

INTEL·LIGÈNCIA ARTIFICIAL I NATURAL: DE LA DIVERSITAT A LA GENERALITAT

José Hernández-Orallo

Universitat Politècnica de València

<http://josephorallo.webs.upv.es/>

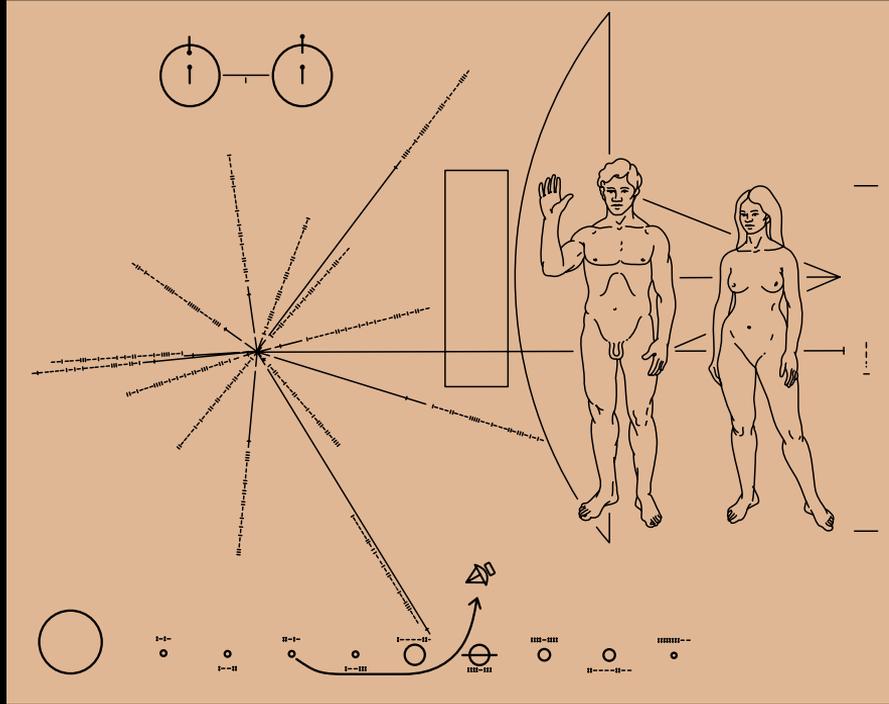
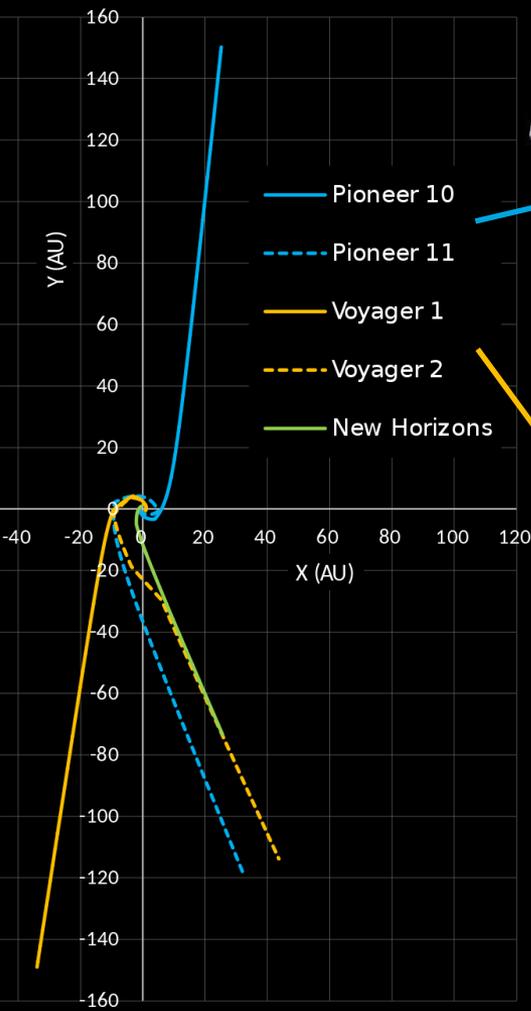


UNIVERSITAT
POLITÈCNICA
DE VALÈNCIA

A satellite view of Earth at night, showing the curvature of the planet and the dense network of city lights across the continents. The lights are concentrated in the Northern Hemisphere, particularly in North America and Europe. The text "THE ANTHROPOCENE" is overlaid in the center in a bold, white, sans-serif font.

THE ANTHROPOCENE

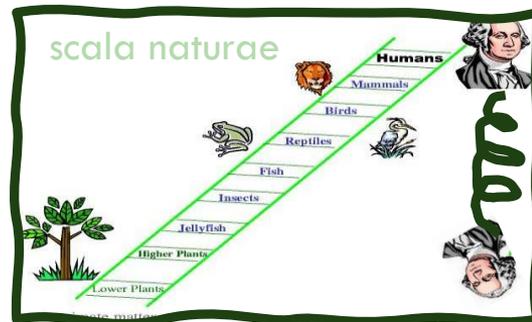
Trajectories of distant spacecraft from launch to 2030



