Manchester Metropolitan University

Artificial Intelligence Principles

6G7V0011 - 1CWK100

Dr. Peng Wang

Email: p.wang@mmu.ac.uk

Department of Computing and Mathematics

Wednesday, Oct. 2nd, 2024

Outline

Outline

What's AI?

Manchester Metropolitan University

Artificial Intelligence (AI) is becoming popular nowadays both in academia and industry. What does it mean to you?

Will you cluster the following as AI?

- A clock/watch
- A plane
- A vehicle
- etc.

What about the following?

- A smart watch
- An auto-pilot aircraft
- An autonomous vehicle
- etc.

Reasons?

The Terminator

Manchester Metropolitan University

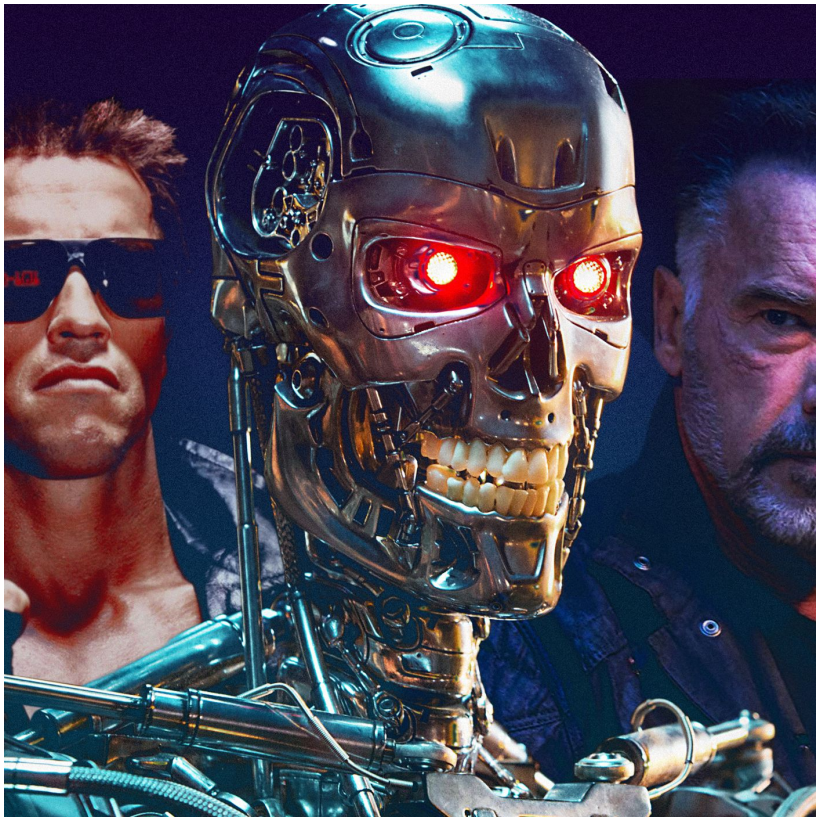

The 1984 science fiction action film The Terminator shows concern that robots(cyborgs) may harm/kill us!

WALL-E

Manchester Metropolitan University

The 2008 film WALL-E shows a robot cleans the ruins/garbage in a inhabitable and deserted earth, and ...

WALL-E

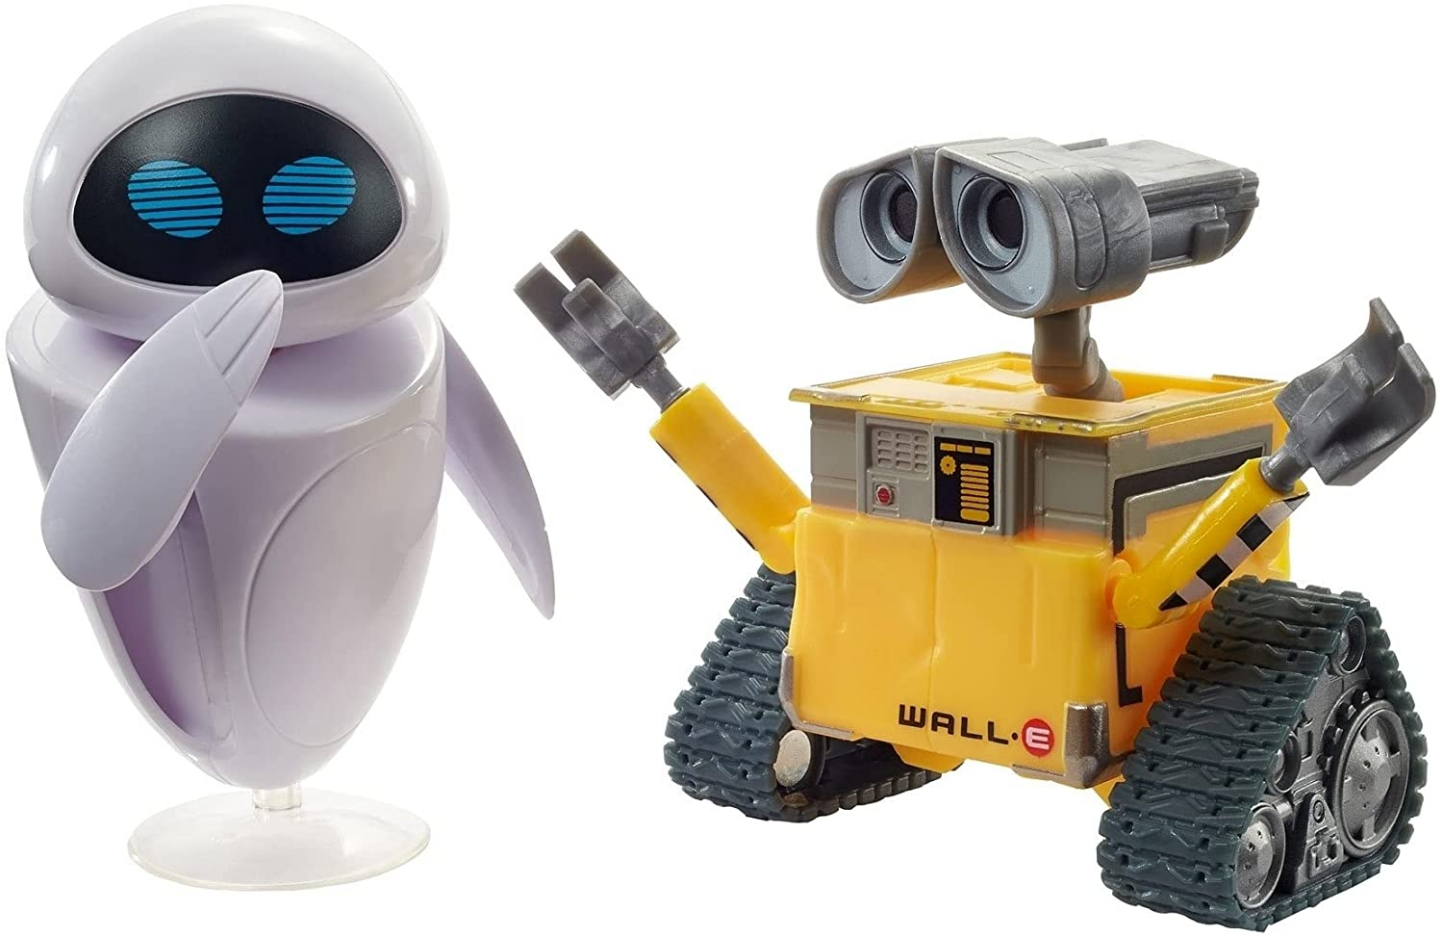

got a mate!

This romantic science fiction film shows the opposite. Robots may help us, and may have emotional needs just like humans.

More Films on AI

(a) Cyborg, 1989

(b) Bicentennial Man, 1999

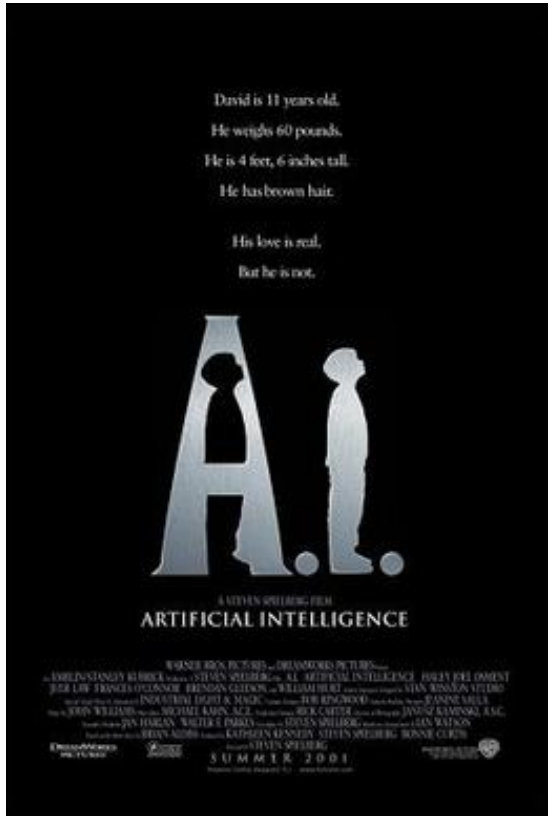

(c) Artificial Intelligence, 2001

Early days' film tend to 'consider' AI as threats, while nowadays' film tend to think 'coexistence' with AI.

Note: This is my personal view!

Autonomous Vehicles - Waymo

Waymo, Google's autonomous vehicle takes us to our destination.

Also refer to this link for details.

ADAS - Tesla

Manchester Metropolitan University

ADAS - Advanced Driver-Assistance Systems

Tesla is ambitious in autonomous vehicle industry as well. More sales, up to June 2021, over 1 million Model 3 sold!

Delivery UGVs

Manchester Metropolitan University

UGV - Unmanned Ground Vehicle

Amazon's Scout

Delivery UAVs

Manchester Metropolitan University

UAV - Unmanned Aerial Vehicle

Drones are not only used in shooting TikTok videos, but also delivering stuff!

Ethical Issues

Manchester Metropolitan University

Drones and autonomous vehicles are used in wars... AI is not only about technology, we need to think about ethical concerns as well!

What Do They Have in Common?

Manchester Metropolitan University

What Do They Have in Common?

Manchester Metropolitan University

- They are designed for specific purposes
  1. Their performance is measured/assessed by matching certain criteria.
- They operate in specific environments
  1. They know completely/partially the environments
  2. Built-in knowledge or a priori knowledge of the environment
- They are equipped with sensors, e.g. cameras, LiDAR for perceiving the environment
- They perform actions to change/affect the environments

Anything else?

Intelligence vs Artificial Intelligence (AI)

Manchester Metropolitan University

What's Intelligence?

Dimension 1: *Is it about thought or behavior?*

1. Some consider intelligence to be a property of *internal thought processes* and *reasoning*
2. Some focus on intelligent *behavior*, an external characterisation

Dimension 2: *Is it about fidelity to human or rationality?*

1. In terms of fidelity to *human* performance
2. Intelligence is abstractly, formally defined as *rationality* – loosely speaking, doing the 'right thing'

Table 1: Two dimensional definition of AI

	Human	Rational
Thought	Thinking humanly	Thinking rationally
Behavior	Acting humanly	Acting rationally

Note: The two dimensional definition of AI also implying the four ways we approach AI.

Acting like a Human

1. **Acting humanly:** The Turing test approach

Turing Test

"Designed by Alan Turing in 1950 to sidestep the philosophical vagueness of the question '*Can a machine think?*' "

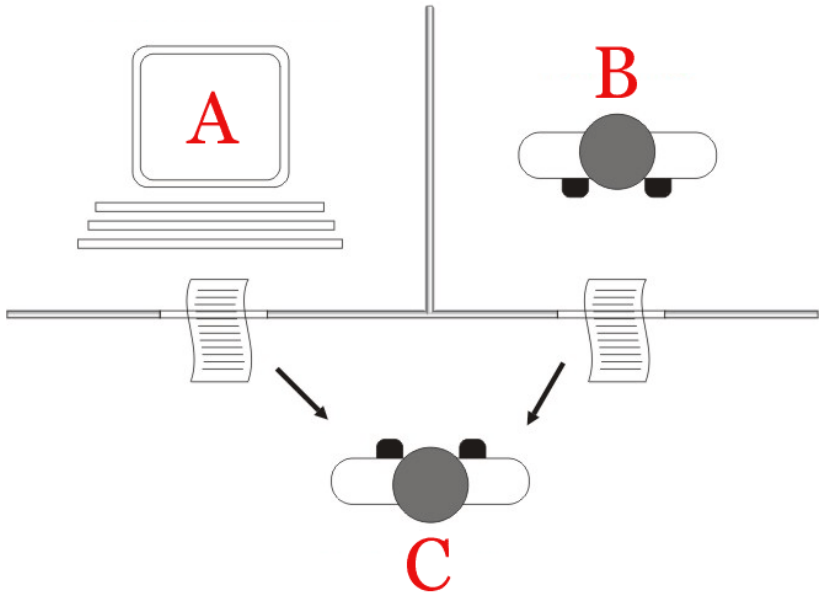

Figure 1: An example of the Turing Test. Figure from Wikipedia.

