

II. State of AI 4/2023



Włodzisław Duch

Neurocognitive Laboratory, Center for Modern Interdisciplinary Technologies,
& Dept. of Informatics, Faculty of Physics, Astronomy & Informatics,
Nicolaus Copernicus University, Toruń, Poland

Search: Wlodzislaw Duch

Klub Nauki i Biznesu Dell Technologies, Warszawa 27.03.2023

AI at the warp speed



Part I.

1. Short history: 3 AI waves, from GOFAI to GANs and LLMs.
2. Recently: superhuman AI results in many applications.
3. Yesterday: transformers, foundational models, language/vision.
4. Today: tools for Artificial General Intelligence.

Part II. Surprises!

1. Emergence and sparks of AGI in GPT-4.
2. State-of-the-art and beyond.
3. AI minds and human brains.

News in [my YouTube ML](#) library, and in my [Flipboard](#).

Recently: superhuman AI

Superhuman AI



Reasoning: 1997–Deep Blue wins in chess;
2016 –AlphaGo wins in Go; 2017 Alpha GoZero 100:0.

Open Games: 2017–Poker, Dota 2; 2019-Starcraft II,
2022 Stratego, Diplomacy – what is left?

Perception: speech, vision, recognition of faces, images,
personality traits, political and other preferences ...

Robotics: 2020 Atlas robot (Boston Dynamics) backflip
and parkour, autonomous vehicles, airplanes.

