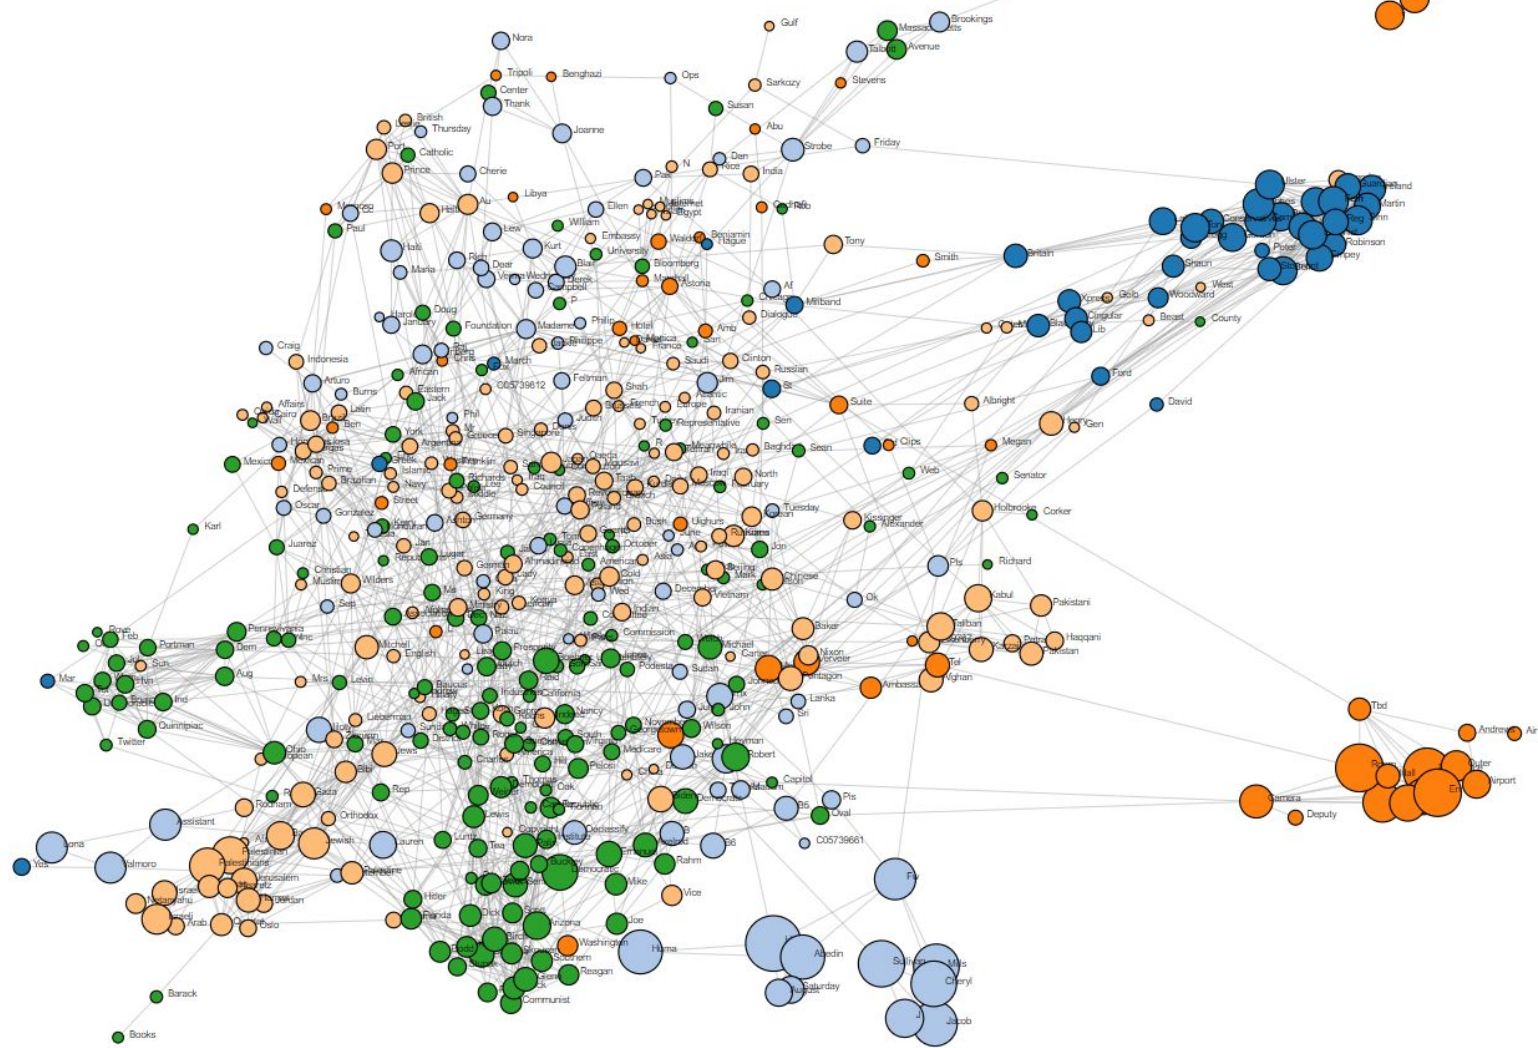


The Road to ChatGPT in 5 Breakthroughs

0. Language





1. Eliza

Welcome to

EEEEEE	LL	IIII	ZZZZZZ	AAAAA
EE	LL	II	ZZ	AA AA
EEEEE	LL	II	ZZZ	AAAAAAA
EE	LL	II	ZZ	AA AA
EEEEEE	LLLLLL	IIII	ZZZZZZ	AA AA

Eliza is a mock Rogerian psychotherapist.

The original program was described by Joseph Weizenbaum in 1966.

This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?

YOU: Men are all alike.

ELIZA: What is the connection, do you suppose ?

YOU: They're always bugging us about something or other.

ELIZA: Can you think of a specific example ?

YOU: Well, my boyfriend made me come here.

ELIZA: Is it important to you that your boyfriend made you come here ?

YOU: He says I'm depressed much of the time.

ELIZA: I am sorry to hear that you are depressed.

YOU: It's true. I am unhappy.

Welcome to

EEEEEE I I TTTT 777777 AAAAA

function ELIZA GENERATOR(user sentence) returns response

Let w be the word in sentence that has the highest keyword rank

if w exists

Let r be the highest ranked rule for w that matches sentence

response \leftarrow Apply the transform in r to sentence

if w = 'my'

future \leftarrow Apply a transformation from the 'memory' rule list to sentence

Push future onto the memory queue

else (no keyword applies)

Either

response \leftarrow Apply the transform for the NONE keyword to sentence

Or

response \leftarrow Pop the oldest response from the memory queue

Return response

Eliza is a mock
The original pro
This implementat

m in 1966.

ELIZA: Is something
YOU: Men are al
ELIZA: What is the
YOU: They're al
ELIZA: Can you thi
YOU: Well, my bo
ELIZA: Is it impor
YOU: He says I'r
ELIZA: I am sorry
YOU: It's true. I am unhappy.

r.

ou come here ?