Acting like a Human

Figure 2: A step further on the Turing test. Figure from Image Credit.

Think like a Human

Manchester Metropolitan University

2. **Thinking humanly:** The cognitive modeling approach

- To know the **theory** of mind?
  1. Introspection - catch our own thoughts as they go by
  2. Psychological experiments - observing a person in action
  3. Brain imaging - observing brain in action
- Express the **theory** as a computer program
  1. Input-output behavior of the computer program matches corresponding human behavior
  2. Newell and Simon developed the mathematical theorem-proving system called the Logic Theorist (LT), which thinks like a human while solving mathematical theorem-proving problems.
  3. Later on, they developed the General Problem Solver (GPS) in 1961 - they were concerned with compare the sequence and timing of its reasoning steps to those of human subjects solving the same problem.

'Laws of Thought Approach'

Manchester Metropolitan University

3. **Thinking rationally:** The 'laws of thought' approach

- **Logic** to govern the operation of the mind
- Aristotle's **syllogisms**
  1. Correct inference: correct conclusion given correct premises
  2. Example: 'Socrates is a man; all men are mortal; therefore, Socrates is mortal'
- A precise notation for statements about objects in the world and the relations among them was developed in the 19th century
- By 1965, programs could solve any solvable problem described in logical notation
- Requires knowledge of the world to be certain - seldom achieved in reality

'Laws of Thought Approach'

Manchester Metropolitan University

3. **Thinking rationally:** The 'laws of thought' approach

- Obstacles
  1. Lack the ability to represent uncertain knowledge, e.g., 'It is cloudy, is it going to rain?'
  2. The gap between solving a problem 'in principle' and 'in practice', e.g. matrix inversion when dimension increases.
- Solutions
  1. Probability fills the gap between logic and uncertainty. makes rigorous reasoning with uncertain information do-able.
  2. So far, 'think rationally' is possible. Lack the ability of 'rational action' generation.

Rational Agent Approach

Manchester Metropolitan University

4. **Acting rationally:** The rational agent approach

Agent

"An **agent** is anything that can be viewed as perceiving its **environment** through **sensors** and acting upon that environment through **actuators**."

Rational Agent

"An **rational agent** is an agent that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome."

Note: More details will be given in the next week!

Rational Agent Approach

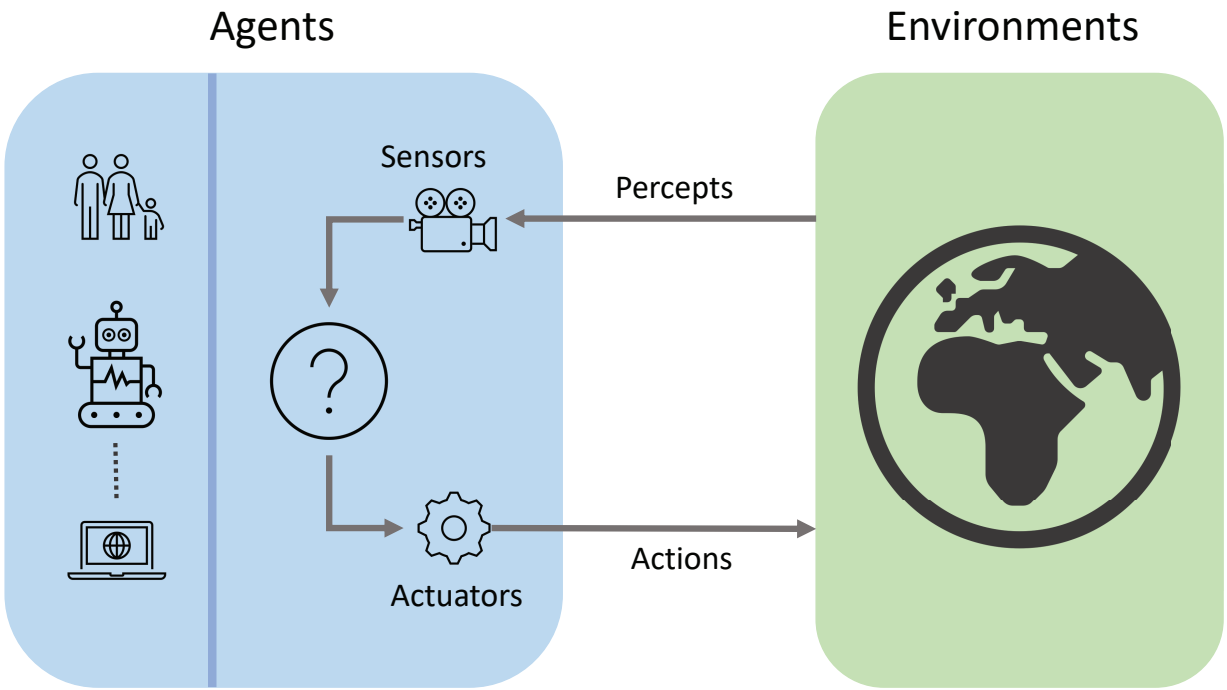

Figure 3: Agents interact with environments.

Refer to this videos to learn more!

Rational Agent Approach

Manchester Metropolitan University

Examples:

- A **human agent** has eyes, ears, and other organs for sensors and hands, legs, vocal tract, and so on for actuators.
- A **robotic agent** might have cameras and infrared range finders for sensors and various motors for actuators.
- A **software agent** receive file contents, network packets, and human input (keyboard/mouse/touchscreen/voice) as sensory inputs and acts on the environment by writing files, sending network packets, and displaying information or generating sounds.

Summary

Manchester Metropolitan University

We have learned:

- The definition of AI

1. 'The pursuit of human-like intelligence must be in part an empirical science related to psychology, involving observations and hypotheses about actual human behavior and thought processes'
2. 'A rationalist approach, on the other hand, involves a combination of mathematics and engineering, and connects to statistics, control theory, and economics.'

A brief history of AI

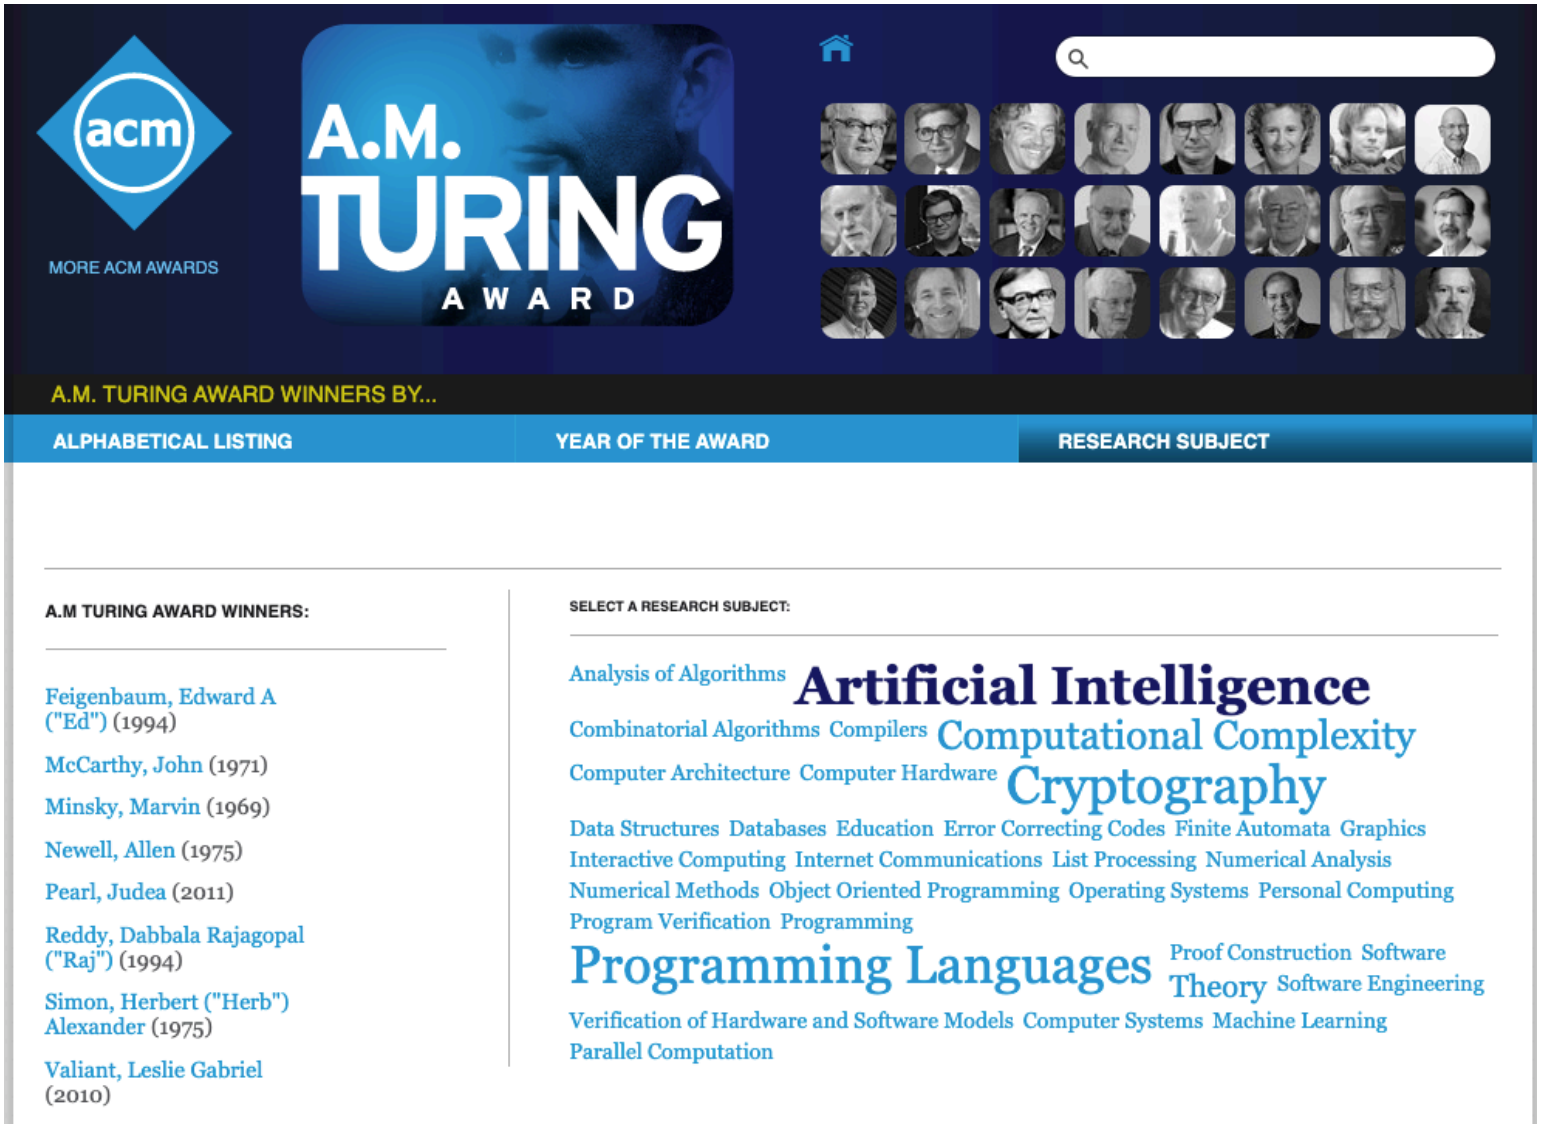

Figure 4: Turing trophy winners in AI.

A brief history of AI

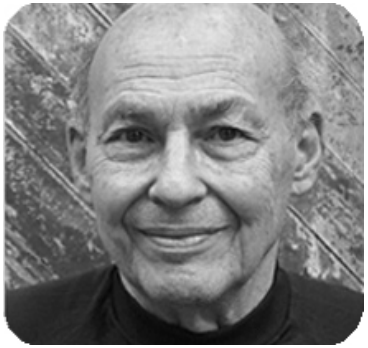

(a) Marvin Minsky

(b) John McCarthy

(c) Ed Feigenbaum

(d) Raj Reddy

(e) Judea Pearl

(f) Geoffrey E Hinton

(g) Yoshua Bengio

(h) Yann LeCun

People that created and witnessed the pivotal moments of AI!

A brief history of AI

Manchester Metropolitan University

- Marvin Minsky (1969) and John McCarthy (1971)
  - Foundation of AI based on representation and reasoning
- Ed Feigenbaum and Raj Reddy (1994)
  - Enconding human knowledge with expert systems to solve real-world problems.
- Judea Pearl (2011)
  - Probabilistic reasoning techniques to deal with uncertainty in a principled manner.
  - Bayesian Network.
- Yoshua Bengio, Geoffrey E Hinton, and Yann LeCun (2011)
  - Multilayer neural networks (Deep learning).