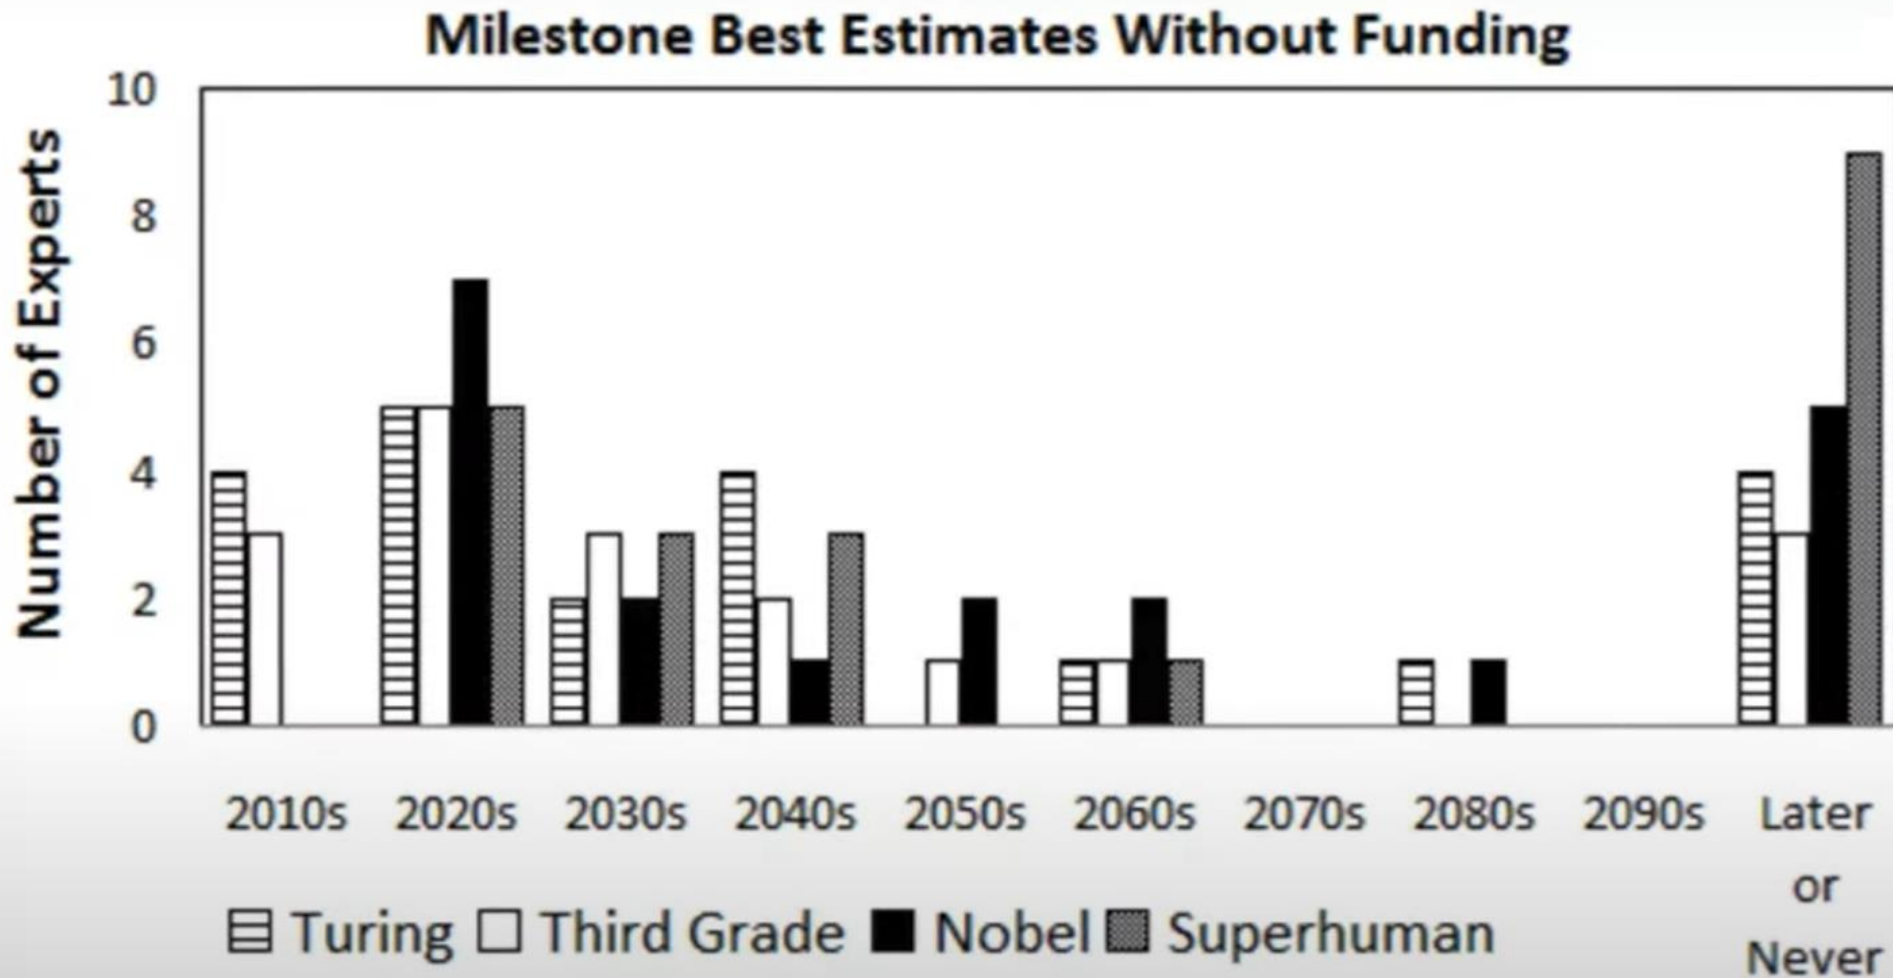
Automation of science: 2015-AI uncovers genetic and
signaling pathways of flatworm regeneration.
2020 AlphaFold 2, now 600 M protein structures.

Creativity and imagination: DeepArt, Midjourney,
Dall-E, AIVA and music composers, GAN revolution.

Language: 2011–IBM Watson wins in Jeopardy;
2018–Watson Debater wins with professionals.
2020: BERT answers questions from SQuAD database.

Cyborgization: BCI, brain optimization, coming?

