Artificial Intelligence beyond Maching Learning, Deep Learning

François Jézéquel
Orange Fab France
General Member Representative at LF AI Governing Board

LF AI Day Europe

June 22nd 2020

The dizzying scope of Artificial Intelligence

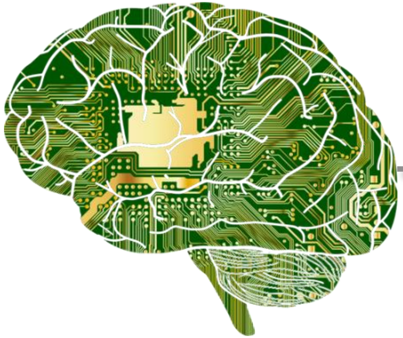

Source: from
Artificial Intelligence: A Modern Approach,
4th Edition
Stuart Russell, Peter Norvig

Problem Solving
- Solving Problems by Searching
- Search in Complex Environments
- Adversarial Search and Games
- Constraint Satisfaction Problems

Knowledge and Reasoning
- Logical Agents
- First-Order Logic
- Inference in First-Order Logic
- Knowledge Representation
- Automated Planning
- Quantifying Uncertainty

Uncertain Knowledge and Reasoning
- Probabilistic Reasoning
- Probabilistic Reasoning over Time
- Probabilistic Programming
- Making Simple Decisions
- Making Complex Decisions

Learning
- Multiagent Decision Making
- **Learning from Examples**
- Learning Probabilistic Models
- **Deep Learning**

Communicating, Perceiving, and Acting
- **Reinforcement Learning**
- Natural Language Processing
- **Deep Learning** for Natural Language Processing

Robotics

AI - Machine Learning impact on research and economy

Number of AI papers on arXiv, 2010-2019
Source: arXiv, 2019.

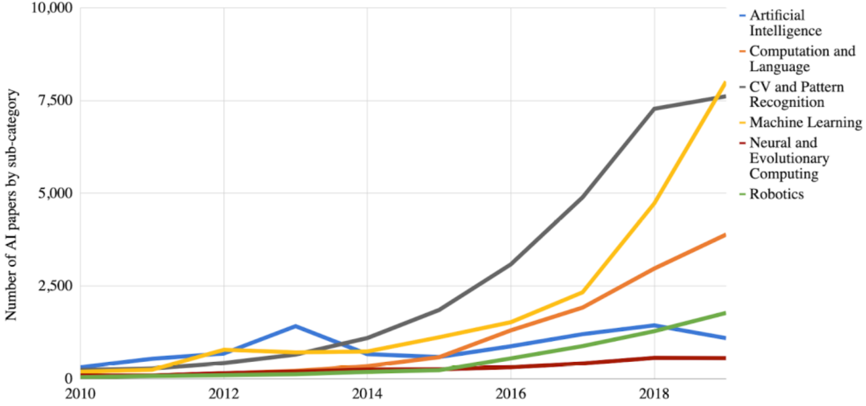

Share of Total Online Job Postings, USA, 2010-2019 monthly
Source: BurningGlass, 2019.

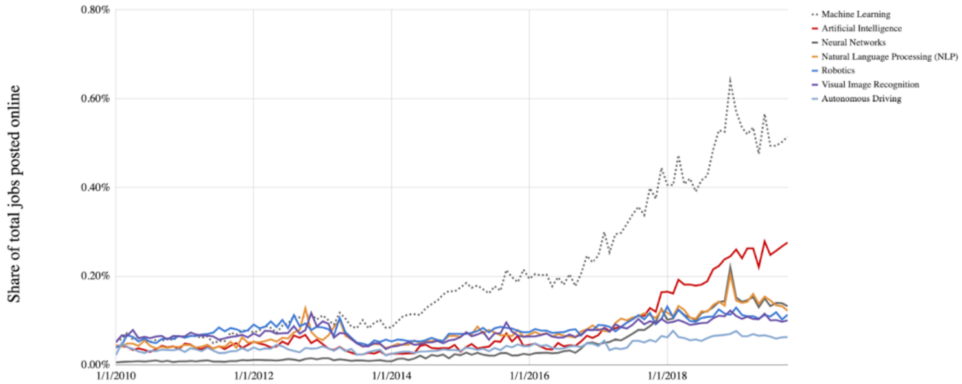

AI software market: **$102.0 Billion** in annual worldwide revenue by 2024
Source: Tractica

Global machine learning market **$36.7 Billion** by 2024
Source: Market Research Future

orange™

Artificial General Intelligence – Definition and capability targets

Artificial general intelligence (AGI) is the representation of generalized human cognitive abilities in software so that, faced with an unfamiliar task, the AI system could find a solution. An AGI system could perform any task that a human is capable of.
Source: SearchEnterprise AI

Source: Gary Marcus: The Road to Artificial General Intelligence – GMIS 2017

orange™

Artificial General Intelligence – level of abstractions and approaches

Structure

Rationale: Intelligence is produced by the human brain. Therefore, to build an intelligent computer means to simulate the brain structure as faithfully as possible.
Background: Neuroscience, biology, etc.
Challenge: There may be biological details that are neither possible nor necessary to be reproduced in AI systems.