A brief history of AI

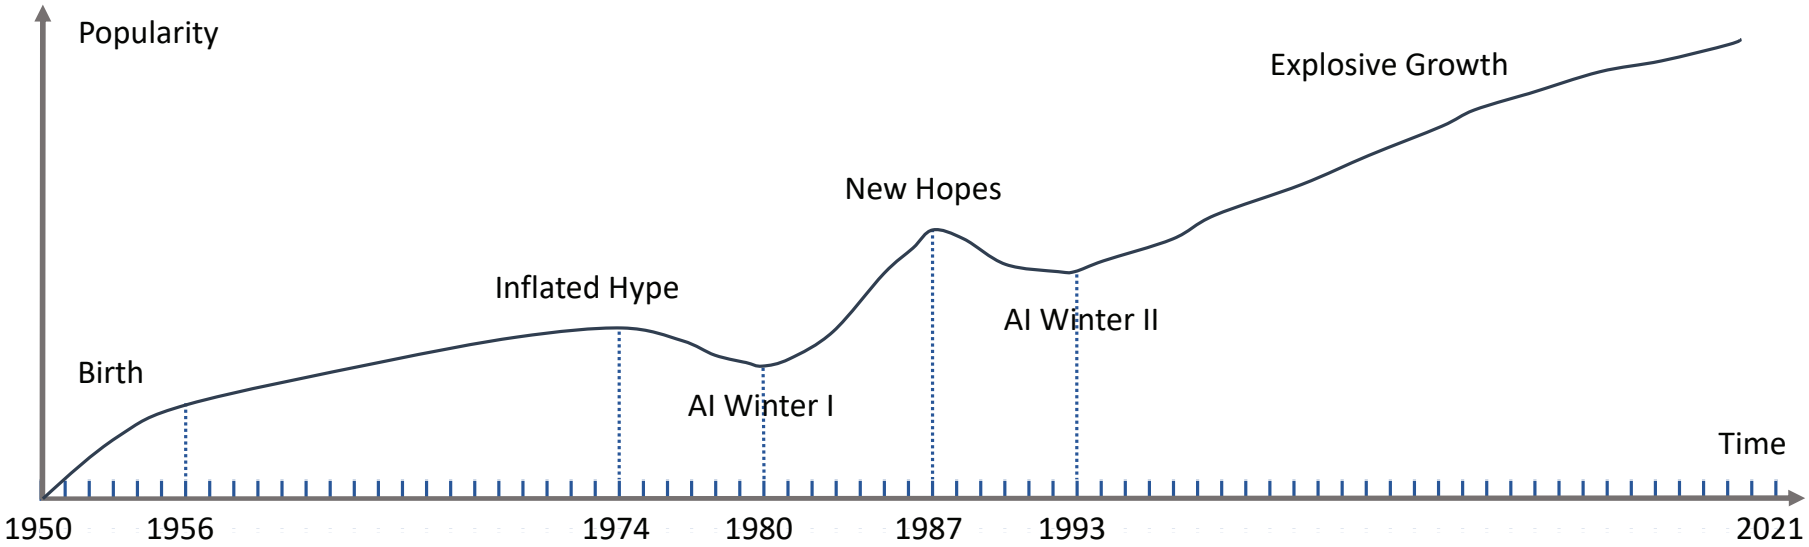

- **1950s-1960s:** First AI boom – the age of reasoning, prototype AI developed, including the perceptron and the first neural network.
- **1980s-1990s:** Second AI boom: the age of knowledge representation (appearance of expert systems capable of reproducing human decision-making)
- **1997:** Deep Blue beats Gary Kasparov
- **2006:** University of Toronto develops Deep Learning
- **2011:** IBM's Watson won Jeopardy
- **2016:** Go software based on Deep Learning beats world's champions

- 1970s: AI winter I
  - Limited applicability of AI leads to funding pullback.
  - 1969: Marvin Minsky and Seymour Papert's perceptron failed to live up to expectations.
  - 19774: The Lighthill report - promise in AI wasn't delivered.
- 1990s: AI winter II
  - Limitations of if-then reasoning become more apparent.
  - 1987: DARPA cut funding for AI research
  - 1990: Expert systems' if-then rules failed to process novel knowledge.

Figure 5: Landscape of AI history.

Refer to this Figure 5 while watching this video!

Read and Discussion

Manchester Metropolitan University

Turing Awards laureates' opinion, see the video!

1. Read 'A brief history of Artificial Intelligence.'
2. Read 'Will there be another AI winter? why and why not?'.

Will there be another AI winter?

An open discussion on whether there will be another AI winter? If so, why?

Outline

AI Tree

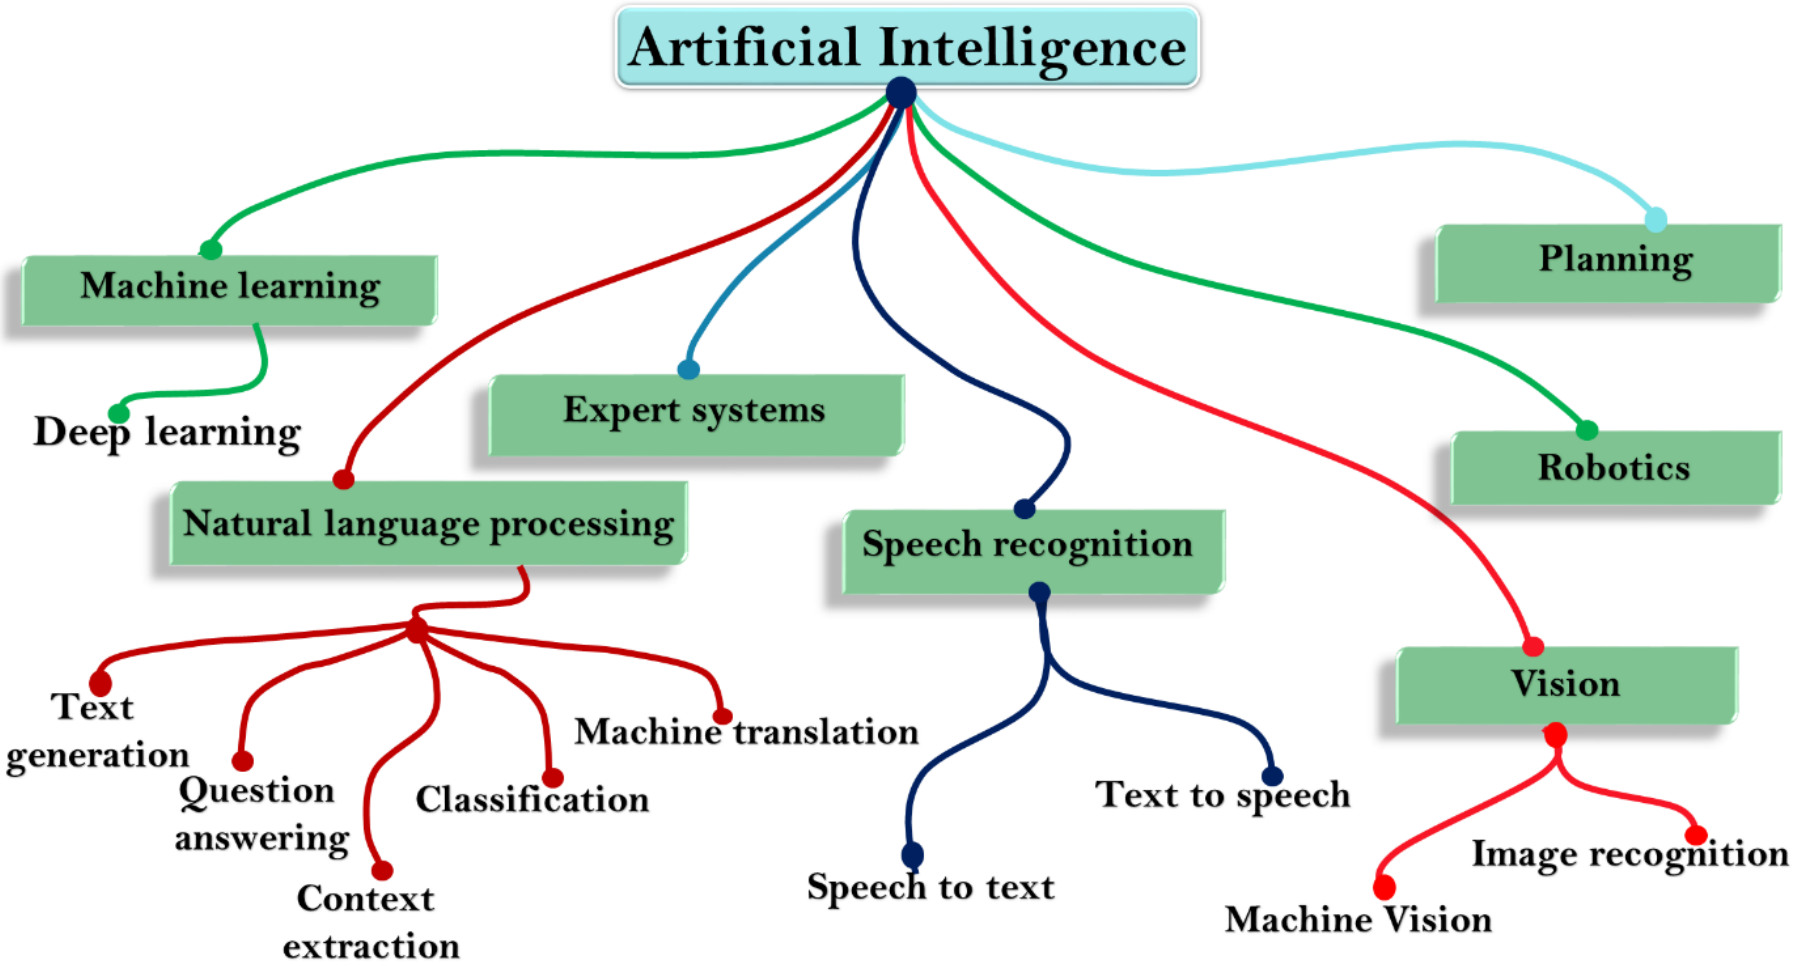

Figure 6: Subareas of AI, Figure Credit

Machine Learning ≠ AI

- The term 'Machine Learning' was coined by Arthur Samuel in 1959.
  - Samuel Checker-playing program
- Common definition by Tom Mitchell
  - **Machine Learning is the study of computer algorithms that improve automatically through experience.**
- Deep Learning ∈ Machine Learning, the most popular sub-field of AI.

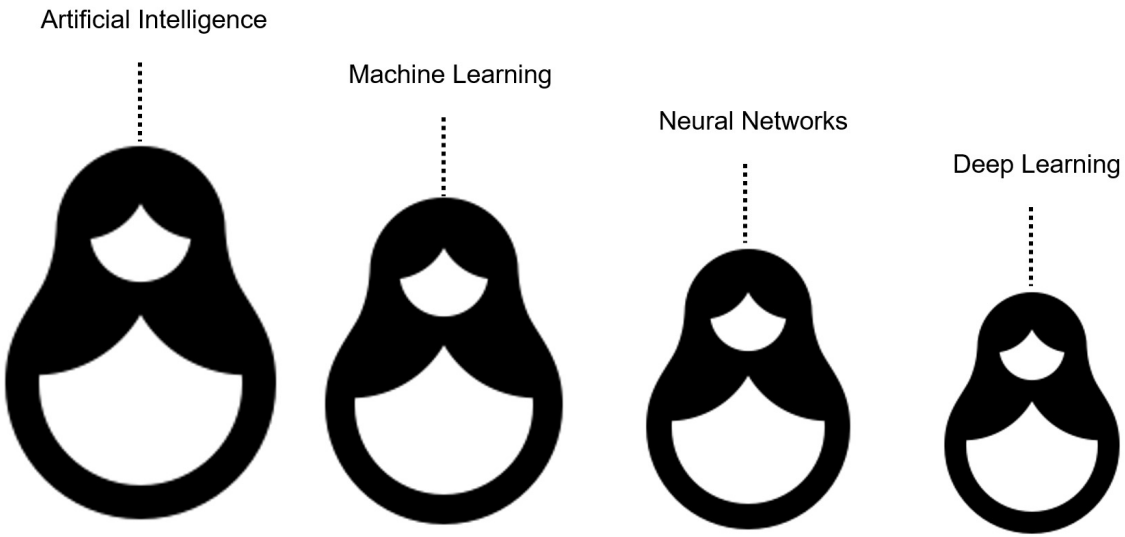

Figure 7: Hierarchy of various concepts, Figure Credit

AI Players - Industrial

Manchester Metropolitan University

(a) Apple

(b) Amazon

(c) facebook

(d) Microsoft

(e) Google

(f) DJI

(g) Tencent

(h) pony.ai

(i) sensetime

(j) intel

(k) Nvidia

(l) qualcomm

(m) OpenAI

(n) DeepMind

(o) ARM

AI Players - Academia

(p) Oxford Uni.

UNIVERSITY OF CAMBRIDGE

(q) Cambridge Uni.

(r) Stanford Uni.

Massachusetts Institute of Technology

(s) MIT

- Artificial Intelligence and Machine learning: NeurIPS, ICML, ICLR, IJCAI, AAAI, etc.
- Computer vision: CVPR, ECCV, ICCV, etc.
- Robotics: ICRA, IROS, etc.