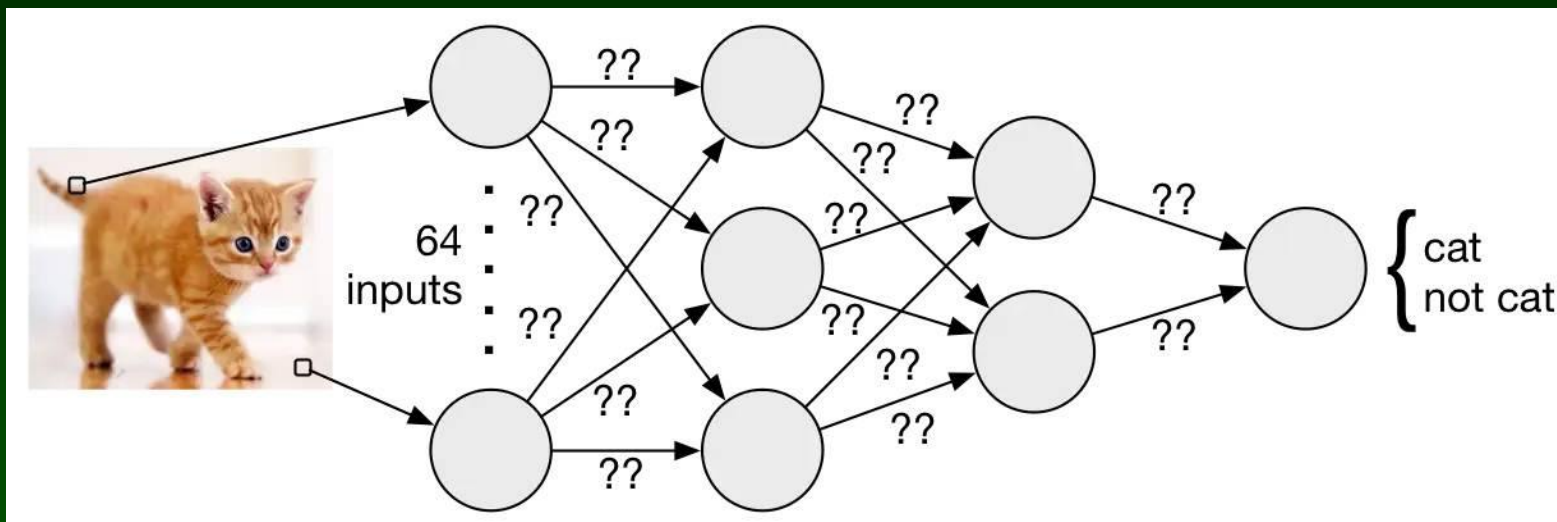
AGI - when?



In 2009 many people thought that superhuman milestone will never be reached.
[AGI-09 Survey – AI Impacts](#)

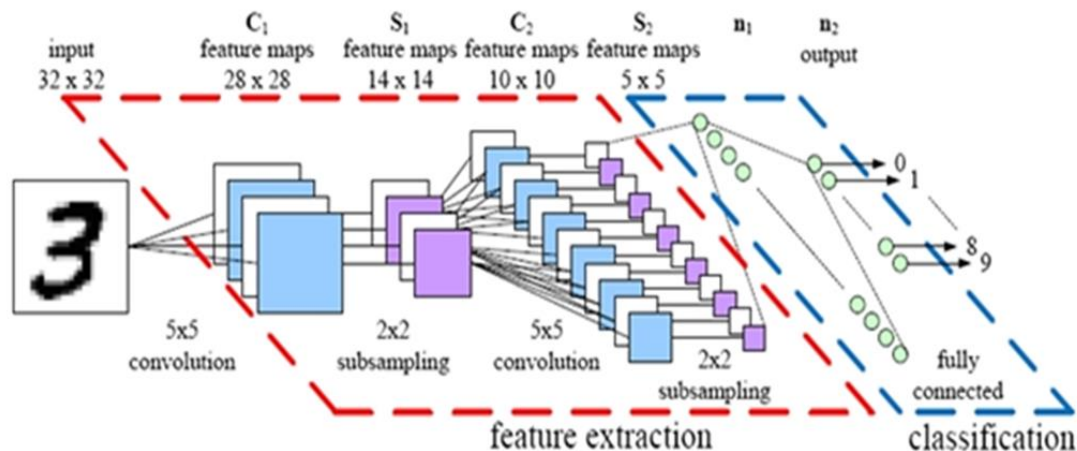
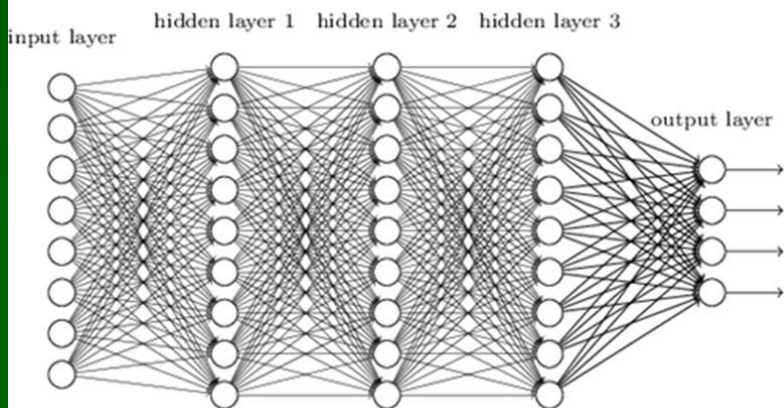
Yesterday: transformers

