

AN INTRODUCTION TO ARTIFICIAL INTELLIGENCE



Marco Gori
University of Siena

OUTLINE

- HISTORY
- SYMBOLIC AND SUB-SYMBOLIC AI
- IS AI CHALLENGING HUMAN INTELLIGENCE?
- RISKS AND GEOPOLITICAL ISSUES
- SLOWDOWN? DECENTRALIZED AI

HISTORY

HISTORY

Toys in the research labs


- 1956 Dartmouth Summer Research Project on AI (Minsky, McCarty, Rochester, Shannon, ...)
- 1959 Widrow ADALINE: application to ASR
- 1962 Rosenblatt's perceptron
- 1969 Minsky & Papert's seminal book
- '70 Anderson's associators, Kohonen's self-organizing maps
- 1982 Hopfield nets
- 1986 Backpropagation and the PDP ...
- 2012 Explosion of deep learning
- 2017 Transformers and Large⁴ Language Models
- **2022 GPT** ...

AI

THE BREAKTHROUGH OF CONNECTIONISM - PDP research group

big collections and
computational resources

the wave which supports AI



IN THE LABS

the role of games

M. Minsky (1968)

“it is not that the games and mathematical problems (puzzles) are chosen because they are clear and simple; rather it is that they give us, for the smallest initial structures, the greatest complexity”

SYMBOLIC AND SUB-SYMBOLIC AI



IT IS NOT FORCE ONLY!



Gioco del 24:

$$25!/2 \approx 7.76 \times 10^{24}$$








CHE EAVBOIS
GRANDE




CHAR RECOGNITION VS CHESS-BOARD

Why do machines need to learn?



*Handwritten
characters: The 2^d
warning!*




~ 0001100000100100000000100000001000000010100001000111110000000011.

THE COGNITIVE PROCESS OF VISION



pixel


SEGMENTATION AND RECOGNITION



Segmentation and
Recognition

The chicken-egg
dilemma

SEGMENTATION AND RECOGNITION (CON'T)



Problems with time warping, different speakers, noise ...

Segmentation: non-just Energy-based thresholding!


FROM NEUROSCIENCE

Neural activity:
about 70 mV, 50 Hz




100 billion neurons, 10,000 connections → **1000 trillion connections**

FROM NEUROSCIENCE (CON'T)




Secrets from biology?


NEURAL MODEL



$$a_i = b_i + \sum_{j=1}^d w_{i,j} x_j,$$
$$y_i = \sigma(a_i) = 1/(1 + e^{-a_i}).$$




LEARNING AS WEIGHT UPDATING



Borrowing from biology:
Learning consists of an appropriate change of the synaptic connections

DECISION PROCESS


preprocessing




π



learning




output decoding




HOW TO MEASURE THE PROGRESS OF LEARNING

Case of classification




SYNAPTIC CONNECTIONS IN POPULAR LLM



... trillions of connections ...

Human brain is made of 1,000 trillion connections
... but the speed is not comparable!

EMERGENCE OF INTELLIGENCE



IS AI CHALLENGING HUMAN INTELLIGENCE?

... Intelligence which emerges from the Web / huge data collections

Fact #1 - AI is 60 years old




John McCarthy (1927-2011)
1956 - Coined the term "Artificial Intelligence"
1958 - Invented LISP
1971 - Received the Turing Award





Marvin Minsky (1927-2016)
1959 - Co-founded the MIT AI Lab
1968 - Advised Kubrick on "2001: A Space Odyssey"
1969 - Received the Turing Award

aws

From "AI labs" to the cinema



?



GAMES

M. Minsky (1968)

“it is not that the games and mathematical problems (puzzles) are chosen because they are clear and simple; rather it is that they give us, for the smallest initial structures, the greatest complexity”

- Puzzles (Sam Loyd Game, Rubik's cube, ...)
- Chess
- Go
- ...

MEDICAL DIAGNOSYS

A patient reports experiencing fatigue, and blood tests reveal a significant hemoglobin deficiency. Additionally, she has gained weight and mentions noticing some changes in his voice. What could be the diagnosis?


GPT-4

The symptoms described — **fatigue, significant hemoglobin deficiency (anemia), weight gain, and voice changes** — may point to **hypothyroidism**, particularly if the voice change is described as **hoarseness** or a **deepening** of the voice.

Possible Diagnosis: Hypothyroidism

ALGORITHM DESIGN AND PROGRAMMING

There are 15 matches on the table, and two players (an opponent and the computer) take turns selecting a number of matches, denoted by a and c . Each player can take a number of matches such that $a \in [1, 3]$, $c \in [1, 3]$. The player who picks up the last match loses.



PYTHON CODE

```
1 #
2 # Nim-like
3 #
4 import random
5 n=int(input("Quanti fiammiferi sono sul tavolo all'inizio? "))
6
7 m=4 # Modulo del gioco: Possiamo prendere 1,2,3 fiammiferi.
8 PlayFirst=input("Vuoi giocare per primo? (s/n) ")
9 if PlayFirst=='s':
10     a=int(input("Quanti fiammiferi prendi? "))
11     n=n-a
12 if (n==0) and (PlayFirst=='s'):
13     win=True
14 else: win=False
15 #
16 while (n>1):
17     if ((n-1) % m == 0):
18         c=random.randint(1,m-1) # Il gioco e' nelle mani
19         dell'avversario
20     else:
21         c=(n-1) % m # Strategia ottima
22         win=True
23     n=n-c
24     print("Io prendo %d fiammiferi" %c)
25     if (n>1): # chiaramente, non devo procedere se n=1
26         a=int(input("Quanti fiammiferi prendi? "))
27         n=n-a
28 if win: print("Ho vinto!");
29 else: print("Questa volta ho perso ...")
```

GPT only made a trivial mistake

CREATIVITY

- Music (Suno AI, Udio, ...)
- Drawing (Dall-e, ...)
- film (Sora, Runway, Veo, ...): an experiment in our lab (with Luigi Perotti): <https://vimeo.com/1063132376?share=copy>
- ...