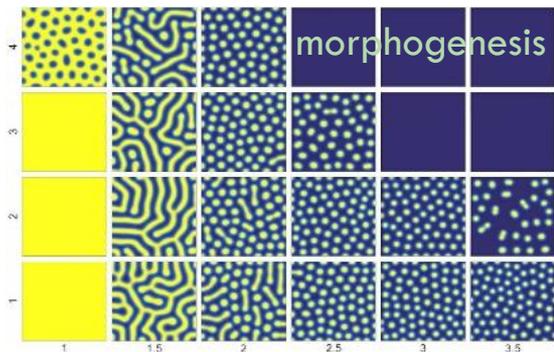
ANTHROPOCENTRISM SHOULD BE OVER



**Aristarchus,
Copernicus:**
Earth is not the
centre of the
universe



**Darwin,
Wallace:**
Humans are
not at the top
of any ladder



**Cajal, Turing,
Church:**
Life and
intelligence are
just computation
(substrate not key)

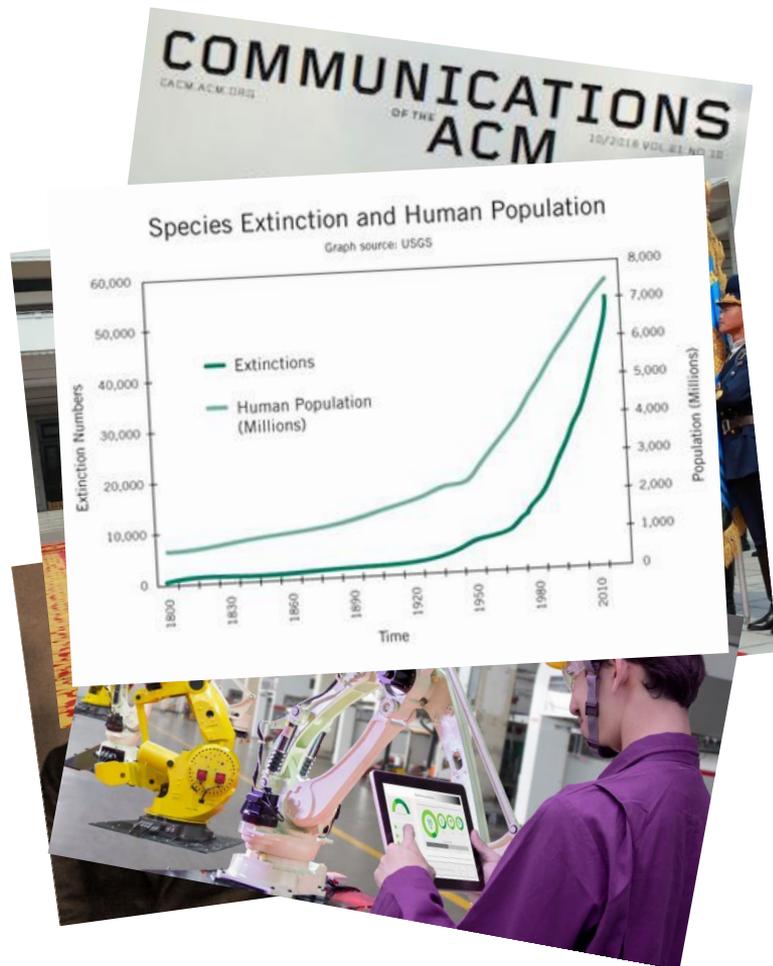


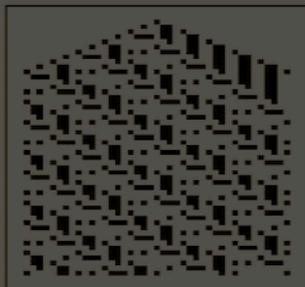
**Shannon,
Solomonoff:**
Compression
and learning
are just
algorithms

BUT IT'S LINGERING ON

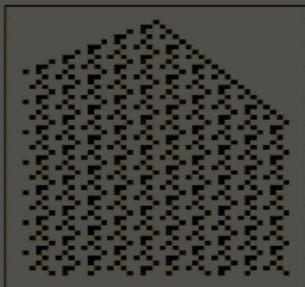
- “Human-level” intelligence as a goal?
- By Turing’s (human) imitation game?
- Human alignment. Whose values?
- Human in the loop. Or “human as a cog”?
- Existential risk. Preserve humans forever?

Should we look at AI
from an anthropocentric
perspective?