1. Industrial players with access to more data
2. Computer power and skillful engineers
3. Staff exchange between industry and academic

Stats of NeurIPS 2020 I

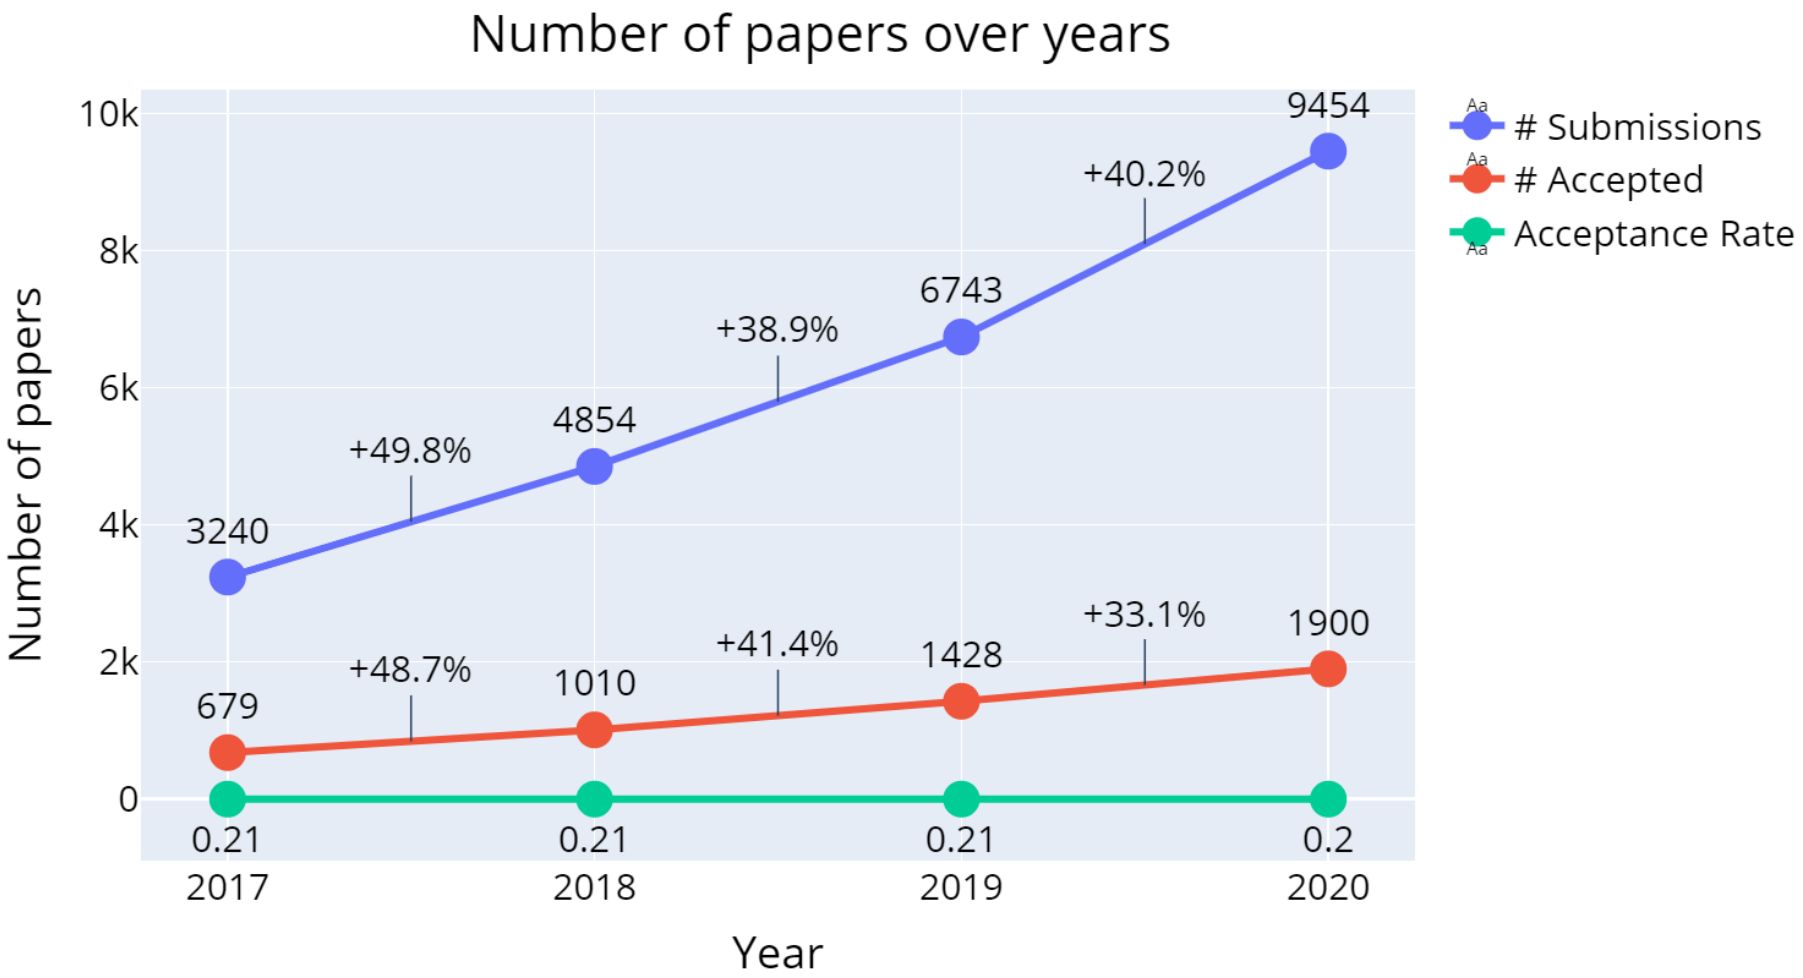

(t) NeurIPS 2020 Figure Credit

Stats of NeurIPS 2020 II

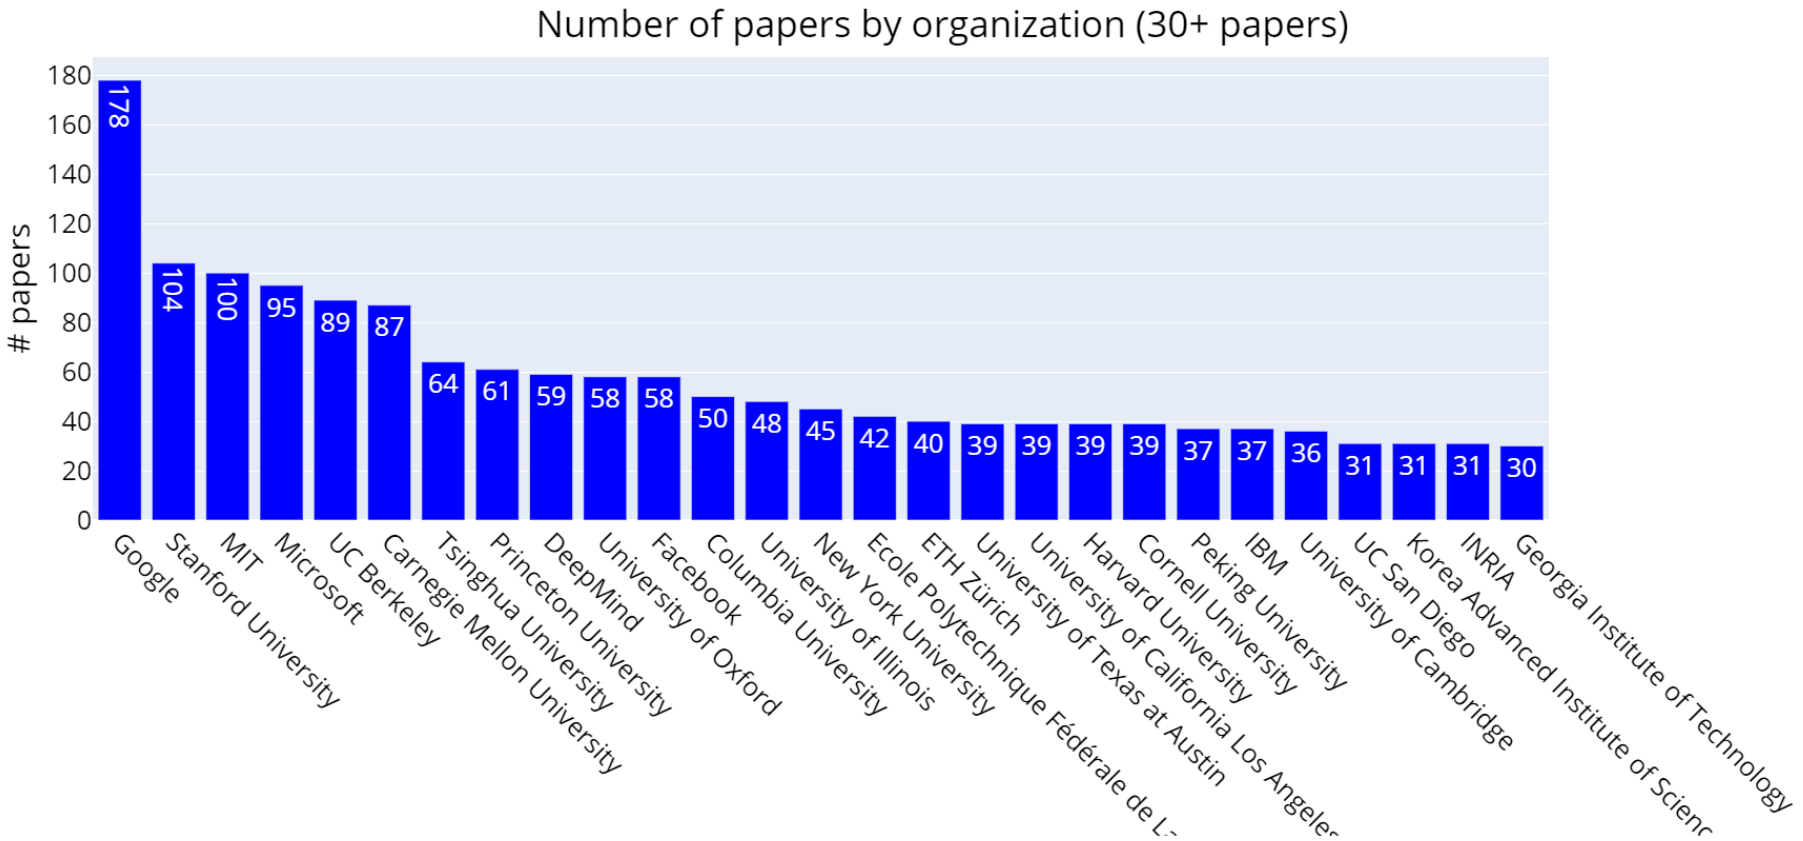

(u) NeurIPS 2020 Figure Credit

Stats of NeurIPS 2020 III

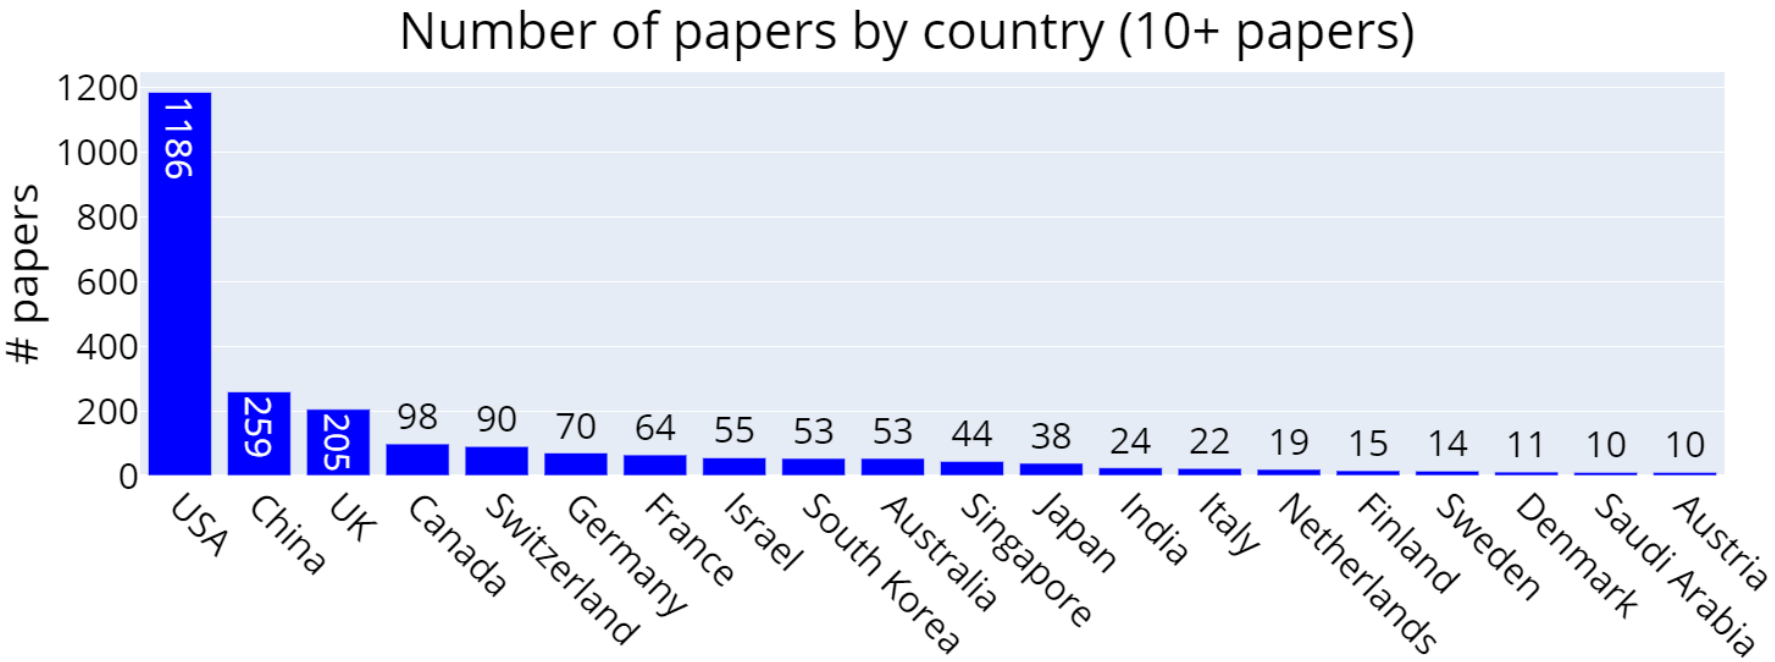

(v) NeurIPS 2020 Figure Credit

Stats of NeurIPS 2020 IV

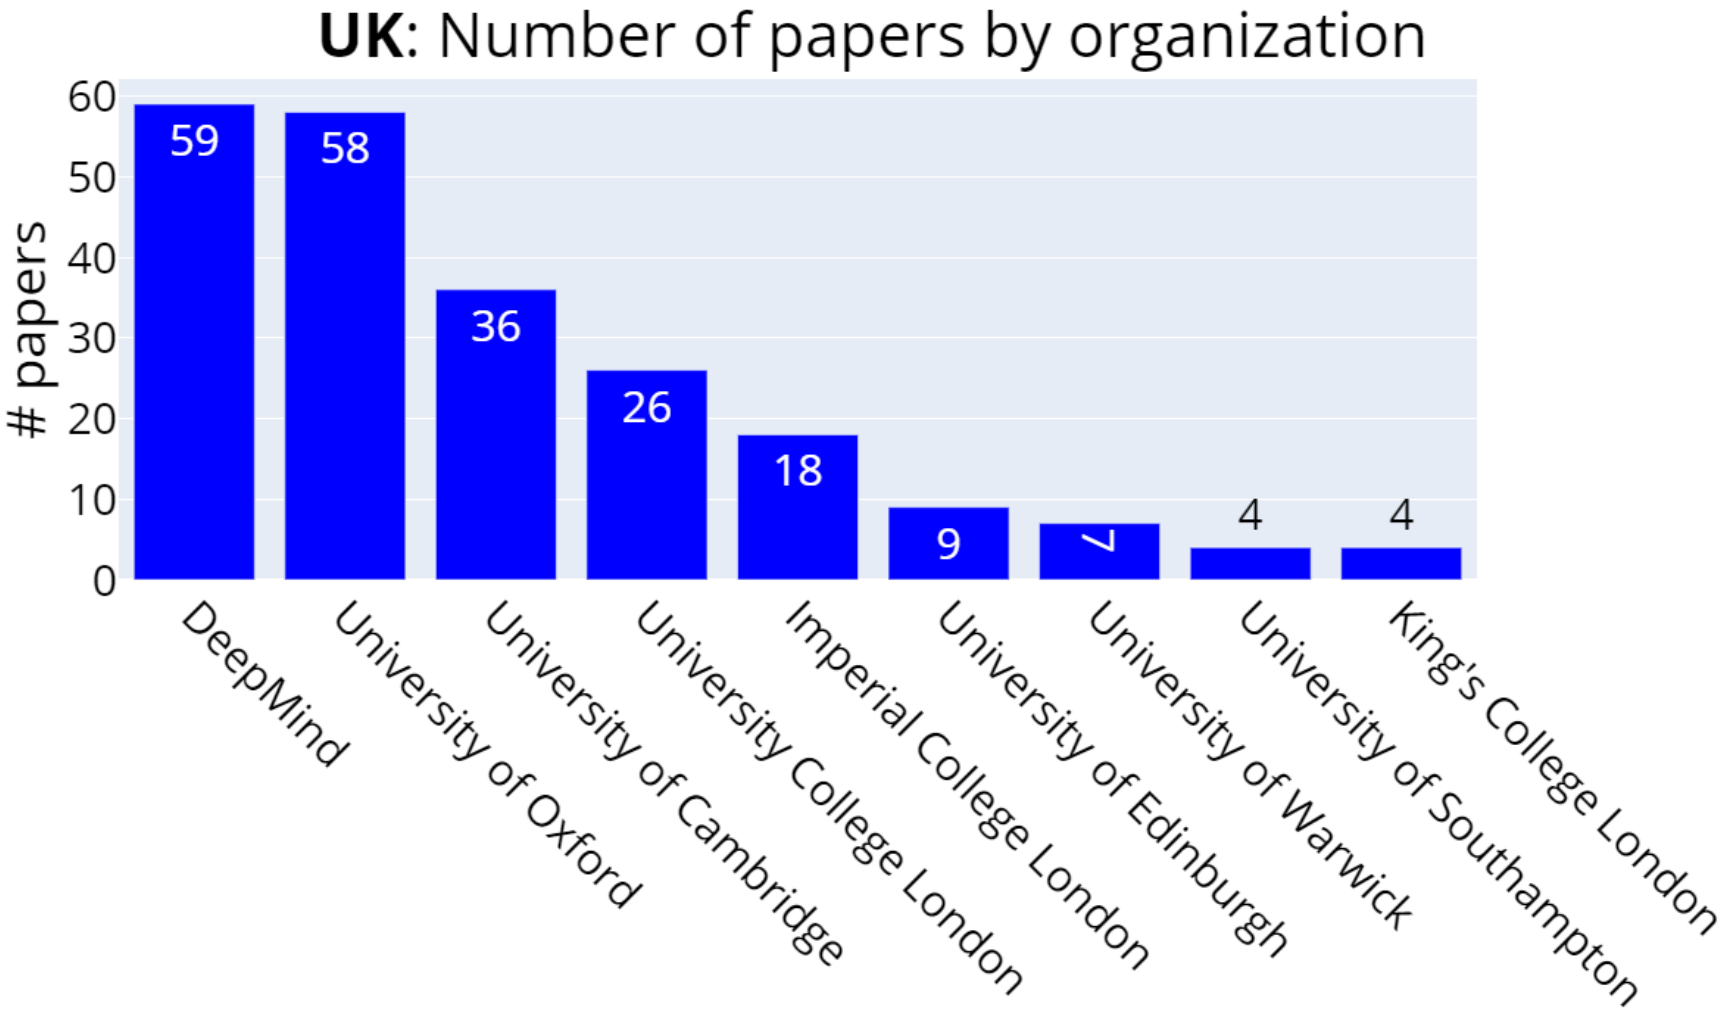

(w) NeurIPS 2020 Figure Credit

Outline

State-of-the-Art AI

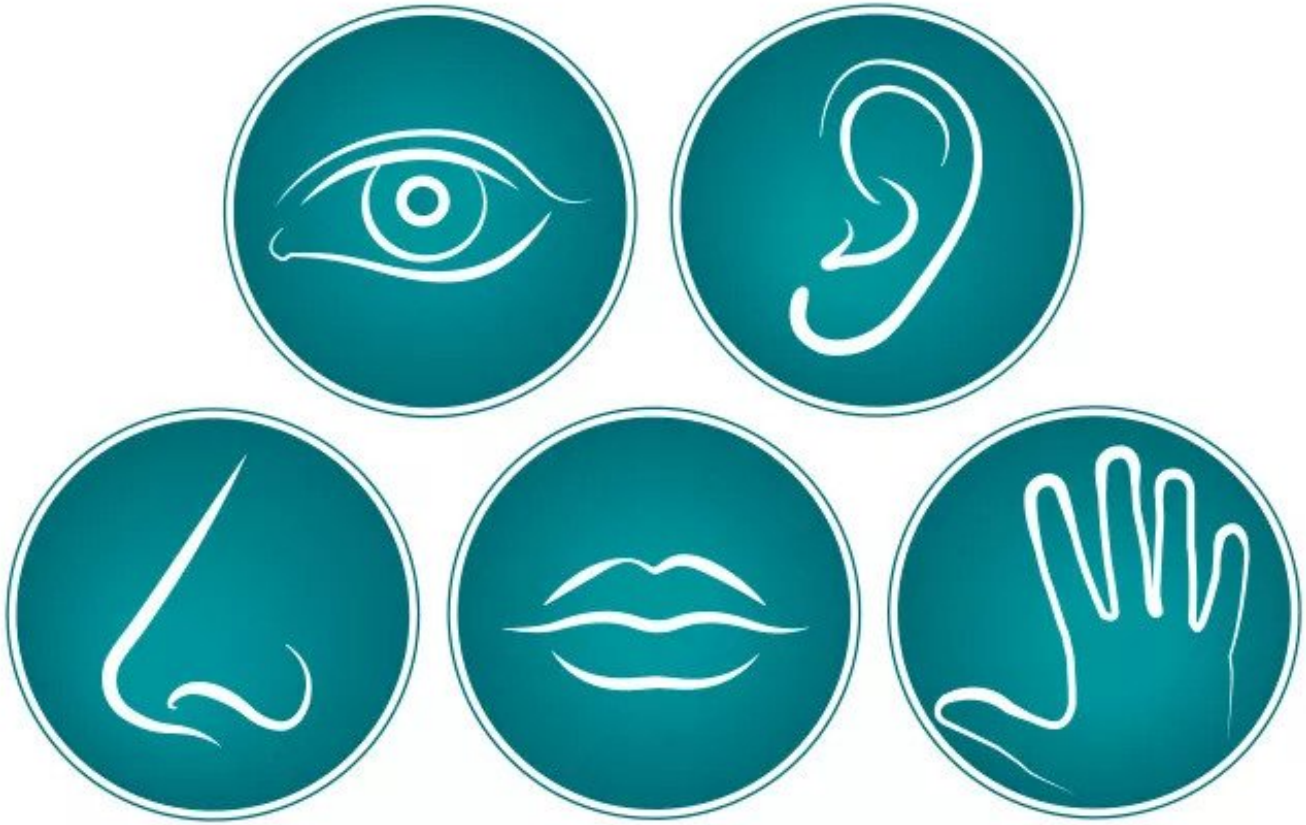

Figure 8: Basic Human Sense, Figure Credit

Basic sense: Sight (Computer Vision), Hearing (Audio/Speech/NLP), Touch (HCI/HRI), Smell, Taste.

Computer Vision

Figure 9: ImageNet: 14,197,122 images, 21841 synsets indexed.

Figure 10: Li Fei-Fei, creator of ImageNet

Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. "Imagenet: A large-scale hierarchical image database." In 2009 IEEE conference on computer vision and pattern recognition, pp. 248-255. IEEE, 2009.

- The ILSVRC evaluates algorithms for object detection and image classification at large scale.
- To allow researchers to compare progress in detection across a wider variety of objects – taking advantage of the quite expensive labeling effort.
- To measure the progress of computer vision for large scale image indexing for retrieval and annotation.