Neural classifiers

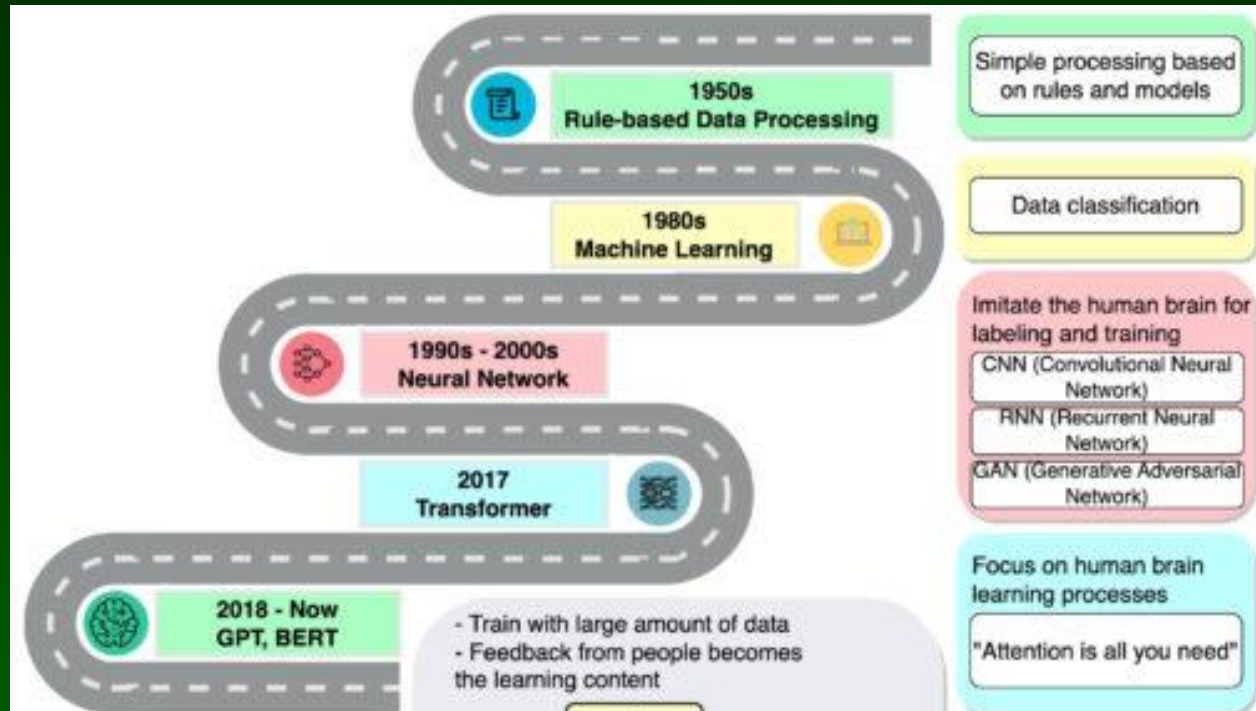


Words, image patches => networks with adjustable parameters
=> training to recognize patterns => object classification, diagnosis.

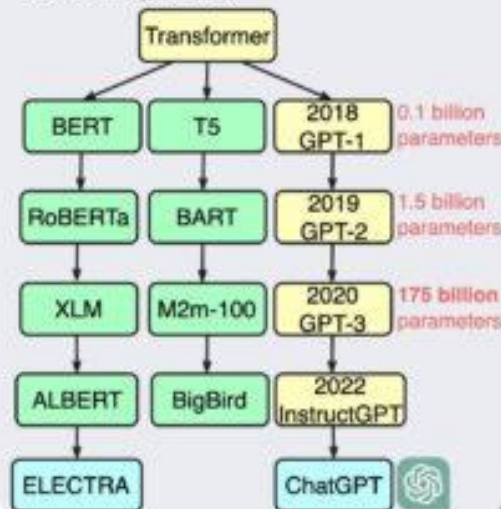
Deep neural network



LLM timeline



- Train with large amount of data
- Feedback from people becomes the learning content

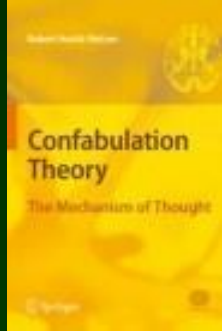


+ trillion parameters!