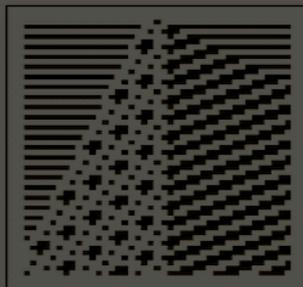




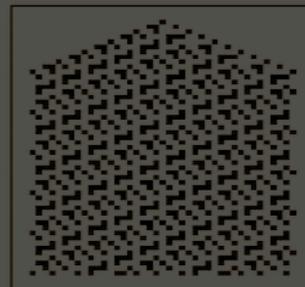
s=4: rule 209218



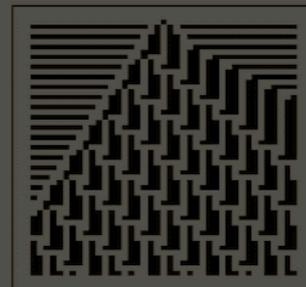
s=5: rule 838



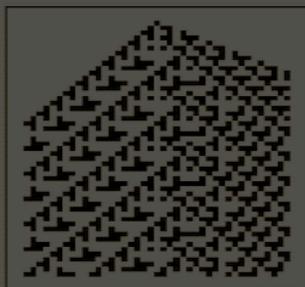
s=6: rule 774857



s=7: rule 209574



s=8: rule 8139237



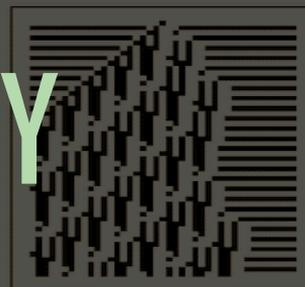
s=9: rule 2061228



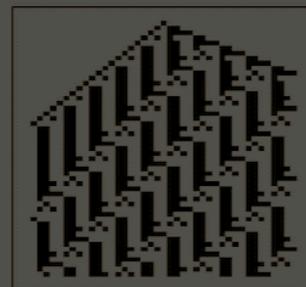
s=10: rule 316715877



s=11: rule 3215526



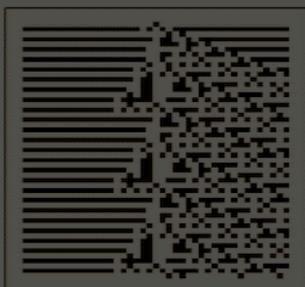
s=12: rule 82866221



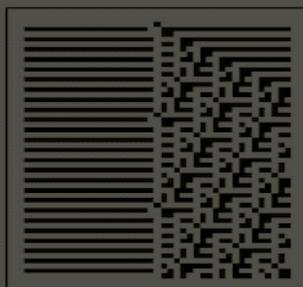
s=13: rule 22107076



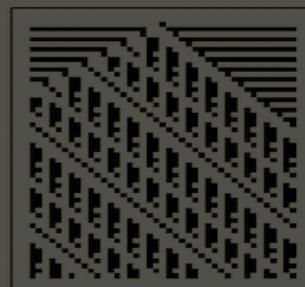
14: rule 1006223533



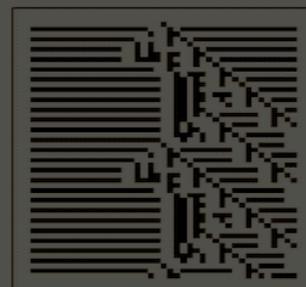
15: rule 37788005



17: rule 72844195



19: rule 273781569

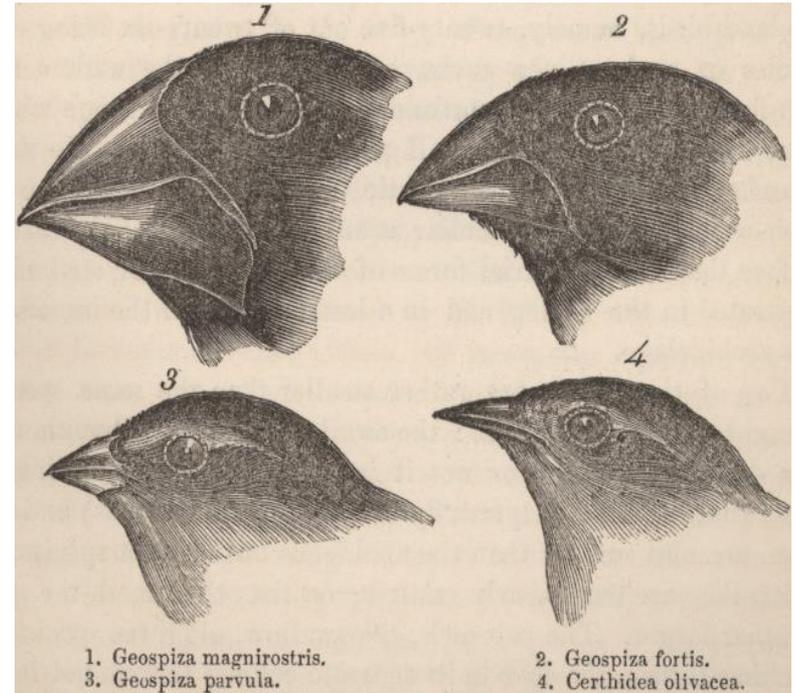


23: rule 22047073

DIVERSITY

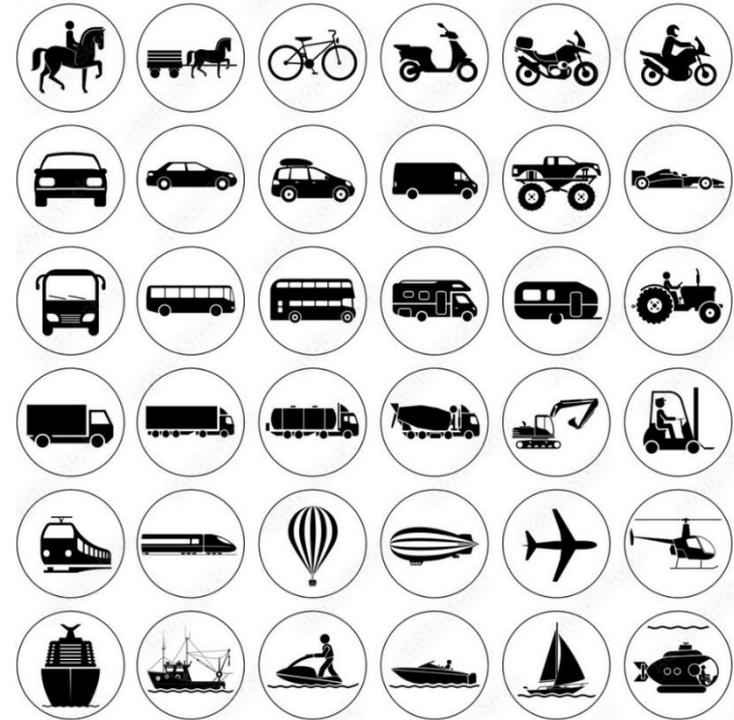
NATURAL DIVERSITY

- Ecosystems lead to niches and **specialisation**.
“for an immediate, local environment”
- Darwin's "survival of the fittest"
(borrowed from economics, Spencer):

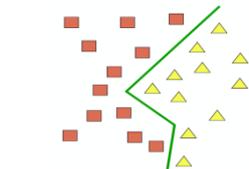


SOCIO-TECHNOLOGICAL DIVERSITY

- Specialisation in science, technology and society pays off.
 - Division of labour
 - Specialise academic disciplines
 - Divide and conquer!



