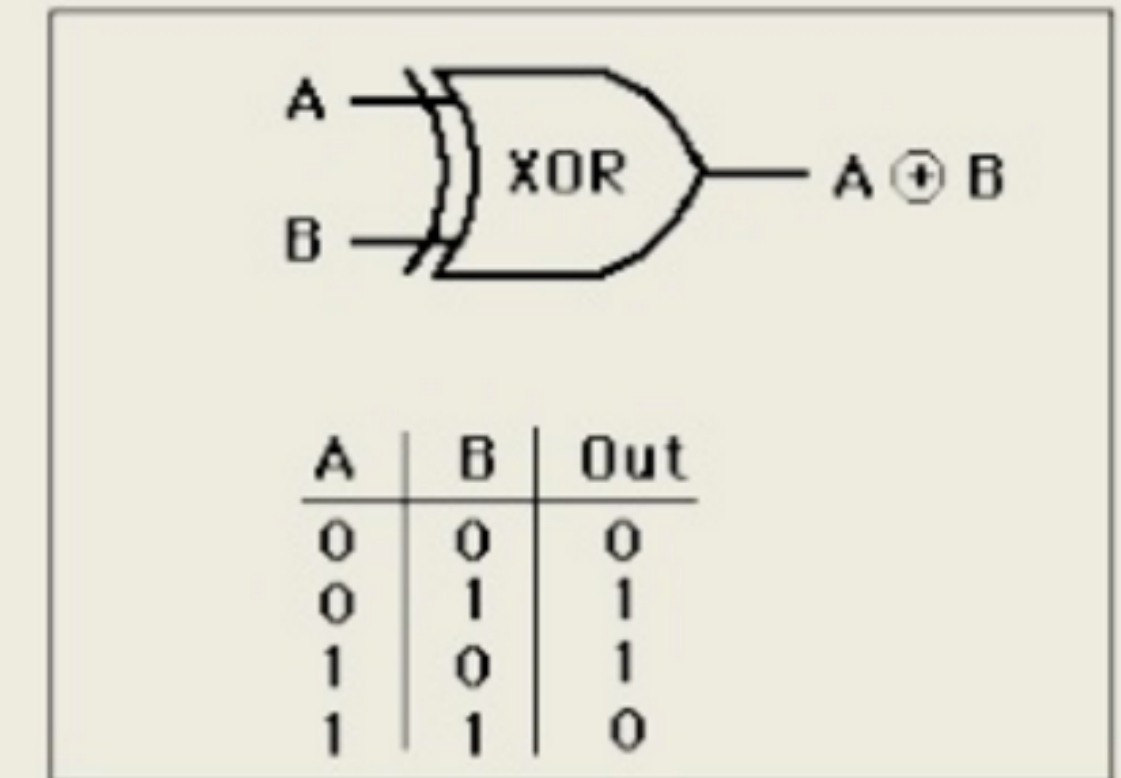
1969: Marvin Minsky and Seymour Pappert published a book - "Perceptron", which cited some limitations in Rosenblatt's technique. They found that **perceptrons can't do the Exclusive OR** logical operation.

Exclusive OR) is a digital logic gate that gives a true (1 or HIGH) output when the number of true inputs is odd. So, if input 0 and 1, or 1 and 0, the output is 1. In other cases, the outputs are 0. This book triggered a long AI Winter

1969: Perceptrons can't do XOR!



<http://www.i-programmer.info/images/stories/BabBag/AI/book.jpg>



<http://hyperphysics.phy-astr.gsu.edu/hbase/electronic/ietron/xor.gif>



Minsky & Papert

<https://constructingkids.files.wordpress.com/2013/05/minsky-papert-71-csolomon-x640.jpg>

Second Generation NNs: Back-propagation

In July 1986, Hinton and David Rumelhart co-published a paper in the journal Nature, "**Learning Representations by Back-propagating errors**", which was the first to systematically and concisely expound the application of **back-propagating algorithms** to neural network models. **Neural network research begins to re-arise.**

Learning representations by back-propagating errors

David E. Rumelhart*, **Geoffrey E. Hinton†**
& **Ronald J. Williams***

* Institute for Cognitive Science, C-015, University of California,
San Diego, La Jolla, California 92093, USA