2. Machine Learning

DATA



LOGIC



OUTPUT

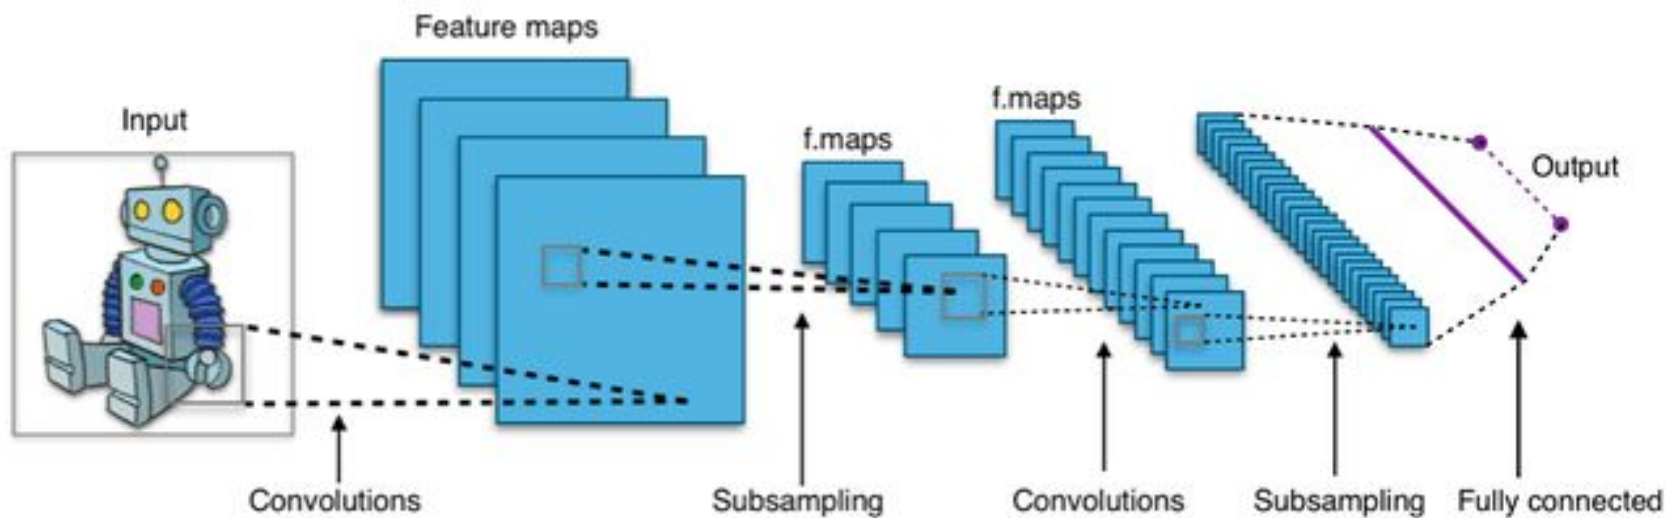
DATA



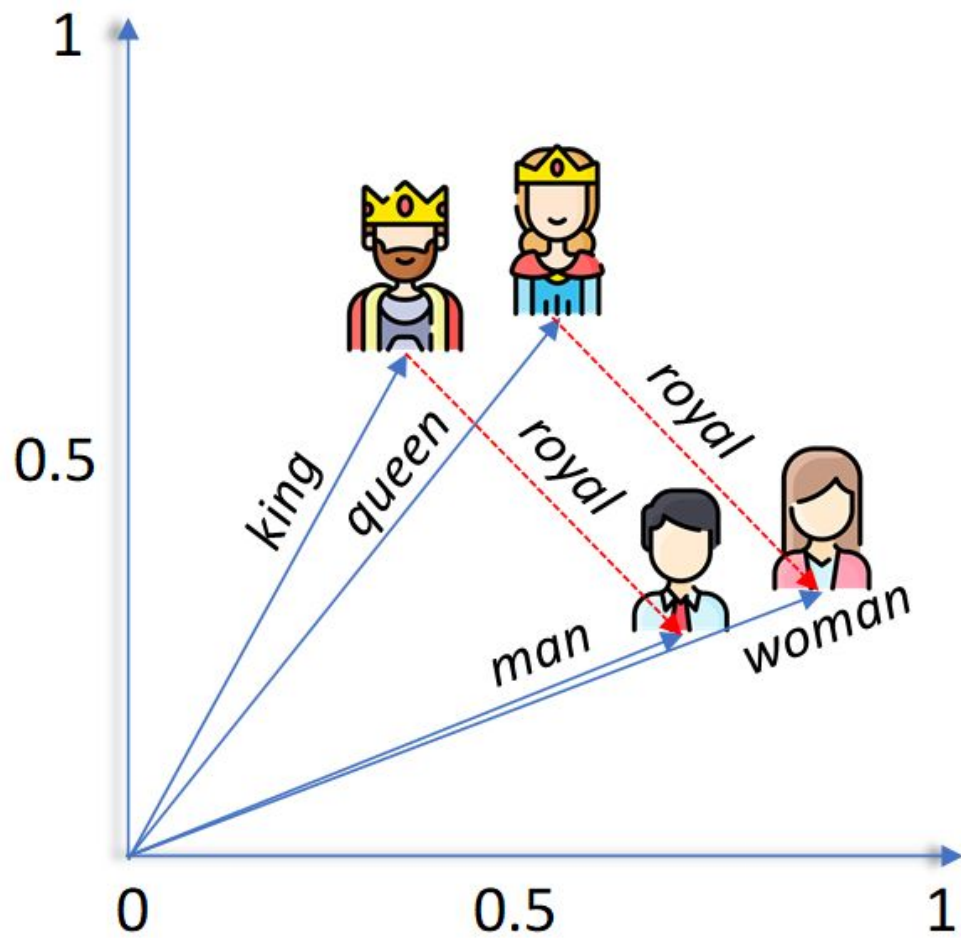
OUTPUT



LOGIC



Probleem:
Taal \neq Wiskunde



3. BERT

1 - Semi-supervised training on large amounts of text (books, wikipedia..etc).

The model is trained on a certain task that enables it to grasp patterns in language. By the end of the training process, BERT has language-processing abilities capable of empowering many models we later need to build and train in a supervised way.