DIVERSITY IN AI



Prediction and estimation

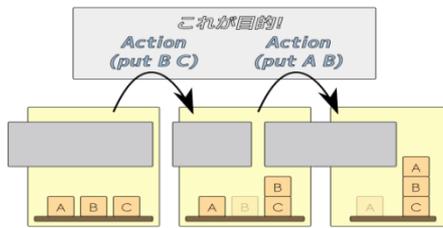
PR: computer vision, speech recognition, etc.



Knowledge-based assistants

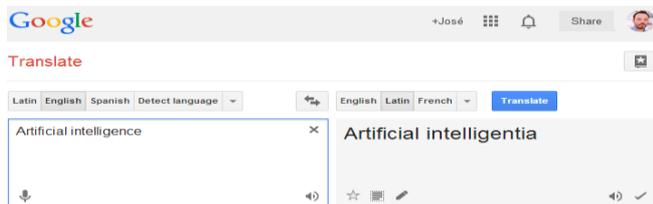


Driverless vehicles



Planning and scheduling

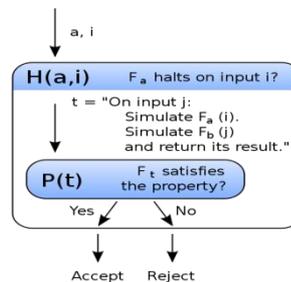
Specific (task-oriented) AI systems



Machine translation



Robotic navigation



Automated deduction

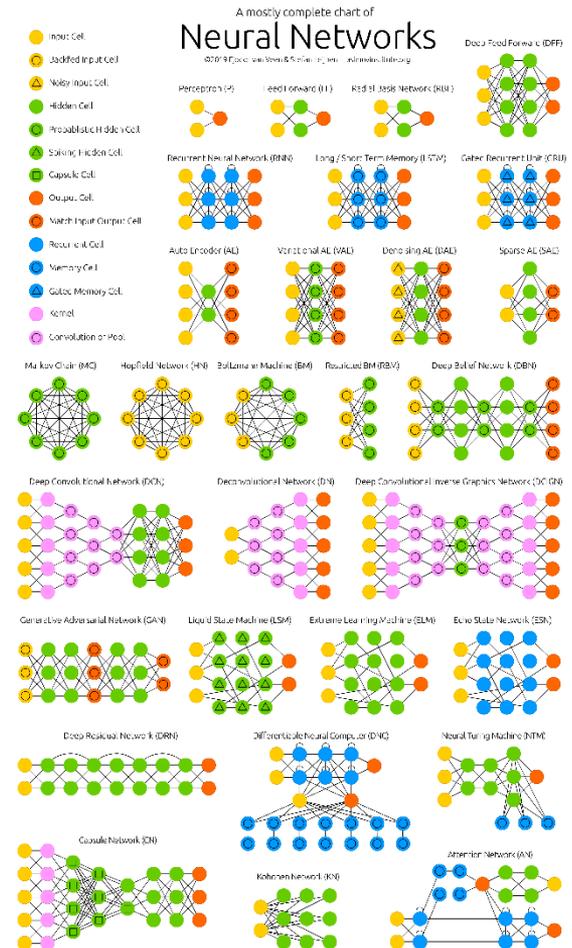


Game playing

All images from wikicommons

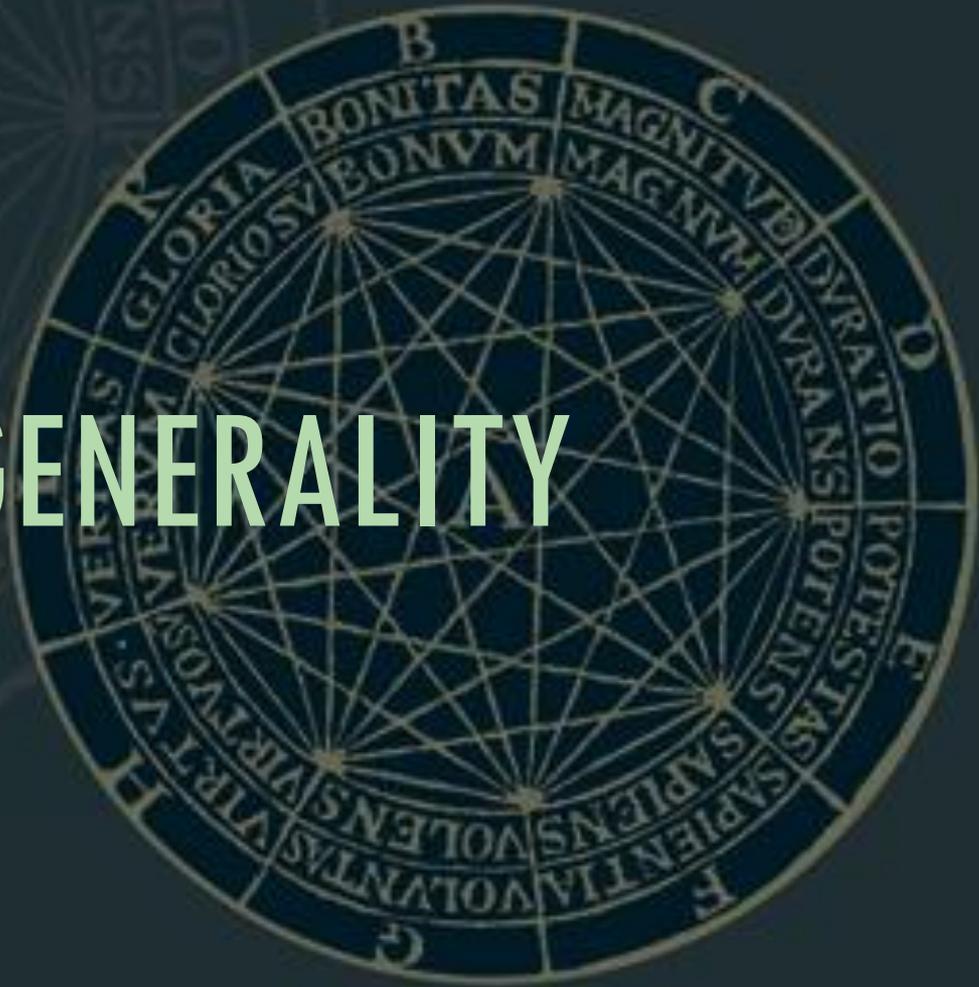
DIVERSITY IN MACHINE LEARNING

- **No-free-lunch theorems** for learning:
 - There are infinitely many niches, no system can be better for all niches, not even on average.
- Specialise ML techniques!
- Embed **learning bias** into the machines!
 - Bayes + specific prior / regulariser



<https://www.asimovinstitute.org/neural-network-zoo/> (2019)

GENERALITY



NATURAL GENERALITY

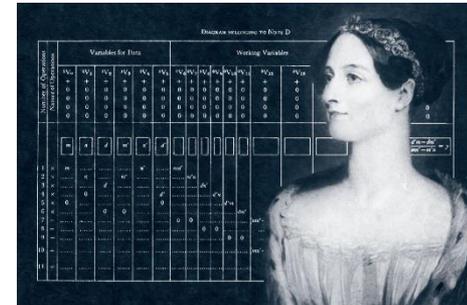
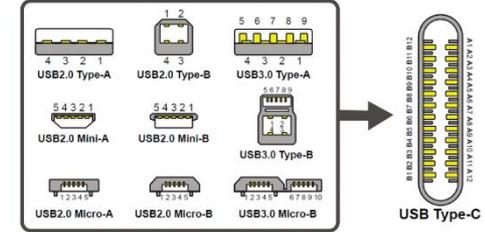
General intelligence is successful behaviour in a wide range of situations (up to a level of difficulty or resources)

- **General intelligence** and the *g* factor
 - *Spearman*: same latent factor explains performance in a range of cognitive tests.
- **Convergent evolution**
 - *General intelligence* is one of these traits!
 - Altricial vs precocial / nurture vs nature
 - Social hypothesis for general intelligence
 - Cultural hypothesis for more general intelligence



SOCIOTECHNOLOGICAL GENERALITY

- One-fits-all:
 - Standardisation or winner-takes-all