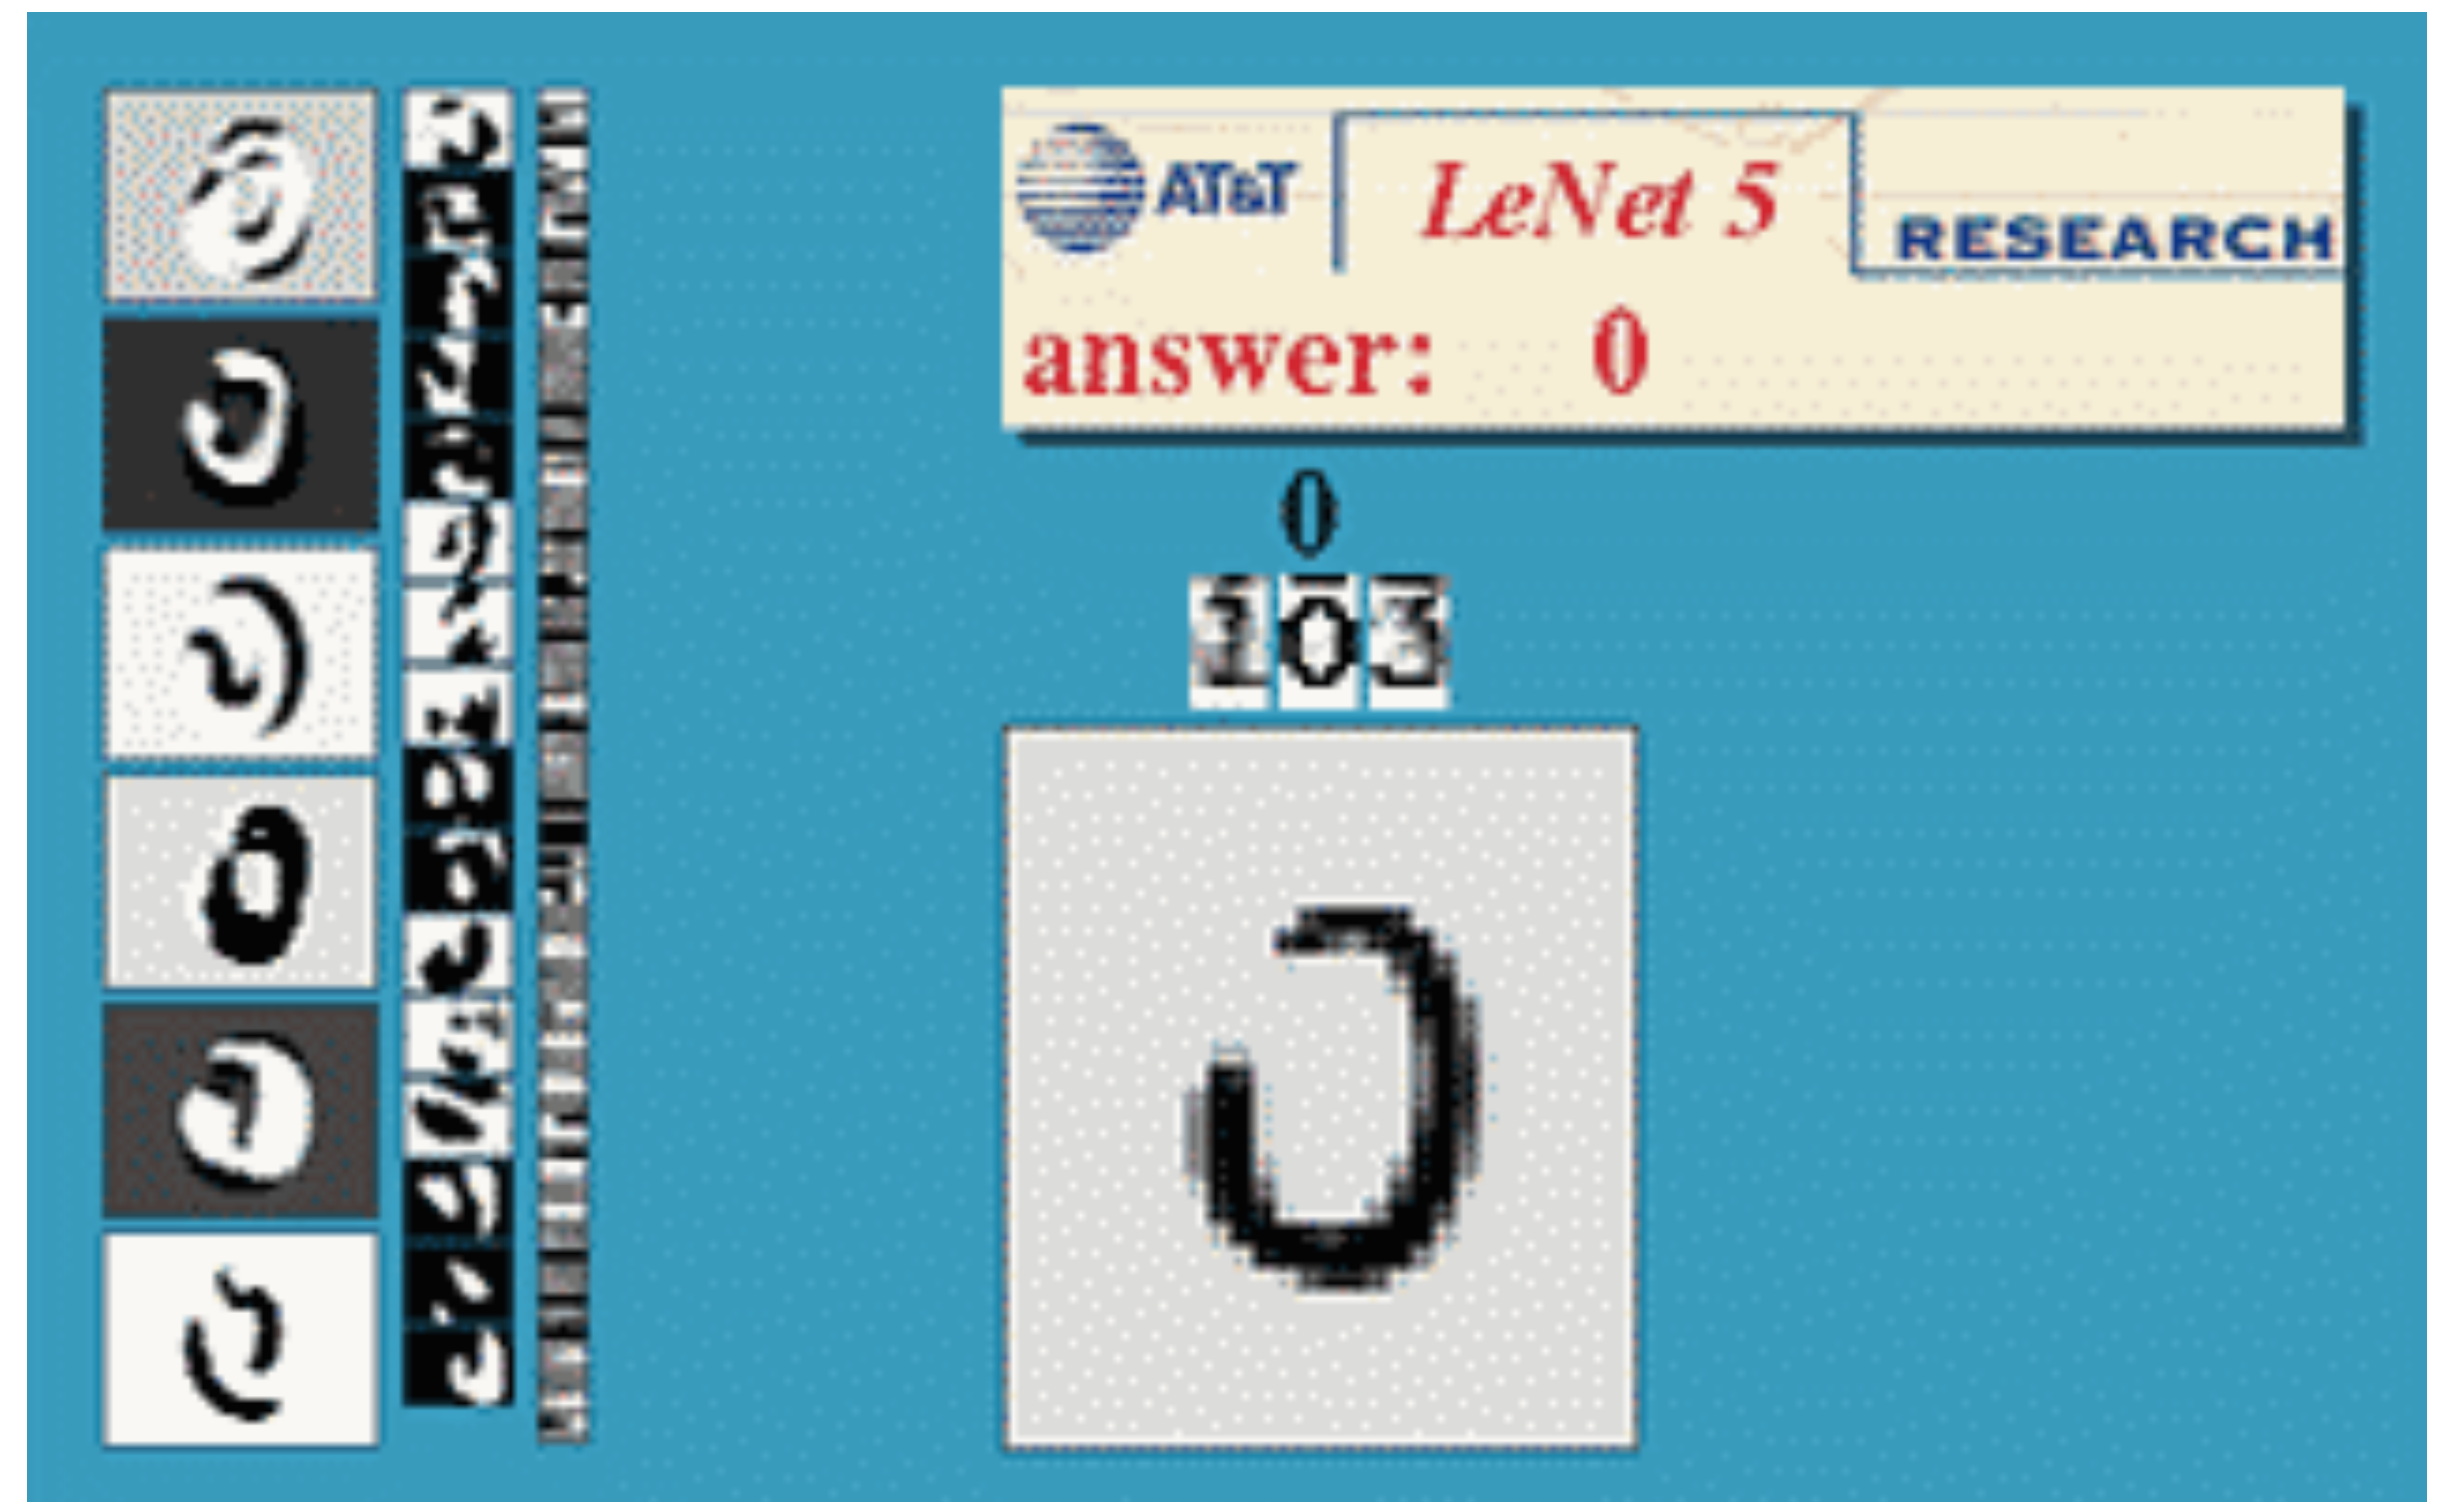
† Department of Computer Science, Carnegie-Mellon University,
Pittsburgh, Philadelphia 15213, USA