Language algorithms



Language models: relation of words in complex network structures. In 2018, to gain a general-purpose “language understanding”, Google created BERT, model pre-trained on a very large text corpus.

- **Bidirectional Encoder Representations from Transformers (BERT).** Transformer-based machine learning technique for (NLP) pre-training.
- English-language BERT: two networks, smaller 110M parameters, larger model with 340M parameters in 24-layers; trained on the BooksCorpus with 800M words, and Wikipedia with 2,500M words. In 2019 BERT worked already in 70 languages.
- BERT model was then fine-tuned for specific NLP tasks such as question answering or semantic information retrieval. Many smaller pre-trained open software models were published in GitHub repository.
- The network learns to predict masked words (images, signals):
Input: the man went to the [MASK1]. He bought a [MASK2] of milk.
Labels: [MASK1] = store; [MASK2] = gallon.
As in Hecht-Nielsen, Confabulation Theory (2007).

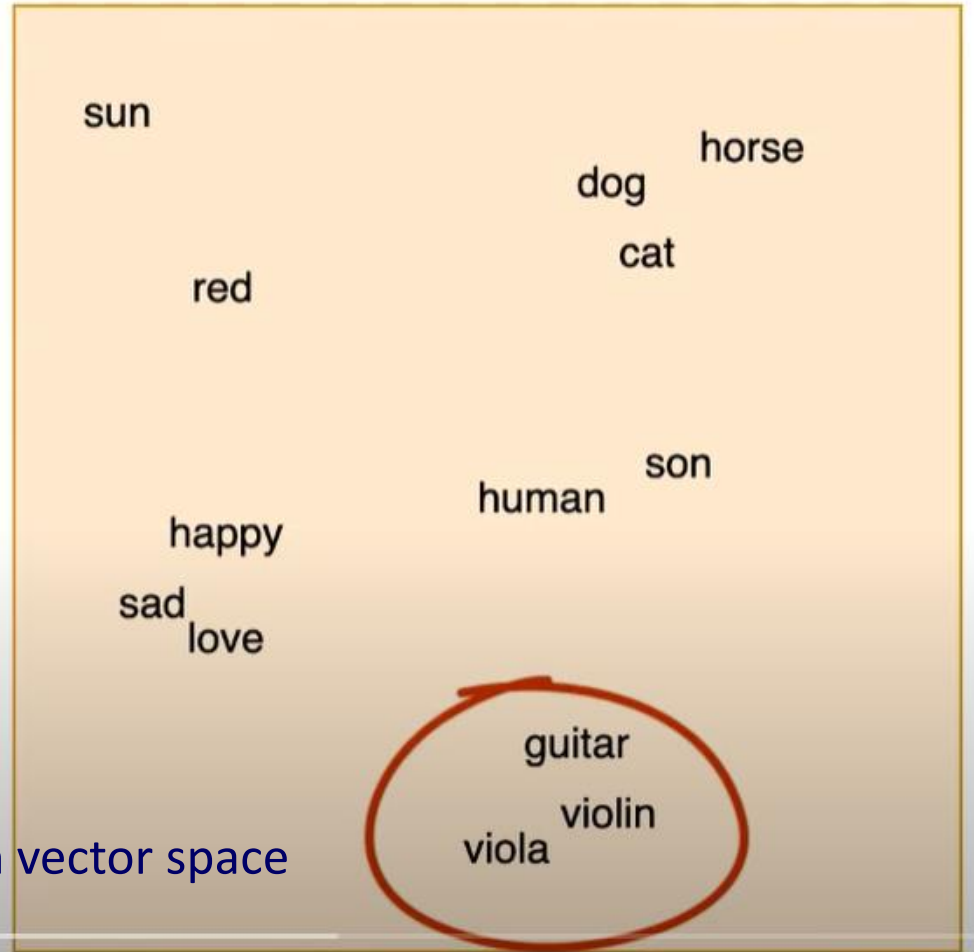
Embeddings

Words => vectors, reflecting their similarity and positions in sentences.