HTM
Vicarious

Behavior

Rationale: Intelligence is displayed in how the human beings behave. Therefore, the goal should be to make a computer to behave exactly like a human.
Background: Psychology, linguistics, etc.
Challenge: There may be psychological or social factors that are neither possible nor necessary to be reproduced in AI systems.

Turing Test
cognitive model

Capability

Rationale: Intelligence is evaluated by problem-solving capability. Therefore, an intelligent system should be able to solve certain practical problem that is currently solvable by humans only.
Background: Computer application guided by domain knowledge
Challenge: There is no defining problems of intelligence, and the special-purpose solutions lack generality and flexibility.

AlphaGo
expert system

Function

Rationale: Intelligence is associated to a collection of cognitive functionality, such as perceiving, reasoning, learning, acting, communicating, problem solving, etc. Therefore the goal is to reproduce these functions in computers in a divide-and-conquer manner.
Background: Computer science
Challenge: The AI techniques developed so far are highly fragmented and rigid, and it is hard for them to work together.

Mainstream AI textbooks
Soar

Principle

Rationale: Intelligence is a form of rationality or optimality. Therefore, an intelligent system should always "do the right thing" according to certain general principles.
Background: Logic, mathematics, etc.
Challenge: There are too many aspects in intelligence and cognition to be explained and reproduced by a simple theory.

AIXI
NARS

Source: Pei Wang Department of Computer and Information Sciences, College of Science & Technology, Temple University