INTELLIGENCE AND CONSCIOUSNESS

- Homo sapiens is no longer alone on Earth in achieving high-level cognitive expressions and possessing language. It is not the strongest — and perhaps not even the most intelligent!
- The Cartesian mind-body dilemma has finally been dissolved; AI-powered computers now support computational processes similar to those of the human brain, and the trillions of synaptic connections that once defined it now also reside in GPUs!
- The robot can achieve self-perception, and everything related to our interpretation of consciousness finds a natural extension in machines.

CONSCIOUSNESS AND RESPONSIBILITY

- Roger Penrose: “Consciousness is not computation!”
- Humans perceive pain. Not by chance, the concept of ethics has developed, and civilized countries establish laws to limit and punish any action that causes pain. Do machines feel pain? Can we transfer the concept of a “physical person” (capable of holding rights and obligations) to that of a “physical machine”? *In reality, software is independent of machines — and it is not mortal!*
- In machines, faulty parts are replaced; in biology, cellular processes inevitably lead to death.
- Kurzweil’s singularity: Are cyborgs truly entities supported by solid scientific arguments? What does the progressive replacement of biological neurons with artificial ones mean? Why, after decades, do we still not have prosthetics for the blind?
- I humbly argue that responsibility should remain - as far as possible - with humans, who should use machines merely as cognitive levers!

RISKS AND GEOPOLITICAL ISSUES

TAKING AI SAFETY SERIOUSLY

International Scientific Report on the Safety of Advanced AI

- Artificial General Intelligence (AGI)
- Artificial Super Intelligence (ASI)

Yoshua Bengio:

... while we are racing towards AGI or even ASI, nobody currently knows how such an AGI or ASI could be made to behave morally, or at least behave as intended by its developers and not turn against humans ...

... Entities that are smarter than humans and that have their own goals: are we sure they will act towards our well-being?

TAKING AI SAFETY SERIOUSLY (Yoshua Bengio)

International Scientific Report on the Safety of Advanced AI

- For those who think that we may reach AGI but not ASI (analog vs digital computation)
- For those who think that AGI and ASI will be kind to us (An AI with a self-preservation goal would resist being turned off)
- For those who think that corporations will only design well-behaving AIs and existing laws are sufficient
- For those who think that we should accelerate AI capabilities research and not delay benefits of AGI (invest wisely in research to make sure we can appropriately control those risks)
- For those concerned that talking about catastrophic risks will hurt efforts to mitigate short-term human-rights issues with AI
- For those concerned with the US-China cold war

TAKING AI SAFETY SERIOUSLY (Yoshua Bengio)

International Scientific Report on the Safety of Advanced AI

- For those who think that international treaties will not work
- For those who think the genie is out of the bottle and we should just let go and avoid regulation
- For those who think that open source AGI code and weights are the solution
- For those who think worrying about AGI is falling for Pascal's wager
- For those who discard x-risk for lack of reliable quantifiable predictions

**How can - humans - control
AI-powered autonomous weapons?**

SLOWDOWN?

Pause Giant AI Experiments: An Open Letter

22 March, 2023

DECENTRALIZED AI

THE PROGRESS ON WEB SEARCH



In oriental folklore dragons not only enjoy awesome grace and beauty, they are endowed with immense wisdom. But in the west they are often portrayed as evil—St. George vanquishes a fearsome dragon, as does Beowulf—or, sometimes, friendly—Puff.

GENERATIVE AI: DRAGONS EVOLVE ...

outside academic research labs

un instant de déjà-vu



The milk: it is not just
Web search or GMail ...

Privacy issues?
Transparency?
Trustworthy AI?
Geopolitical issues?
...

Condemned to the sole intelligence that emerges from brute force?

In nature, intelligence emerges only from interactions, not from collections of data!

A UNIVERSE OF SOCIAL LEARNING AGENTS

On-board intelligence ... just like in Nature!




TOWARD DECENTRALIZED AI



COLLECTIONLESS AI PROTOCOL:
INTERACTIONS, NOT DATA COLLECTIONS

TOWARD DECENTRALIZED AI


Specialized Agents: We'll be seeing a lot of symbolic AI again!



A NEW WEB OF AGENTS

Privacy by design

Peer to peer communication protocols



The decentralization of computation unlocks a new goldmine of methodologies for advancing AI within a framework that respects privacy.

CONCLUSIONS

- Consider seriously the risks discussed in the last **International Scientific Report on the Safety of Advanced AI** - chaired by Yoshua Bengio
- Agency: precautionary principle?
- The Web has lead to centralized AI on cloud (huge data collections - the gold - in the dragon's belly)
- The proposed slowdown is difficult to implement:
The challenge of Decentralized and Collectionless AI
(interaction is the new gold)