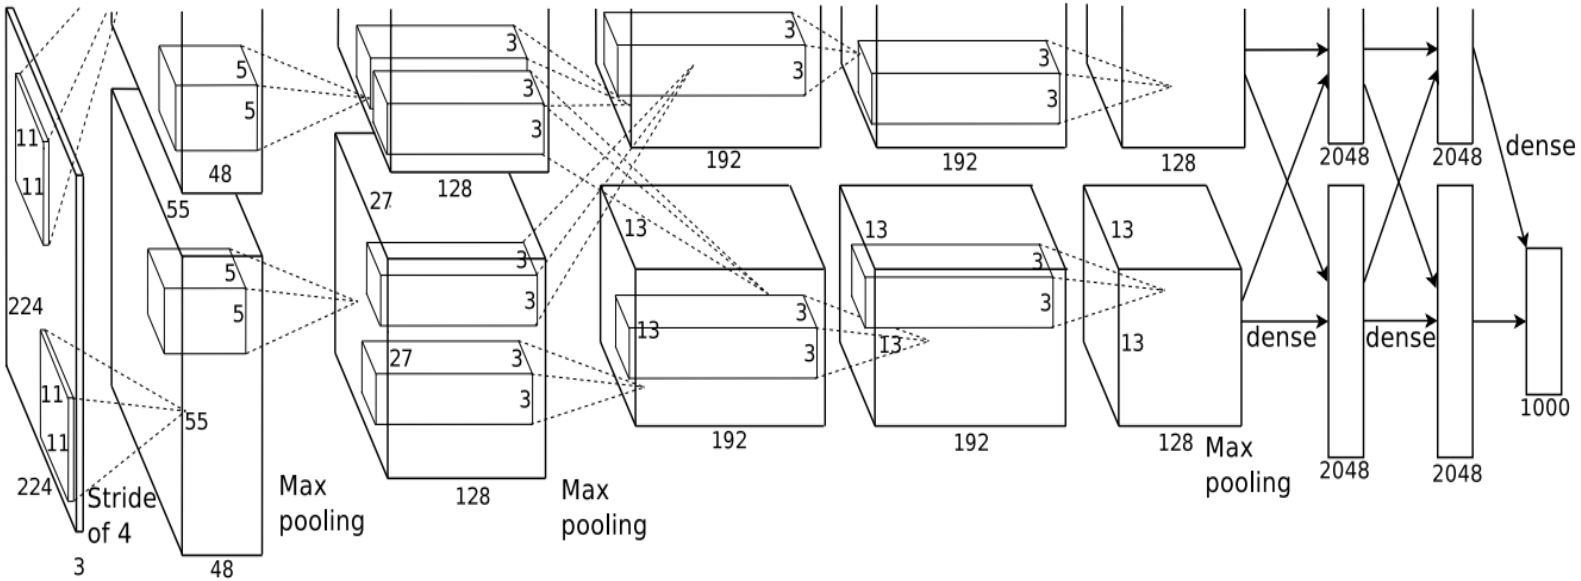

Figure 11: ALEXNET: ILSVRC 2012 winner, Credit

ILSVRC - What they do?

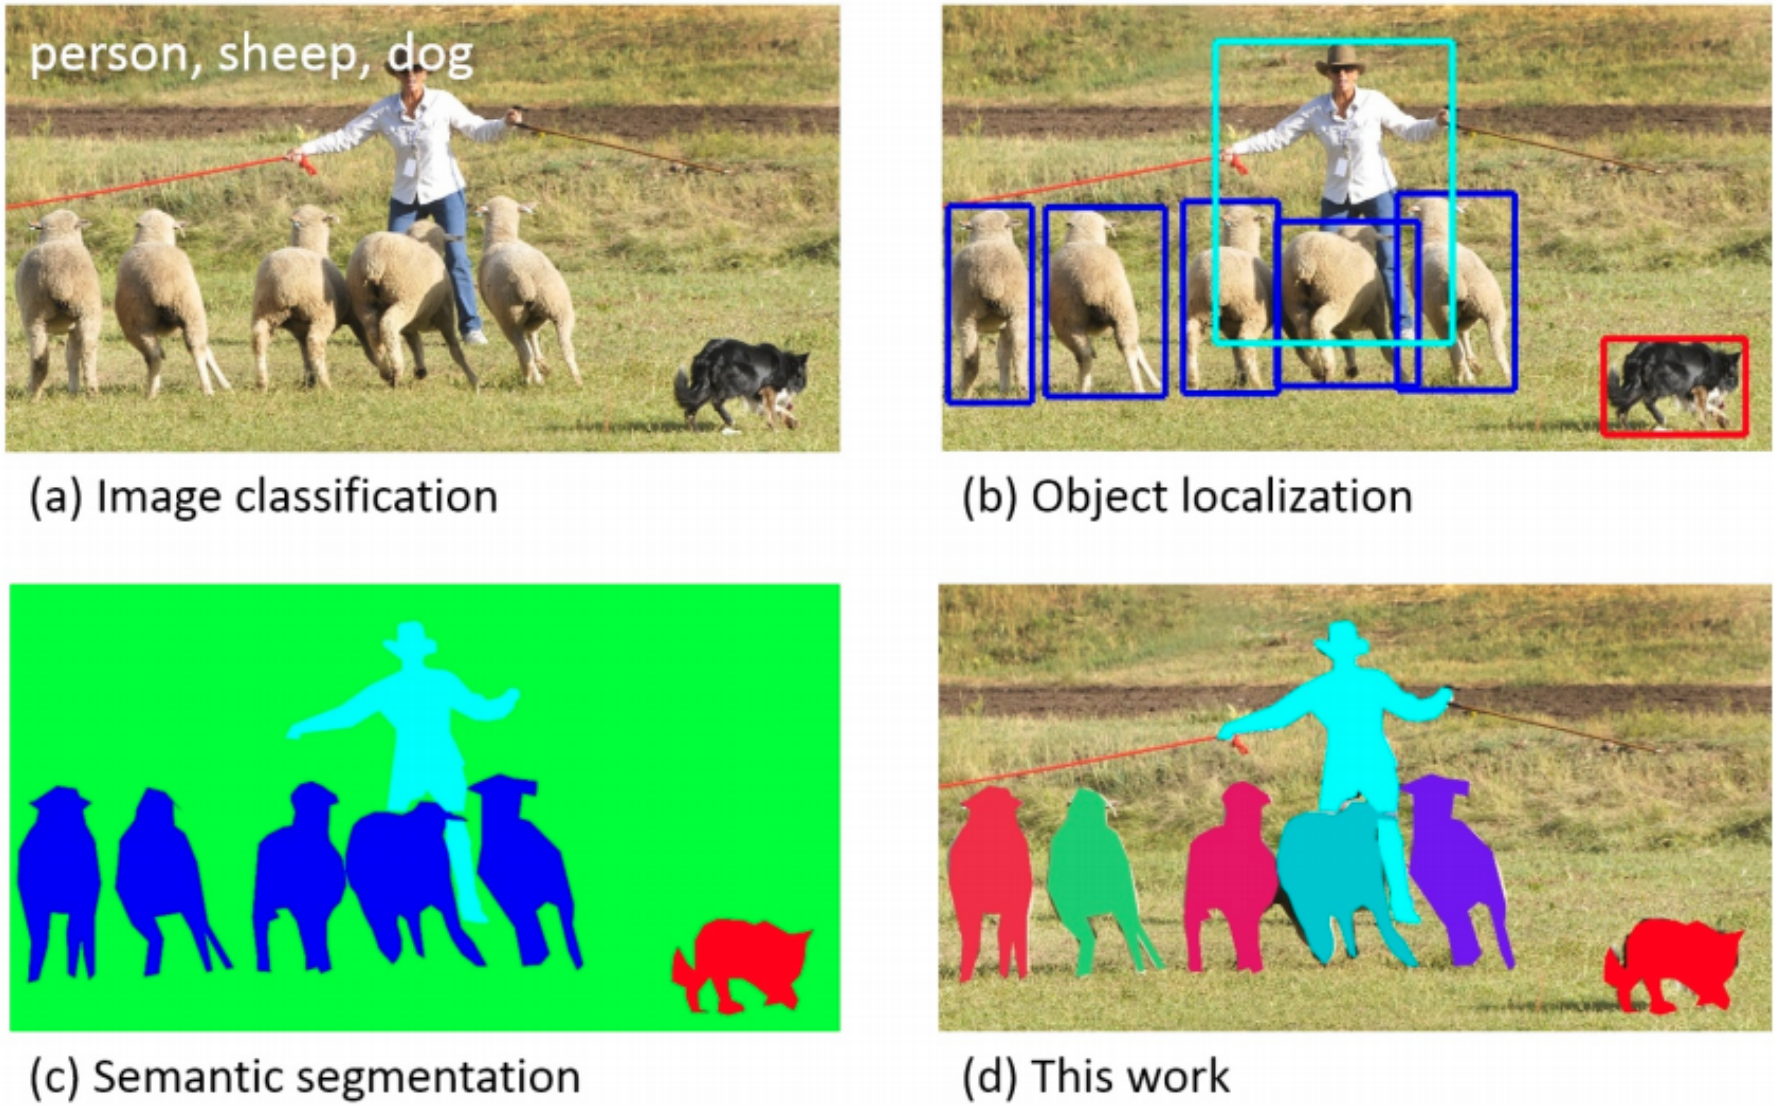

Figure 12: Examples of what ILSVRC does?

ILSVRC Winners

Manchester Metropolitan University

Model	Top-1 (val)	Top-5 (val)	Top-5 (test)
SIFT + FVs [7]	—	—	*26.2%*
1 CNN	40.7%	18.2%	—
5 CNNs	38.1%	16.4%	**16.4%**
1 CNN*	39.0%	16.6%	—
7 CNNs*	36.7%	15.4%	**15.3%**

Figure 13: Deep learning (CNN) outperforms handcrafted features, Credit

Year	Winner	Error rate
2012	ALEXNET	15.3
2013	ZFNET	11.2
2014	GoogLeNet	6.67
	VGG NET	6.67
2015	ResNet	3.57
2016	ResNeXt	4.1
2017	SENet	2.251
2018	PNASNet-5	3.8

Table 2: Winners of ILSVRC

Computer vision nowadays

Computer vision nowadays

Manchester Metropolitan University

What computer vision can do up to now?

Figure 14: fackbook FAIR's amazing work

1. Object detection
2. Segmentation
3. Pose tracking
4. etc.

Computer vision beyond I

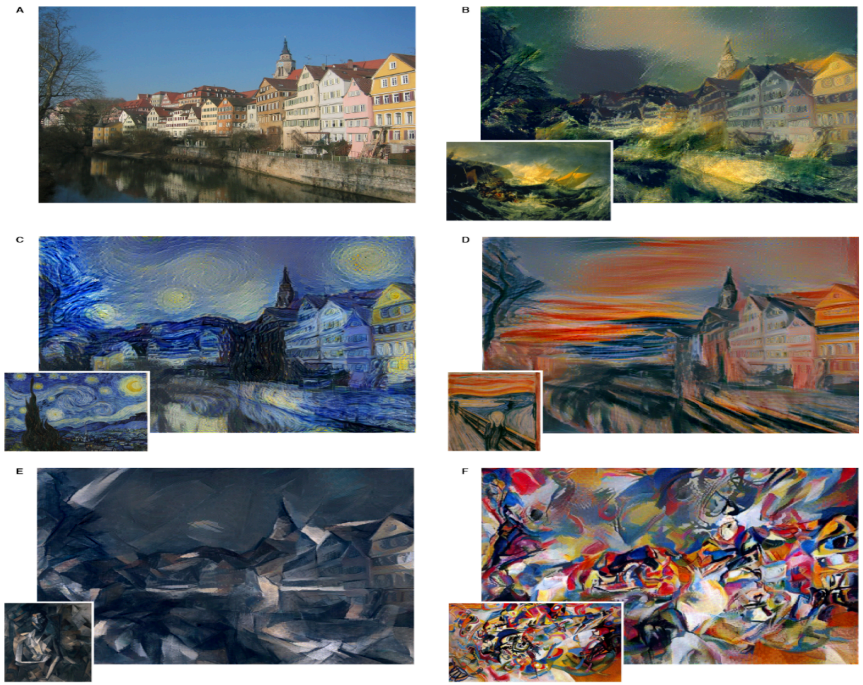

Figure 15: Style Transfer

Figure 16: Image colorisation

Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "A neural algorithm of artistic style." arXiv preprint arXiv:1508.06576 (2015).

Zhang, Richard, Phillip Isola, and Alexei A. Efros. "Colorful image colorization." In European conference on computer vision, pp. 649-666. Springer, Cham, 2016.

Computer vision beyond II

Figure 17: Image reconstruction

Liu, Guilin, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. "Image inpainting for irregular holes using partial convolutions." In Proceedings of the European Conference on Computer Vision (ECCV), pp. 85-100. 2018.

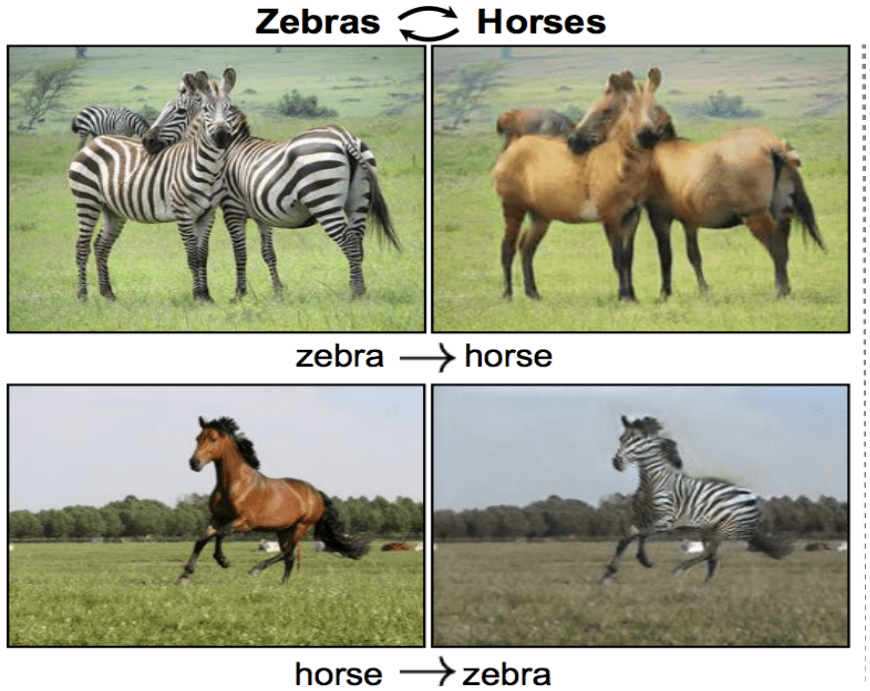

Figure 18: Image synthesis

Zhu, Jun-Yan, Taesung Park, Phillip Isola, and Alexei A. Efros. "Unpaired image-to-image translation using cycle-consistent adversarial networks." In Proceedings of the IEEE international conference on computer vision, pp. 2223-2232. 2017.