Artificial General Intelligence – Some leading organizations

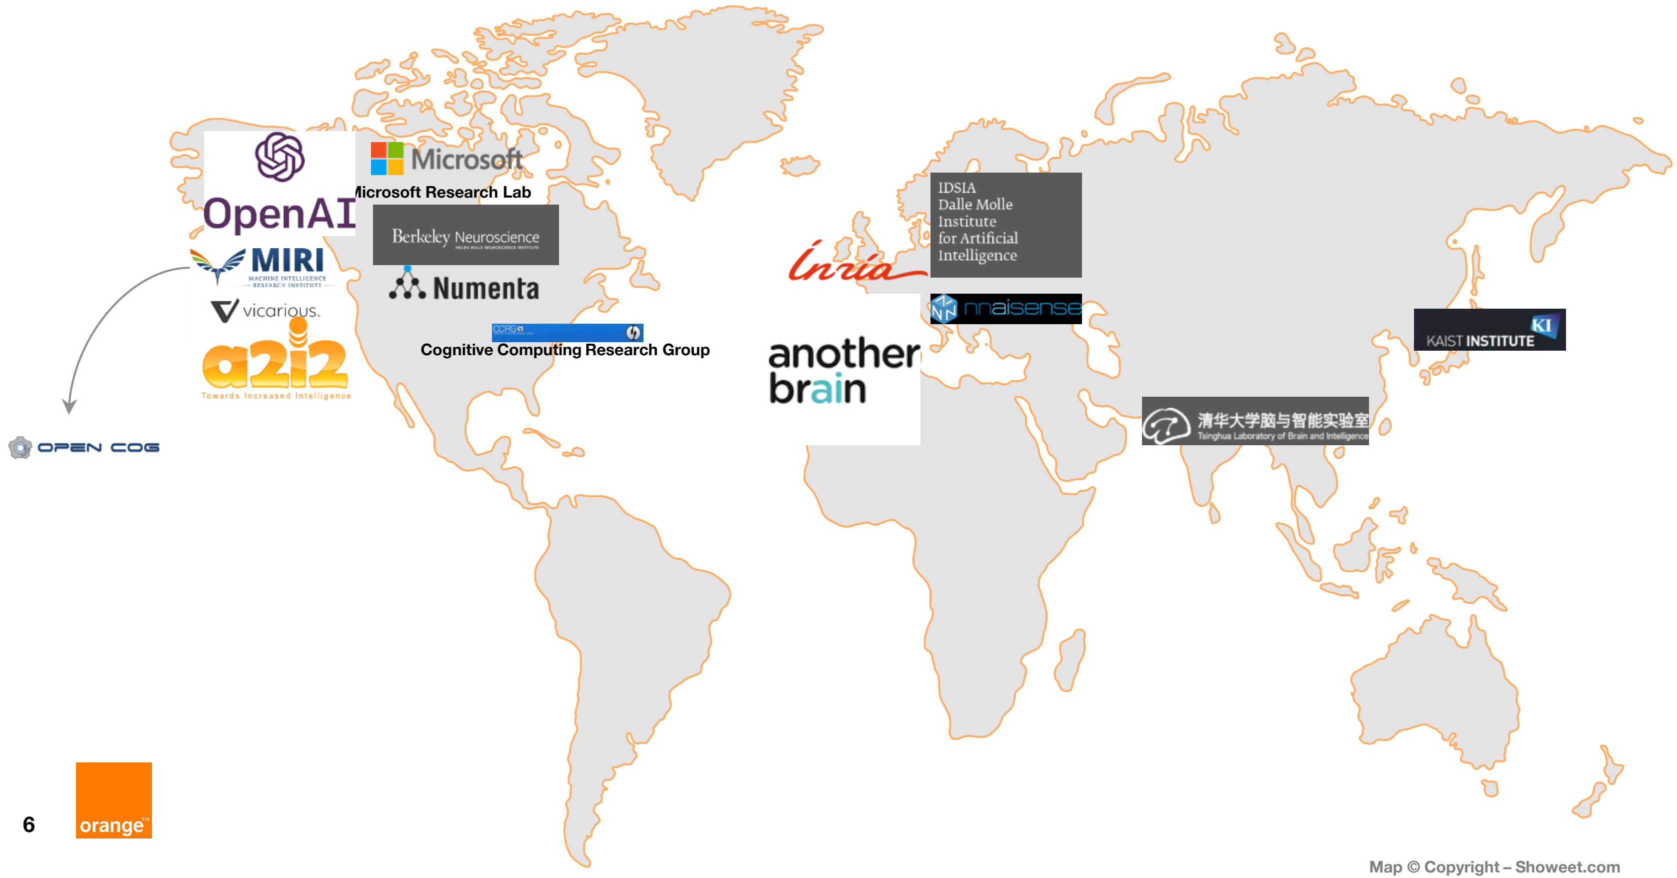

Map © Copyright – Showeet.com

Let’s conduct a simple experiment…

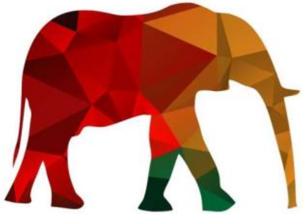


… to attempt a glimpse at how the neocortex works

- **The neocortex has a hierarchical architecture**

- **The neocortex follows a memory-prediction model**
  Your brain with its massive memory capacity is constantly predicting what you see, hear or feel, mostly in ways you are unconsious of. When you already have seen 2 elephants the next elephant recognition is accelerated

- **The neocortex creates invariant representations**
  Your brain is able to recognize an elephant representation whatever the angle of vision, the texture, the color

- **The neocortex associates the different senses**
  The trumpeting information flowed up the auditory hierarchy to an association area that connectes vision with hearing to confirm that you are seeing an elephant

From neuroscience foundations to intelligent machines?

(Credit: Johns Hopkins Medicine)

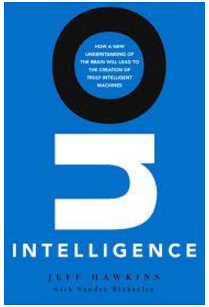

Numenta

https://github.com/numenta

Vernon B. Mountcastle
(1918–2015), once dubbed "the Jacques Cousteau of the [cerebral] cortex", is widely considered the father of modern neuroscience, thanks to his outstanding achievements in the 1950s when he pioneered neurophysiological studies of the cortex.

Jeff Hawkins
(1957-) is the American founder of Palm Computing and Handspring where he invented the PalmPilot and Treo, respectively. He has since turned to work on neuroscience full-time, founding the Redwood Center for Theoretical Neuroscience (formerly the Redwood Neuroscience Institute) in 2002 and Numenta in 2005. Hawkins is the author of *On Intelligence*

Advancing Machine Intelligence with Neuroscience
In the brain, cortical networks are sparsely connected and extremely dynamic. As many as 30% of the connections in the neocortex turn over every few days. Numenta is investigating ways to create highly sparse networks that learn their structure dynamically through training

Mountcastle, V. (1978). "An organizing principle for cerebral function: the unit model and the distributed system," in The Mindful Brain, eds G. Edelman and V. Mountcastle (Cambridge, MA: MIT Press), 7–50.
Mountcastle, V. B. (1997). The columnar organization of the neocortex. Brain 120, 701–722. doi: 10.1093/brain/120.4.701

orange

Some suggested perspectives to be explored

- **AI domain-related initiatives**

  Knowledge Representation and Reasoning - Grakn is an intelligent database: a knowledge graph engine to organise complex networks of data and make it queryable.

  Robotics - Robot Operating System (ROS) - a set of software libraries and tools to build robot applications - https://www.ros.org/

- **Federation of vertical oriented projects**

  Healthcare example

  Healthcare KR&R - OpenClinical information service, a body of resources on advanced knowledge management methods, technologies and applications for healthcare, e.g. Health ontologies

  Healthcare robotics - Boston Dynamics open-sourced its health care robotics toolkit on GitHub

- **AGI approaches**

  Neuro science-based example

  Hierarchical Temporal Memory (HTM) - a theory of intelligence based strictly on the neuroscience of the neocortex implemented in Numenta Platform for Intelligent Computing

Thank you