We describe a new learning procedure, **back-propagation**, for networks of neurone-like units. The procedure repeatedly adjusts the weights of the connections in the network so as to minimize a measure of the difference between the actual output vector of the net and the desired output vector. As a result of the weight adjustments, internal 'hidden' units which are not part of the input or output come to represent important features of the task domain, and the regularities in the task are captured by the interactions of these units. The ability to create useful new features distinguishes back-propagation from earlier, simpler methods such as the perceptron-convergence procedure¹.

<https://www.nature.com/articles/323533a0.pdf?origin=ppub>

Convolutional Neural Networks

In the late 1990s, Yann Lecun used a technique called “**Convolutional Neural Networks (CNN)**” for developing a commercial software to read handwritten numbers on bank checks. This **check recognition system** occupied nearly 20% of the US market in the late 1990s.



LeNet-5, convolutional neural networks
<http://yann.lecun.com/exdb/lenet/index.html>

41 Deep Belief Networks

ARTICLE

A fast learning algorithm for deep belief nets

Authors:  [Geoffrey E. Hinton](#),  [Simon Osindero](#),  [Yee-Whye Teh](#) [Authors Info & Affiliations](#)

Publication: Neural Computation • July 2006 • <https://doi.org/10.1162/neco.2006.18.7.1527>

<https://dl.acm.org/doi/10.1162/neco.2006.18.7.1527>

In 2006, Geoffrey Hinton introduced **Deep Belief Networks**, also introduced **layer-wise pretraining** technique, opened current deep learning era.

High-dimensional word embeddings and attention

In 2000, Yoshua Bengio authored the landmark paper, “**A Neural Probabilistic Language Model**,” that introduced high-dimension **word embeddings** as a representation of word meaning. Bengio’s insights had a huge and lasting impact on natural language processing tasks. His group also introduced a form of **attention** mechanism which led to breakthroughs in machine translation and form a key component of sequential processing with deep learning.