Robotics

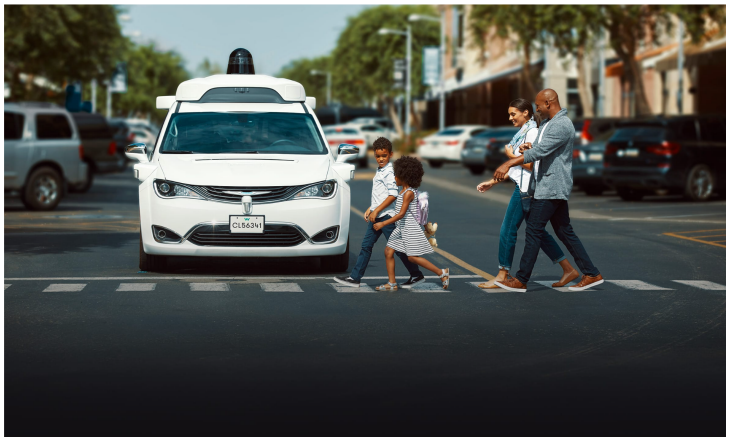

(a) Waymo

(b) DJI FPV

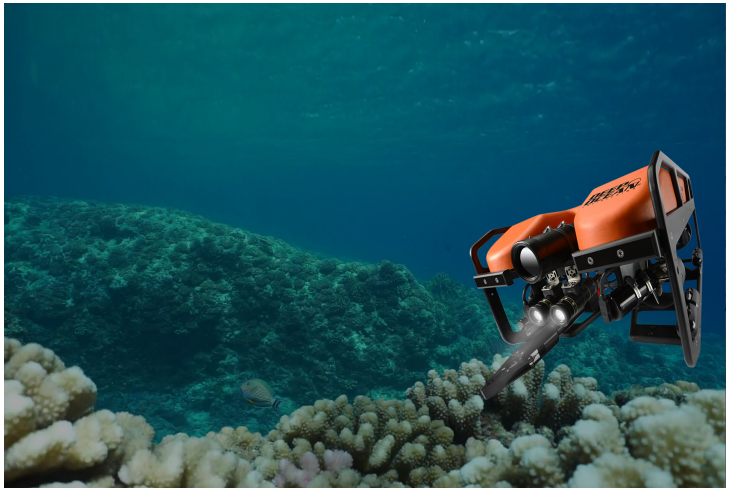

(c) Deep Ocean Engineering

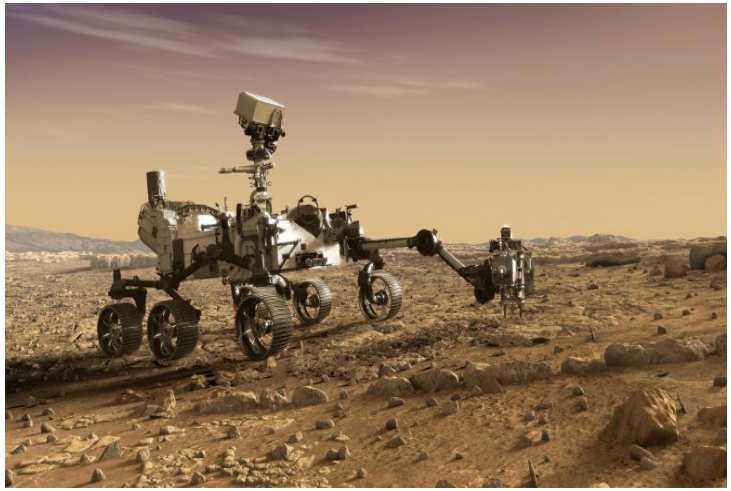

(d) NASA Mars Rover

Robotics - Beyond

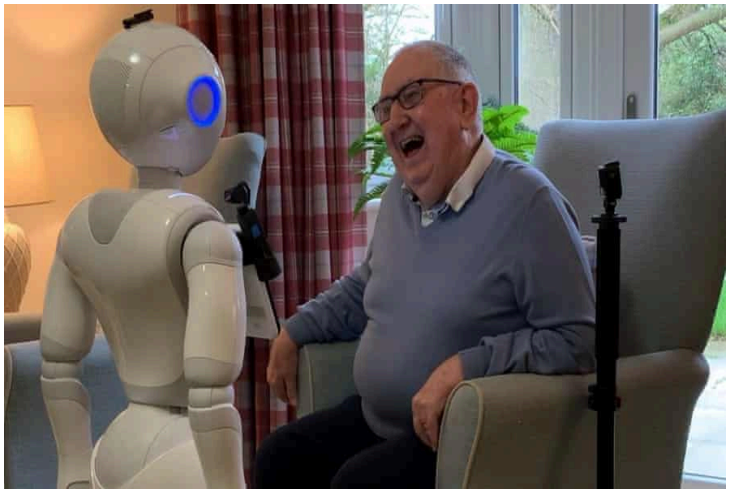

(e) Robot Carer

(f) Cobot

(g) Robotic surgery

(h) Assitants

RoboCup - Fancy competitions?

Manchester Metropolitan University

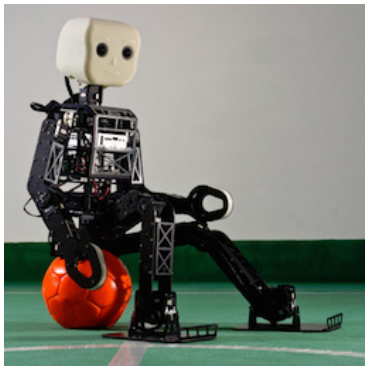

(i) Kidsize humanoid league

(j) Standard platform league

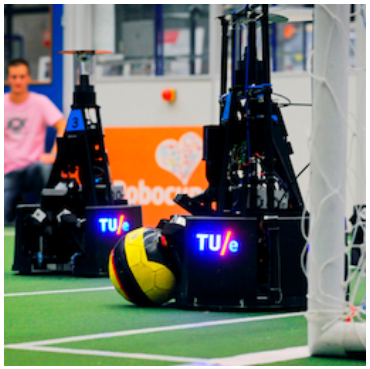

(k) Middle size league

(l) Rescue robot league

(m) Humanoid league

(n) Junior league

(o) Small size league

IEEE Robotics & Automation Society

(p) Organization

Outline

Natural Language Processing

Goal: Deep Understanding

- Requires context, linguistic structure, meanings, etc.
- Different languages/dialects, etc.

Reality: Shallow Matching

- Requires robustness and scale
- Amazing successes, but fundamental limitations

Note: Credits to Dan Klein - UC Berkeley

Speech/Audio Technologies

- Automatic speech recognition (ASR)
- Text-to-speech synthesis (TTS)
- Dialog systems

Language processing technologies

- Question answering
- Machine translation
- Text classification/understanding

Speech/Audio Technologies

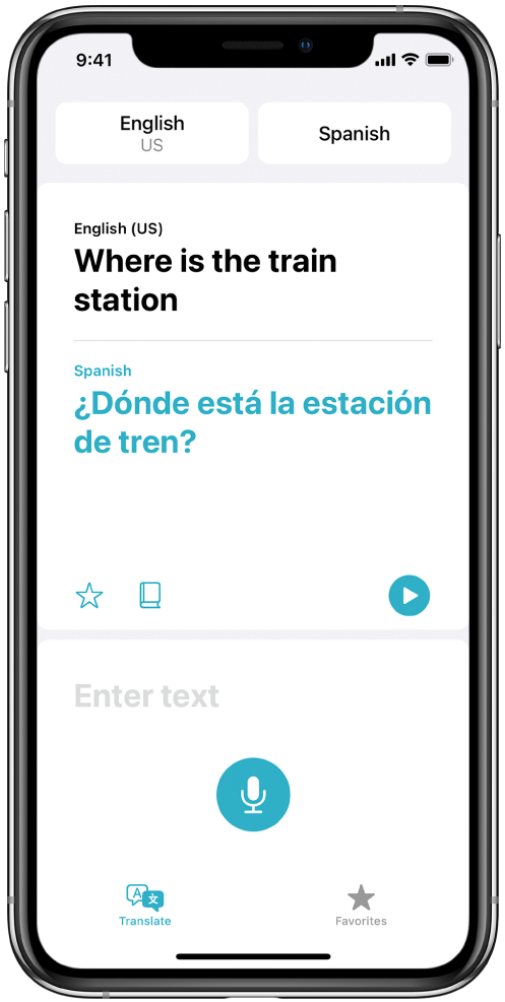

Figure 19: Siri

- Siri's workflow:
  - Speech recognition
  - Language analysis
  - Text to speech