- General-purpose devices
 - Adaptable to many problems
 - Best example:
 - The **programmable computer**



Ada Byron's
"Note D"
algorithm

GENERALITY IN AI

- Lull's Ars Generalis
- Turing's "child machine"
- Simon & Newell's "General Problem Solver"
- McCarthy's dream for generality
- Artificial General Intelligence
- Solomonoff's theory of prediction:
 - Bayes + universal prior (Occam's razor)

1971
Turing
Award
Lecture

Generality in Artificial Intelligence

JOHN MCCARTHY
Stanford University

The Turing Award Lecture given in 1971 by John McCarthy was never published. The postscript that follows, written by the author in 1986, endeavors to reflect the flavor of the original, as well as to comment in the light of development over the past 15 years.

Postscript

My 1971 Turing Award Lecture was entitled "Generality in Artificial Intelligence." The topic turned out to have been overambitious in that I discovered that I was unable to put my thoughts on the subject in a satisfactory written form at that time. It would have been better to have reviewed previous work rather than attempt something new, but such wasn't my custom at that time.

I am grateful to the ACM for the opportunity to try again. Unfortunately for our science, although perhaps fortunately for this project, the problem of generality in artificial intelligence (AI) is almost as unsolved as ever, although we now have many ideas not available in

GENERALITY IN MACHINE LEARNING

- Generality is finally here!!!
- At least since 2020 (GPT-3)
- Made possible by:
 - *the transformers:*