Journal of Machine Learning Research 3 (2003) 1137–1155

Submitted 4/02; Published 2/03

A Neural Probabilistic Language Model

Yoshua Bengio
Réjean Ducharme
Pascal Vincent
Christian Jauvin

BENGIOY@IRO.UMONTREAL.CA
DUCHARME@IRO.UMONTREAL.CA
VINCENTP@IRO.UMONTREAL.CA
JAUVINC@IRO.UMONTREAL.CA

*Département d’Informatique et Recherche Opérationnelle
Centre de Recherche Mathématiques
Université de Montréal, Montréal, Québec, Canada*

<https://mila.quebec/wp-content/uploads/2019/08/bengio03a.pdf>

Generative Adversarial Networks



Since 2010, Bengio's papers on generative deep learning, in particular the **Generative Adversarial Networks (GANs)** developed with Ian Goodfellow, have spawned a revolution in computer vision and computer graphics. In one fascinating application of this work, computers can actually create original images, reminiscent of the creativity that is considered a hallmark of human intelligence.

Generative Adversarial Nets

Ian J. Goodfellow, Jean Pouget-Abadie*, Mehdi Mirza, Bing Xu, David Warde-Farley,
Sherjil Ozair†, Aaron Courville, Yoshua Bengio‡
Département d'informatique et de recherche opérationnelle
Université de Montréal
Montréal, QC H3C 3J7

<https://arxiv.org/pdf/1406.2661.pdf>


Generative Adversarial Networks



ImageNet competition

At the end of 2012, Geoff Hinton and his PhD students Alex Krizhevsky and Ilya Sutskever took the first place in the **ImageNet image classification competition**, and increased the accuracy rate to 84.7%.

Relying on deep learning, they shocked the machine learning community. Since then, a large number of researchers have begun to enter this field.



In 2013, Google Hires Brains that Helped Supercharge Machine Learning

Table 1: Major milestones that will be covered in this paper

Year	Contributer	Contribution
300 BC	Aristotle	introduced Associationism, started the history of human's attempt to understand brain.
1873	Alexander Bain	introduced Neural Groupings as the earliest models of neural network, inspired Hebbian Learning Rule.
1943	McCulloch & Pitts	introduced MCP Model, which is considered as the ancestor of Artificial Neural Model.
1949	Donald Hebb	considered as the father of neural networks, introduced Hebbian Learning Rule, which lays the foundation of modern neural network.
1958	Frank Rosenblatt	introduced the first perceptron, which highly resembles modern perceptron.
1974	Paul Werbos	introduced Backpropagation
1980	Teuvo Kohonen	introduced Self Organizing Map
	Kunihiko Fukushima	introduced Neocogitron, which inspired Convolutional Neural Network
1982	John Hopfield	introduced Hopfield Network
1985	Hilton & Sejnowski	introduced Boltzmann Machine
1986	Paul Smolensky	introduced Harmonium, which is later known as Restricted Boltzmann Machine
	Michael I. Jordan	defined and introduced Recurrent Neural Network
1990	Yann LeCun	introduced LeNet, showed the possibility of deep neural networks in practice
1997	Schuster & Paliwal	introduced Bidirectional Recurrent Neural Network
	Hochreiter & Schmidhuber	introduced LSTM, solved the problem of vanishing gradient in recurrent neural networks
2006	Geoffrey Hinton	introduced Deep Belief Networks, also introduced layer-wise pretraining technique, opened current deep learning era.
2009	Salakhutdinov & Hinton	introduced Deep Boltzmann Machines
2012	Geoffrey Hinton	introduced Dropout, an efficient way of training neural networks

On the Origin of Deep Learning

47 2018 Turing Award



<https://awards.acm.org/about/2018-turing>

AlphaGo vs. Lee Sedol


4-1

March 2016

ALPHAGO




From AlphaGo to MuZero




MuZero: Mastering Go, chess, shogi and Atari without rules

AlphaFold 2



T1037 / 6vr4
90.7 GDT
(RNA polymerase domain)



T1049 / 6y4f
93.3 GDT
(adhesin tip)