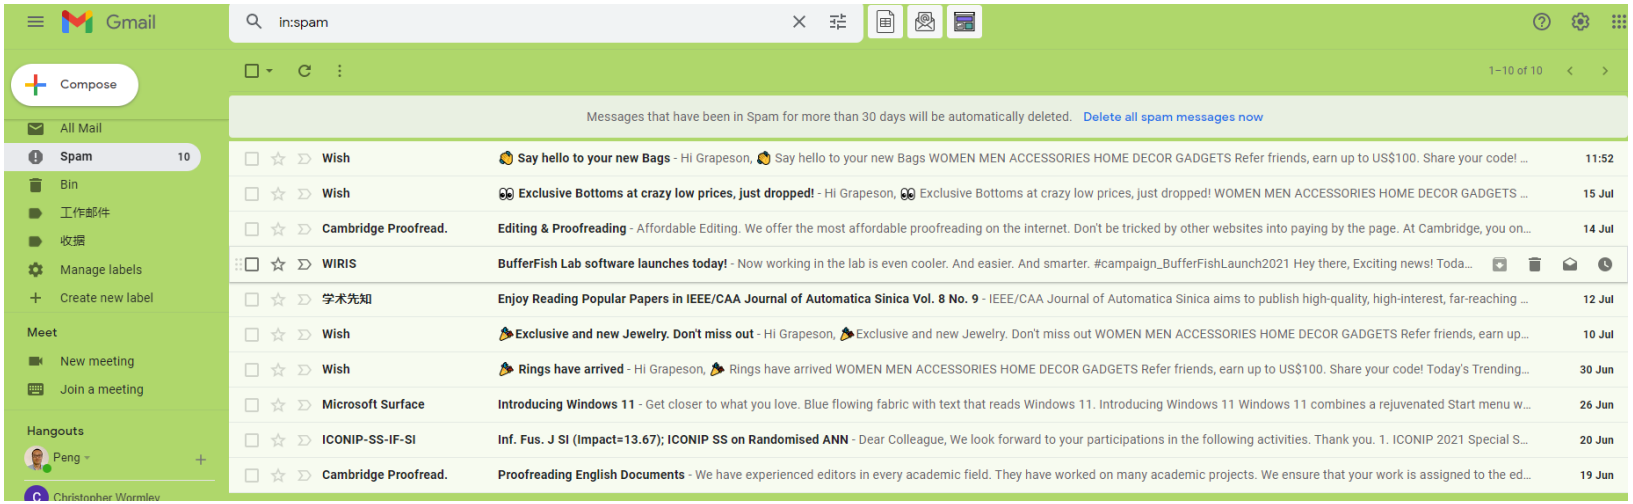

Figure 20: Spam detection

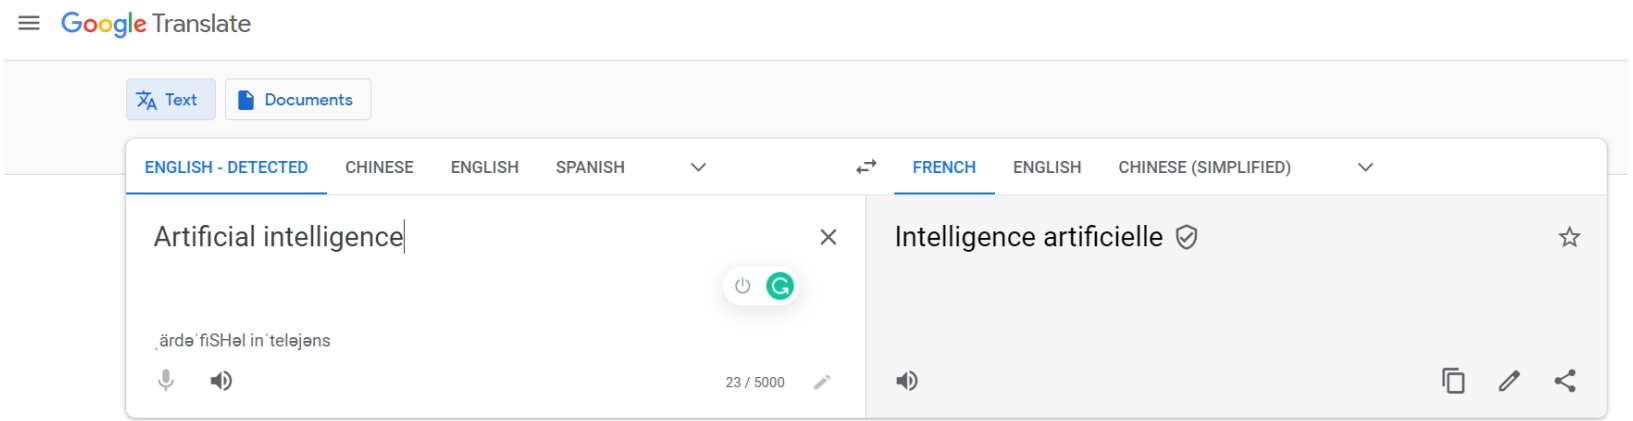

Figure 21: Language detection and translation

Behind the Applications

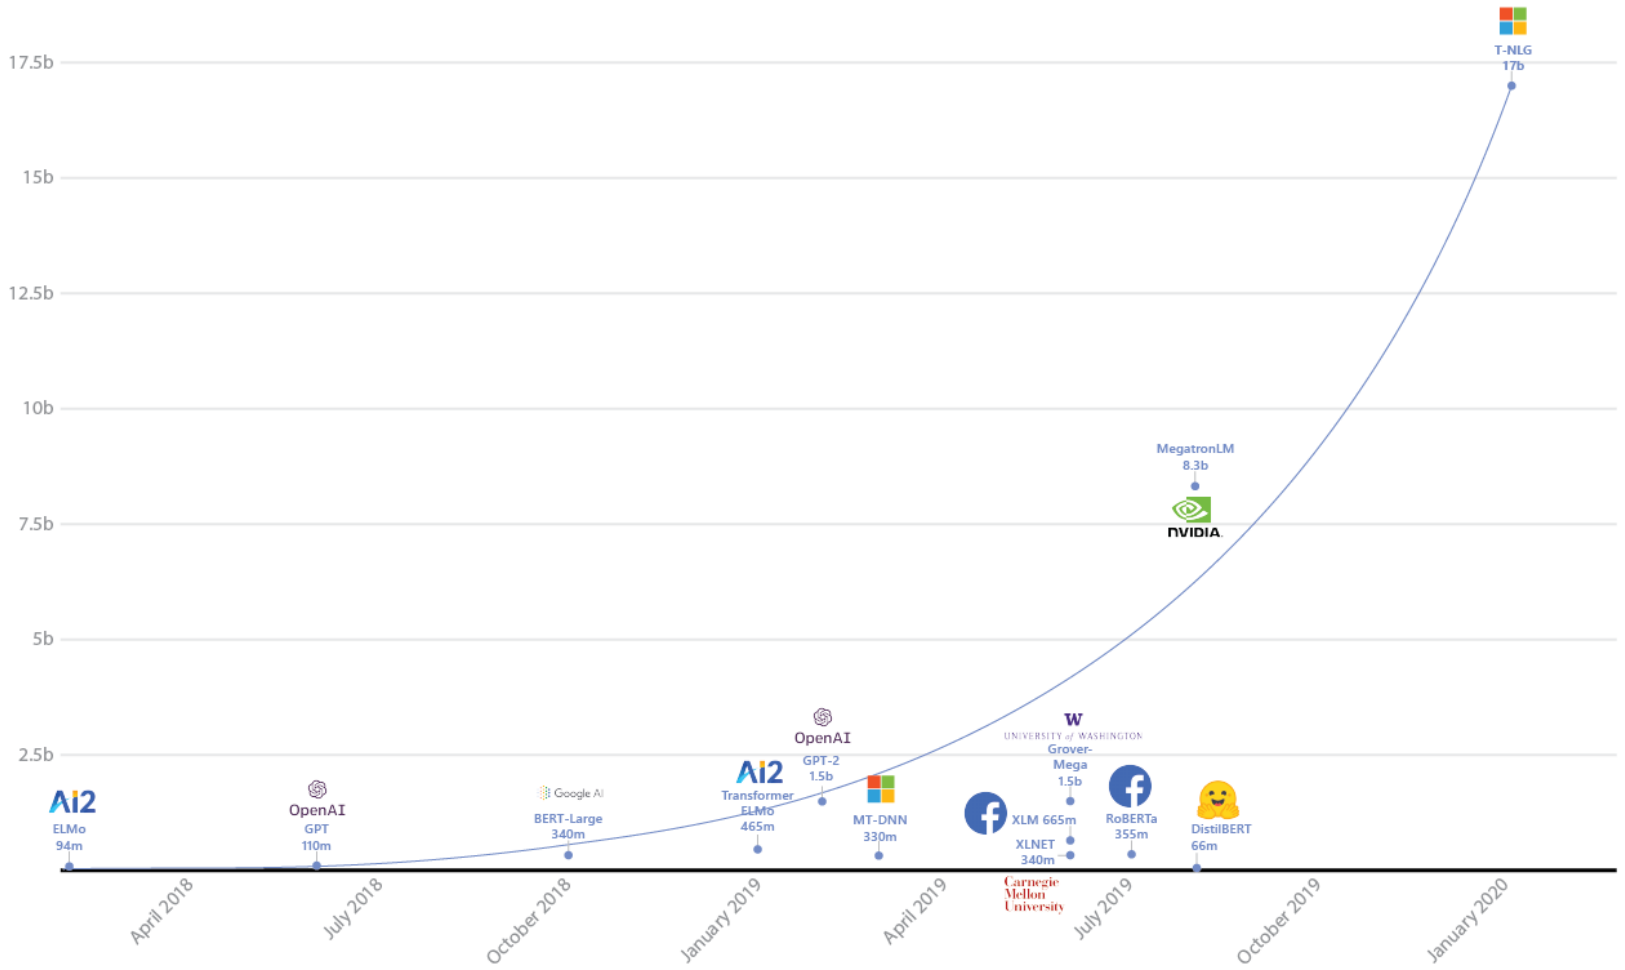

Figure 22: Exponential parameter increment, See SOTA NLP

From 'Context-independent' to 'Context-aware'.

AI in Games

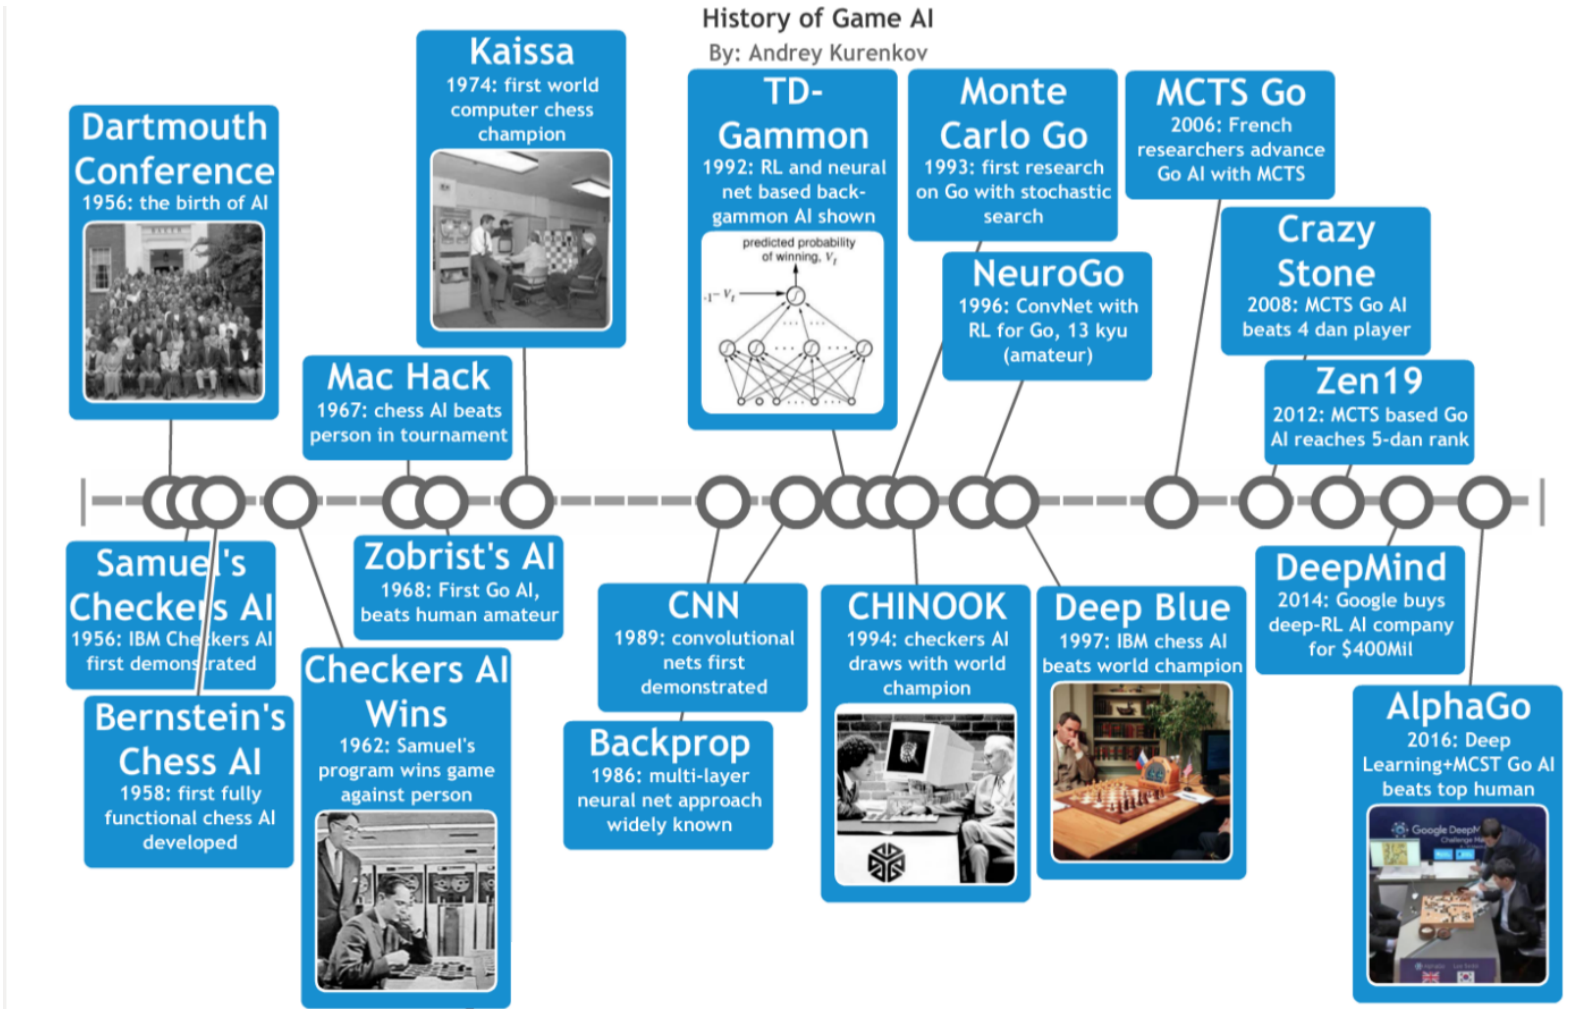

Figure 23: History of AI in Game, Credit

Classic Moment - 1

Figure 24: Deep Blue defeated Kasparov, 11/05, 1997

'Kasparov had won the first game, lost the second and then drawn the following three. When Deep Blue took the match by winning the final game, Kasparov refused to believe it.' Read more.

Classic Moment - 2

Figure 25: AlphaGo defeated Lee Sedol, 9,10,12,13, and 15 - March 2016

'Out of five games, AlphaGo won four games and Lee won the fourth game which made him recorded as the only human player who beat AlphaGo in all of its 74 official games.' Read more.

Deep Blue and AlphaGo

Figure 26: Deep Blue

Figure 27: AlphaGo

- Driven by minmax search
- Run on IBM super computer
- Intelligent via 'Brute-force'

- Driven by Reinforcement Learning
- Run on Google cloud
- Intelligence via learning

To know more about the differnces, read this blog.

Summary

Manchester Metropolitan University

In this session,

- We learned what's intelligence and artificial intelligence, and the difference and connection between them.
- We learned four different definitions of AI, and we emphasised on the 'rational agent' definition.
- We now know the 70 years' history of AI. We know the ups and downs in the history. We are witnessing the big data and deep learning era. We understand the opportunity and challenges as well.
- We learned the key players from industry and academia.
- We learned the state-of-the-art techniques in CV, Robotics, NLP, and Gaming.

Pros and Cons of AI

Manchester Metropolitan University

Discussion on the pros and cons of AI

Artifical superintelligence (ASI) is a good thing or bad thing?