Semi-supervised Learning Step

Model:



Dataset:



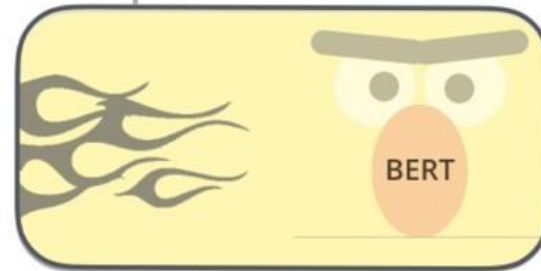
Objective:

Predict the masked word
(language modeling)

2 - Supervised training on a specific task with a labeled dataset.

Supervised Learning Step

Model:
(pre-trained
in step #1)



Classifier

75% Spam
25% Not Spam

Dataset:

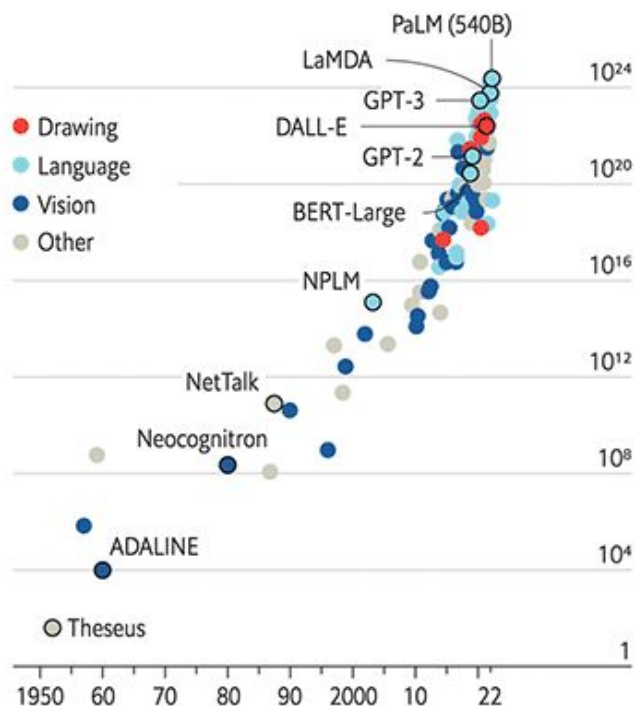
Email message	Class
Buy these pills	Spam
Win cash prizes	Spam
Dear Mr. Atreides, please find attached...	Not Spam