- Experimental result
- Computational prediction

“

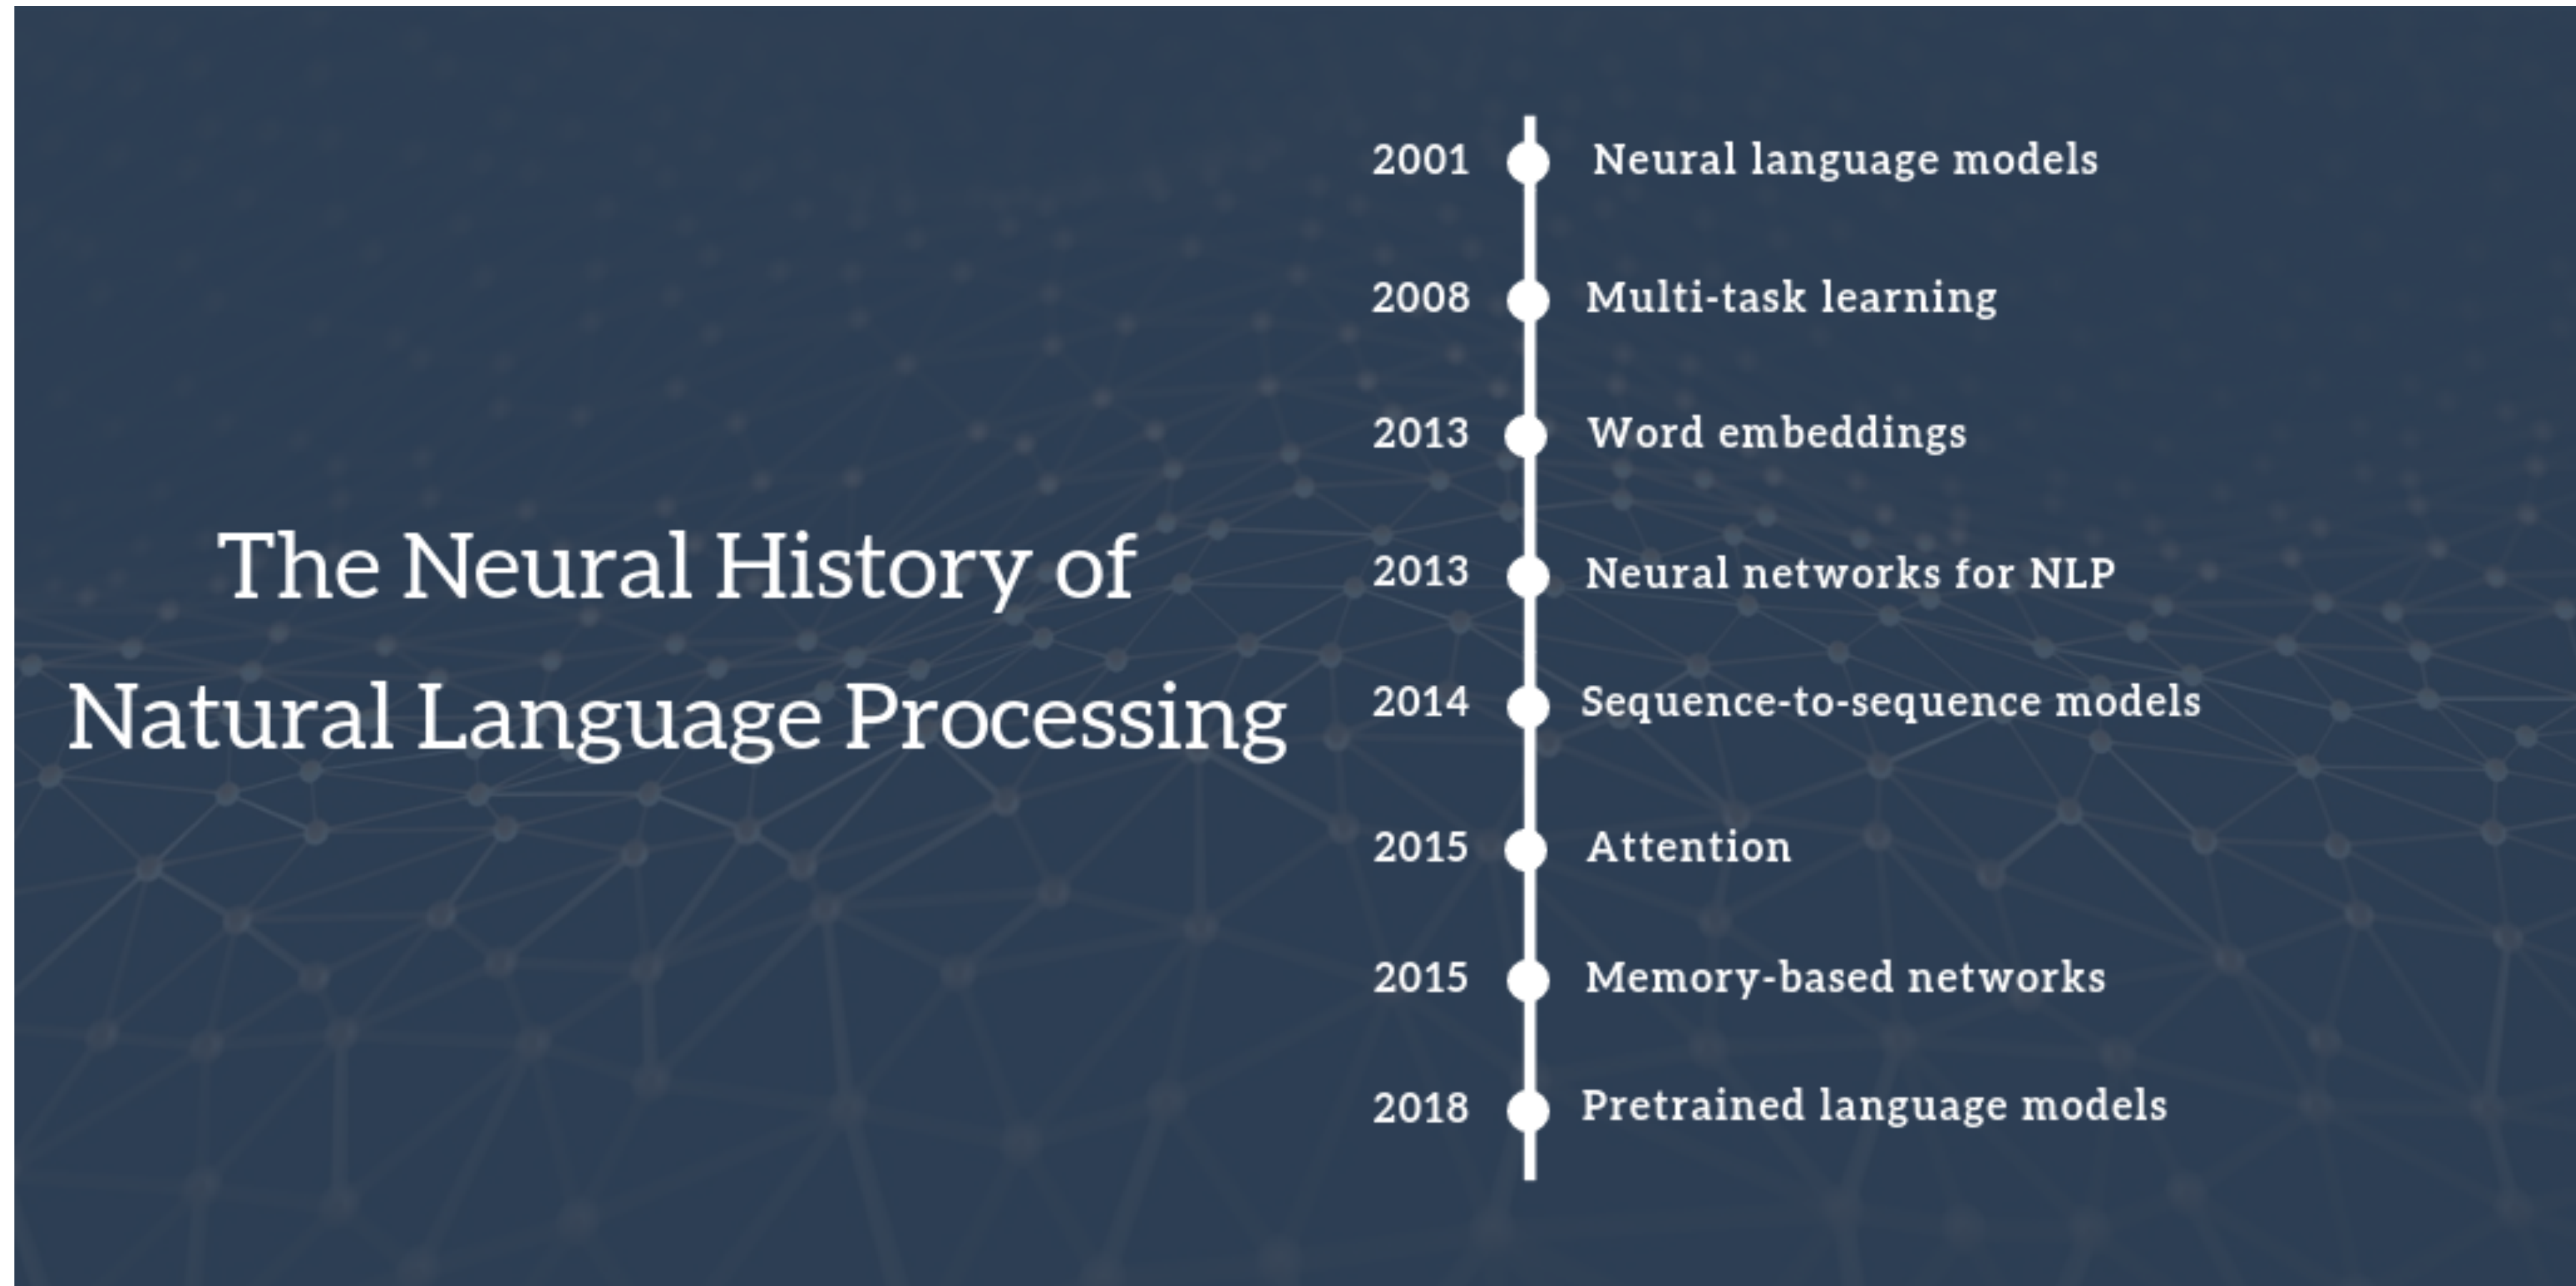
We have been stuck on this one problem – how do proteins fold up – for nearly 50 years. To see DeepMind produce a solution for this, having worked personally on this problem for so long and after so many stops and starts, wondering if we’d ever get there, is a very special moment.