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

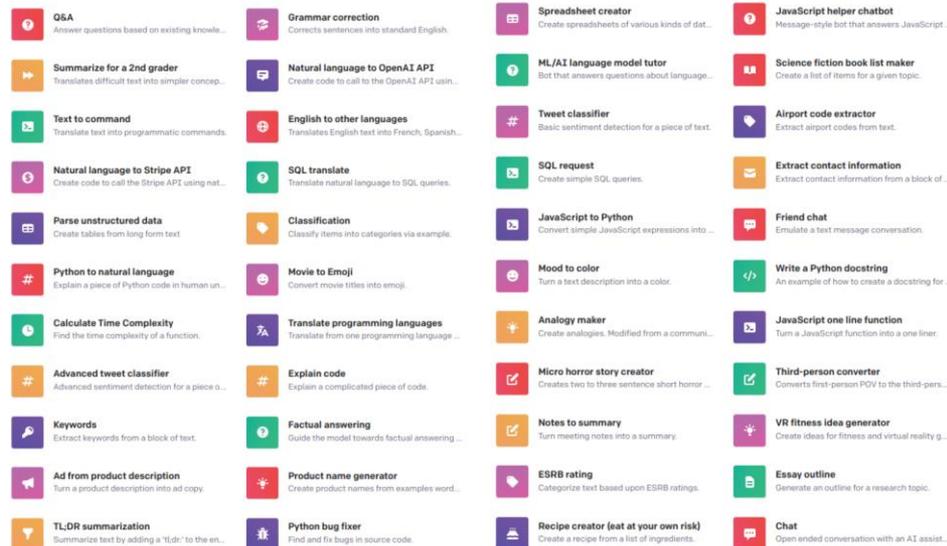
Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

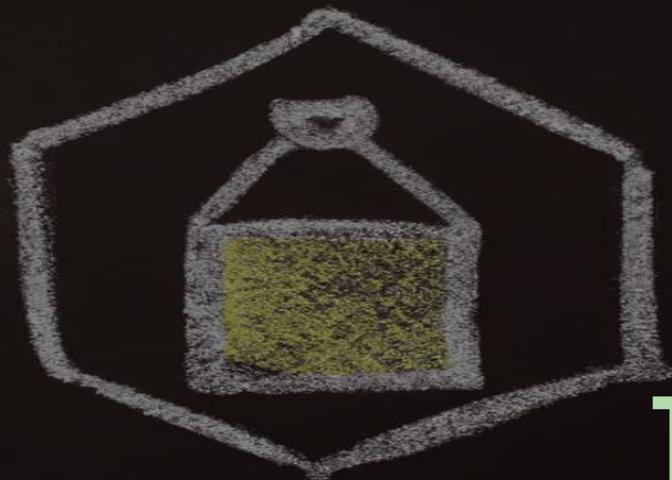
Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Łukasz Kaiser*
Google Brain
lukaszkaizer@google.com

Illia Polosukhin* ‡
illia.polosukhin@gmail.com



Task examples in OpenAI's playground



THE LESSONS



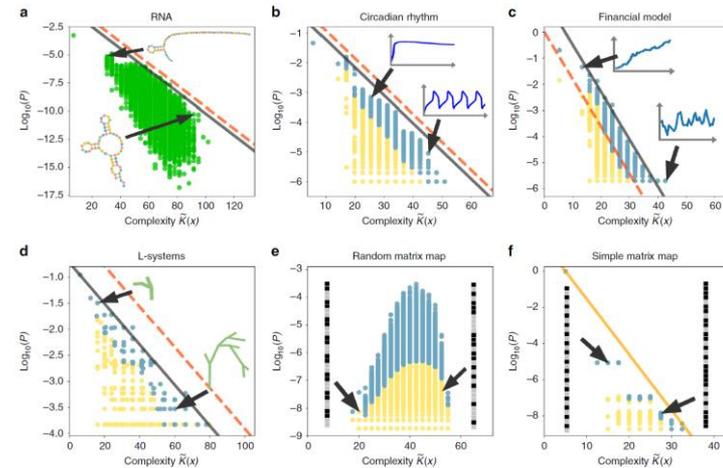
THE SALTY LESSON: “OUR” WORLD IS COMPRESSIBLE

- Physics, life and society as computation.
 - Random inputs lead to low-complexity outputs.
 - No-free-lunch assumptions not true!
 - Lunch is affordable!
- Solomonoff was right.
 - Modern deep learning and transformers have Occam’s razor as a prior!
 - Same prior is all you need!

Input–output maps are strongly biased towards simple outputs

[Kamaludin Dingle](#) ✉, [Chico Q. Camargo](#) & [Ard A. Louis](#) ✉

[Nature Communications](#) 9, Article number: 761 (2018) | [Cite this](#)



THE SWEET LESSON: IMITATION IS A GPT

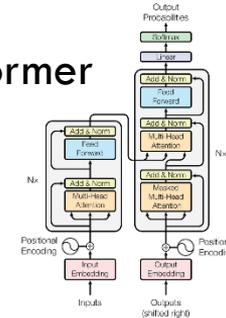
- Shannon (1948): “language model”
- Turing’s imitation game (1950)
- 2020s Surprise!
 - as compression improves, imitation improves, and many tasks become possible.



Human culture



Transformer



A
General-
Purpose
Tool

Figure 1: The Transformer - model architecture.

THE BITTER LESSON : SCALING IS ALL YOU NEED

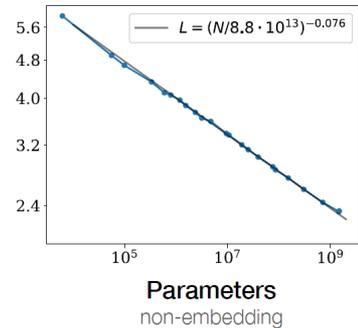
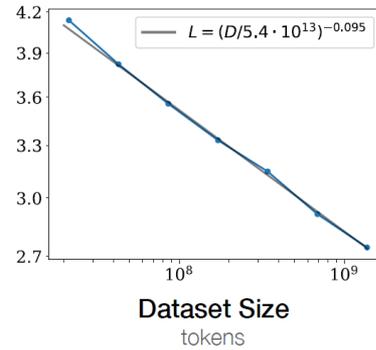
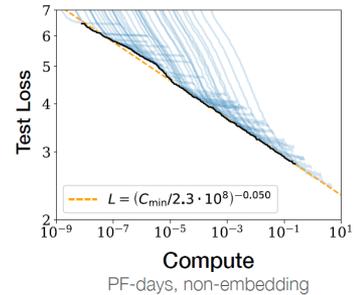
- Rich Sutton’s “bitter lesson” (2019):
 - AI has traditionally built specialised techniques that leverage “human knowledge of the domain, but the only thing that matters in the long run is the leveraging of computation”.
- Scaling laws hold better than Moore’s law!