4. GPT-X

The blessings of scale

AI training runs, estimated computing resources used

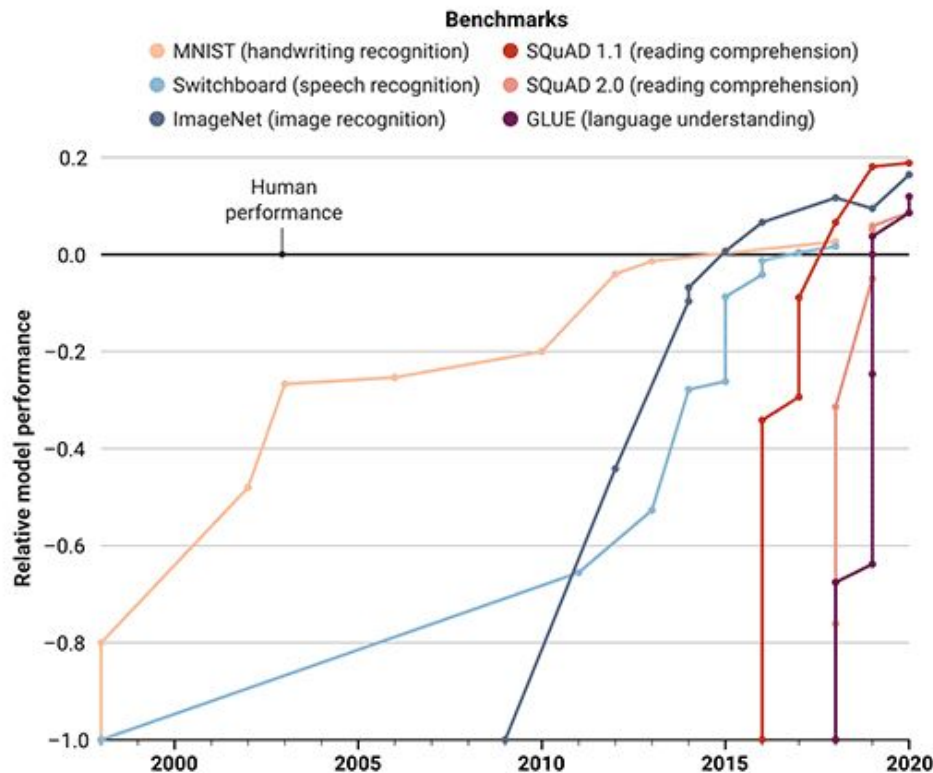
Floating-point operations, selected systems, by type, log scale



Sources: "Compute trends across three eras of machine learning", by J. Sevilla et al., arXiv, 2022; Our World in Data

Quick learners

The speed at which artificial intelligence models master benchmarks and surpass human baselines is accelerating. But they often fall short in the real world.



(GRAPHIC) K. FRANKLIN/SCIENCE; (DATA) D. KIELA ET AL., DYNABENCH: RETHINKING BENCHMARKING IN NLP, DOI:10.48550/ARXIV.2104.14337

5. ChatGPT

Step 1

Collect demonstration data, and train a supervised policy.

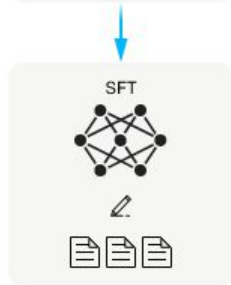
A prompt is sampled from our prompt dataset.



A labeler demonstrates the desired output behavior.



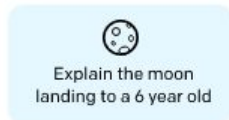
This data is used to fine-tune GPT-3 with supervised learning.



Step 2

Collect comparison data, and train a reward model.

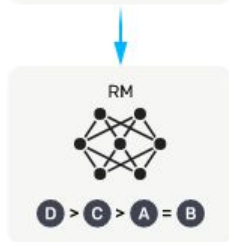
A prompt and several model outputs are sampled.



A labeler ranks the outputs from best to worst.



This data is used to train our reward model.



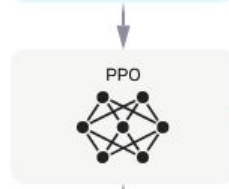
Step 3

Optimize a policy against the reward model using reinforcement learning.

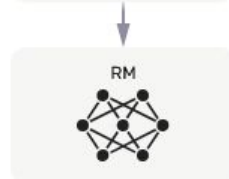
A new prompt is sampled from the dataset.



The policy generates an output.



The reward model calculates a reward for the output.



The reward is used to update the policy using PPO.