PROFESSOR JOHN MOULT
CO-FOUNDER AND CHAIR OF CASP, UNIVERSITY OF MARYLAND

AlphaFold: a solution to a 50-year-old grand challenge in biology

The neural history of NLP

Deep learning has brought great impacts to NLP




<https://runder.io/a-review-of-the-recent-history-of-nlp/>


OpenAI GPT-3

GPT-3, an autoregressive language model with 175 billion parameters, can "generate news articles which human evaluators have difficulty distinguishing from articles written by humans"

Language Models are Few-Shot Learners				
Tom B. Brown*	Benjamin Mann*	Nick Ryder*	Melanie Subbiah*	
Jared Kaplan†	Prafulla Dhariwal	Arvind Neelakantan	Pranav Shyam	Girish Sastry
Amanda Askell	Sandhini Agarwal	Ariel Herbert-Voss	Gretchen Krueger	Tom Henighan
Rewon Child	Aditya Ramesh	Daniel M. Ziegler	Jeffrey Wu	Clemens Winter
Christopher Hesse	Mark Chen	Eric Sigler	Mateusz Litwin	Scott Gray
Benjamin Chess		Jack Clark	Christopher Berner	
Sam McCandlish	Alec Radford	Ilya Sutskever	Dario Amodei	
		OpenAI		



Automated driving




[Waymo 360° Experience: A Fully Autonomous Driving Journey](#)





*The global **Medical Robotics** market was 3980 million US\$ in 2018 and is expected to 12200 million US\$ by the end of 2025, growing at a CAGR of 15.0% between 2019 and 2025*



Microsoft Xiaoice 微软小冰



Xiaomi Xiaoi 小米小爱



Bilibili 冷鸢youya



Baidu Xiaodu 百度小度

The Era of Artificial Intelligence

Network

Manufacturing

Health

Biology

Finance

Retail

HCI

IoT

Assistant

Transport

Smart City

Customer Service

Risks and Benefits of AI

Benefits of AI

- Increase work efficiency
- Work with high accuracy
- Reduce cost of training and operation
- Accelerate scientific discovery
- “First solve AI, then use AI to solve everything else.” — Demis Hassabis, CEO of DeepMind

Risks of AI

- Lethal autonomous weapons
- Biased decision making
- Safety-critical applications: driving cars, managing water supplies, healthcare...
- Superintelligent: **the gorilla problem**
 - What the gorilla problem means is **loss of control**. Gorillas “created us”; but now we humans have taken control and the gorillas have lost control. We are now creating AI and lose control over it. To make working AI, we have to hand over some control. But: how much control is that? Do we cede so much control to the AI that we turn from captains to passengers of our ship?



61 References

- ◎ History of AI:
 - <http://aitopics.org/misc/brief-history>
- ◎ <https://searchenterpriseai.techtarget.com/definition/Turing-test>

Thanks! Q&A