Scaling Laws for Neural Language Models

Jared Kaplan* Johns Hopkins University, OpenAI jared@jhu.edu		Sam McCandlish* OpenAI sam@openai.com	
Tom Henighan OpenAI henighan@openai.com	Tom B. Brown OpenAI tonb@openai.com	Benjamin Chess OpenAI bchess@openai.com	Rewon Child OpenAI rowon@openai.com
Scott Gray OpenAI scott@openai.com	Alec Radford OpenAI alec@openai.com	Jeffrey Wu OpenAI jefrwa@openai.com	Dario Amodei OpenAI danodei@openai.com

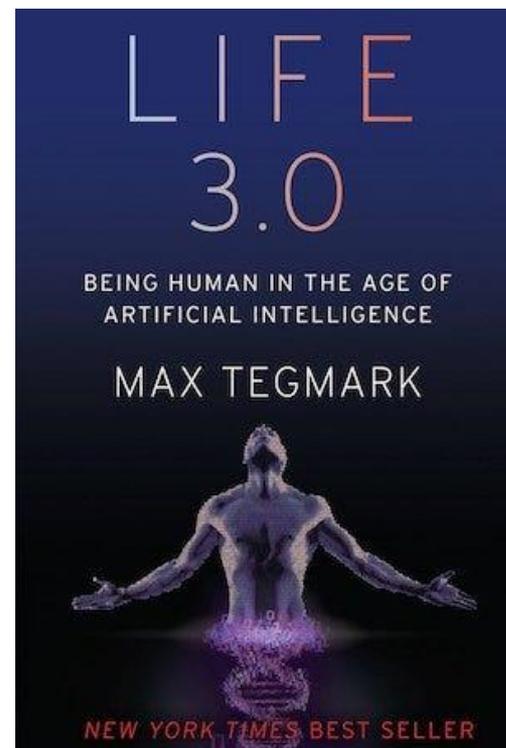
Abstract

We study empirical scaling laws for language model performance on the cross-entropy loss. The loss scales as a power-law with model size, dataset size, and the amount of compute used for training, with some trends spanning more than seven orders of magnitude. Other architectural details such as network width or depth have minimal effects within a wide range. Simple equations govern the dependence of overfitting on model/dataset size and the dependence of training speed on model size. These relationships allow us to determine the optimal allocation of a fixed compute budget. Larger models are significantly more sample-efficient, such that optimally compute-efficient training involves training very large models on a relatively modest amount of data and stopping significantly before convergence.



THE SOUR LESSON : CULTURE REPLICATION IS CHEAP

- With humans, cultural transmission is slow
 - Books, schools, universities, ...
- With AI it's fast (Hinton just realised)
- “Foundation models”
 - Generalise first, specialise next (finetune, prompt)



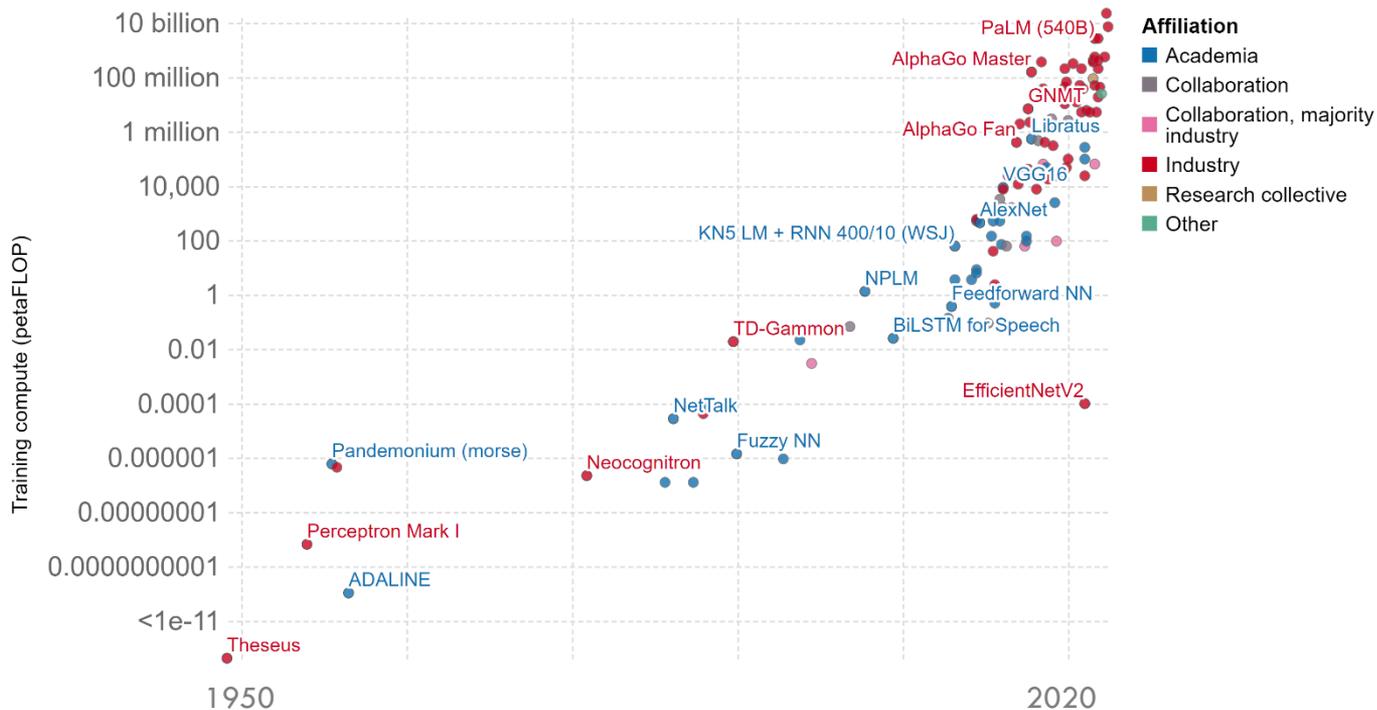


THE MECHANOCENE

* Term adapted from Huw Price's "[the machinocene](#)"

ENERGY-HUNGRY

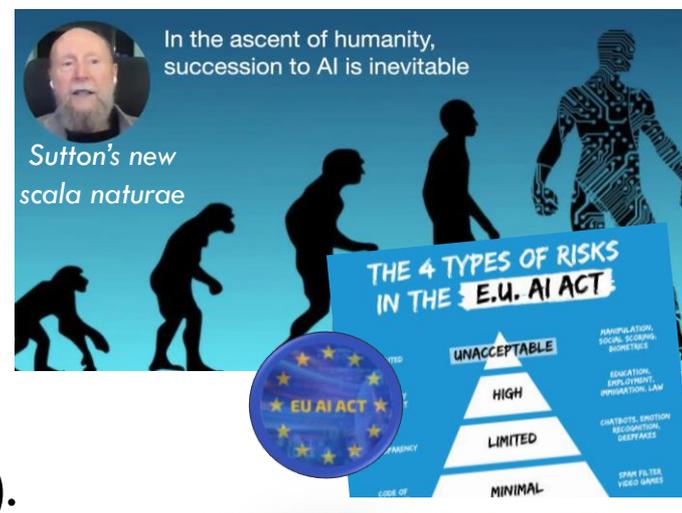
Computation is measured in total petaFLOP, which is 10^{15} floating-point operations¹.



ROOTING FOR THE MACHINES

“I can visualize a time in the future when we will be to robots as dogs are to humans. [...] I am rooting for the machines!” – Shannon (1987).

Only by understanding what intelligence is, will we make AI safe, regulate it well and know who we really are and want to be.



Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.

Signatories:

AI Scientists

Geoffrey Hinton

Emeritus Professor of Computer Science, University of Toronto

Yoshua Bengio

Professor of Computer Science, Université de Montréal

Demis Hassabis

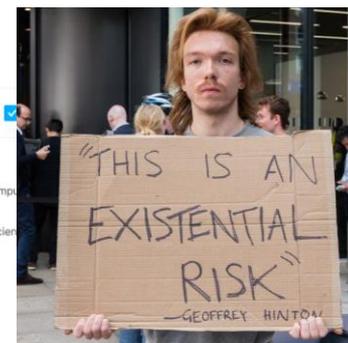
CEO, Google DeepMind

Sam Altman

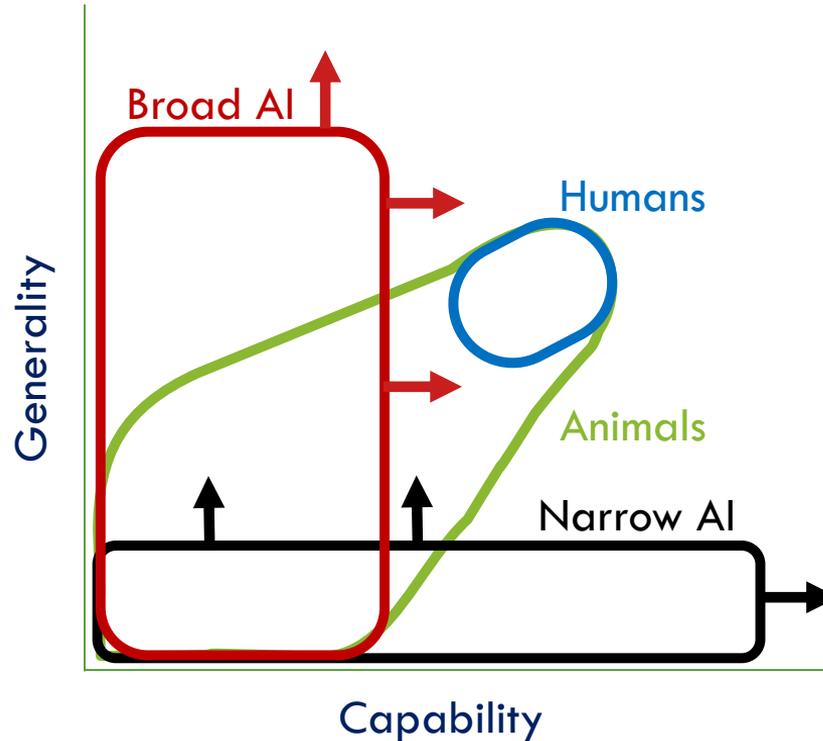
CEO, OpenAI

Dario Amodei

CEO, Anthropic



A NEW BREED : A NEW INTELLIGENCE LANDSCAPE





**THANK
YOU !**