Input Embedding



learns
→



Transformers

Vaswani et al.(2017). *Attention Is All You Need*. arXiv

Attention: given a sequence of tokens (words, image patches), how relevant is each input token to other tokens?

Attention vectors capture contextual relations between words in sentences. For example:

Input: sentence in English;

Output: sentence in Polish language.

Google BERT has used this.

[Simple intro on Youtube.](#)

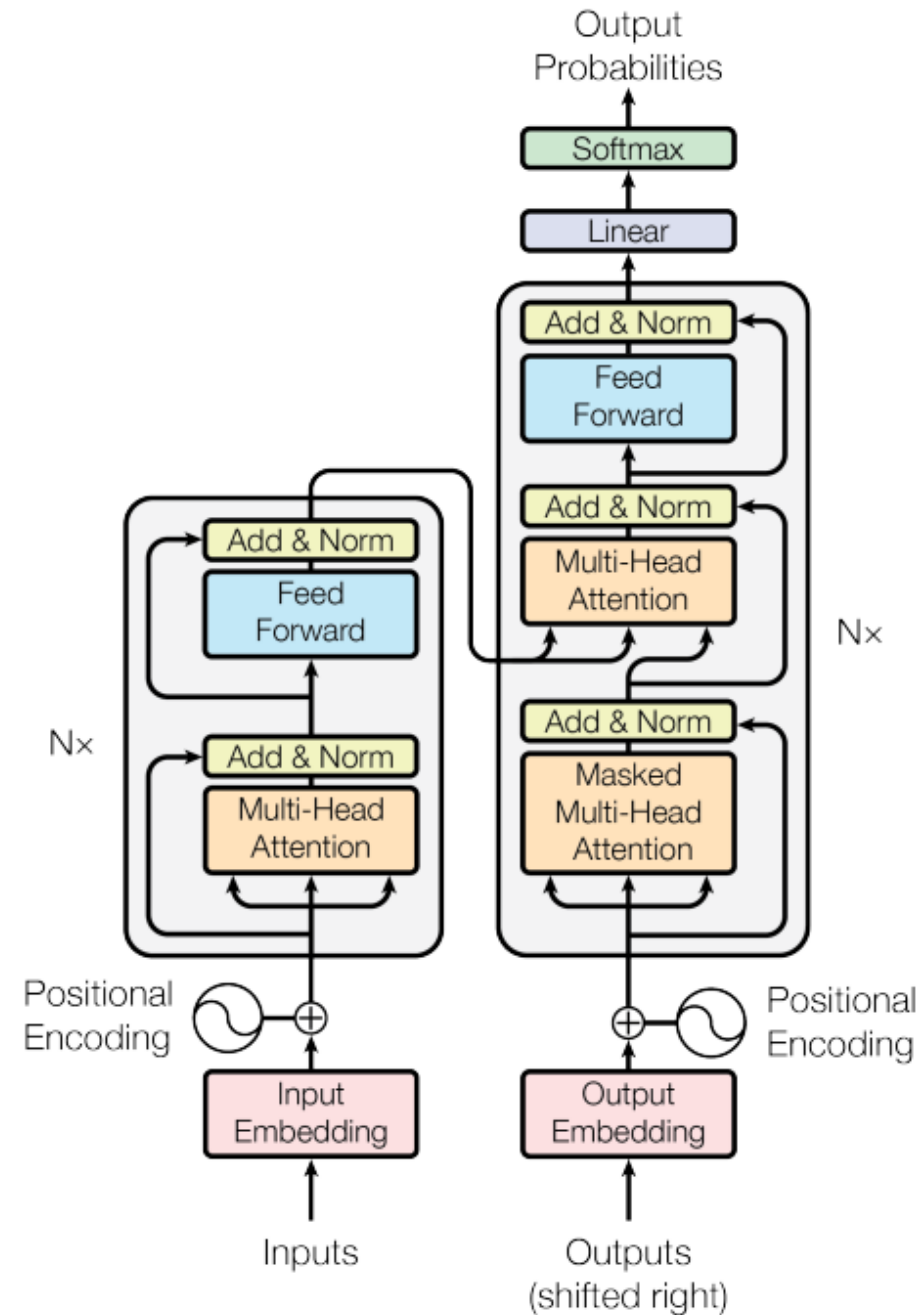
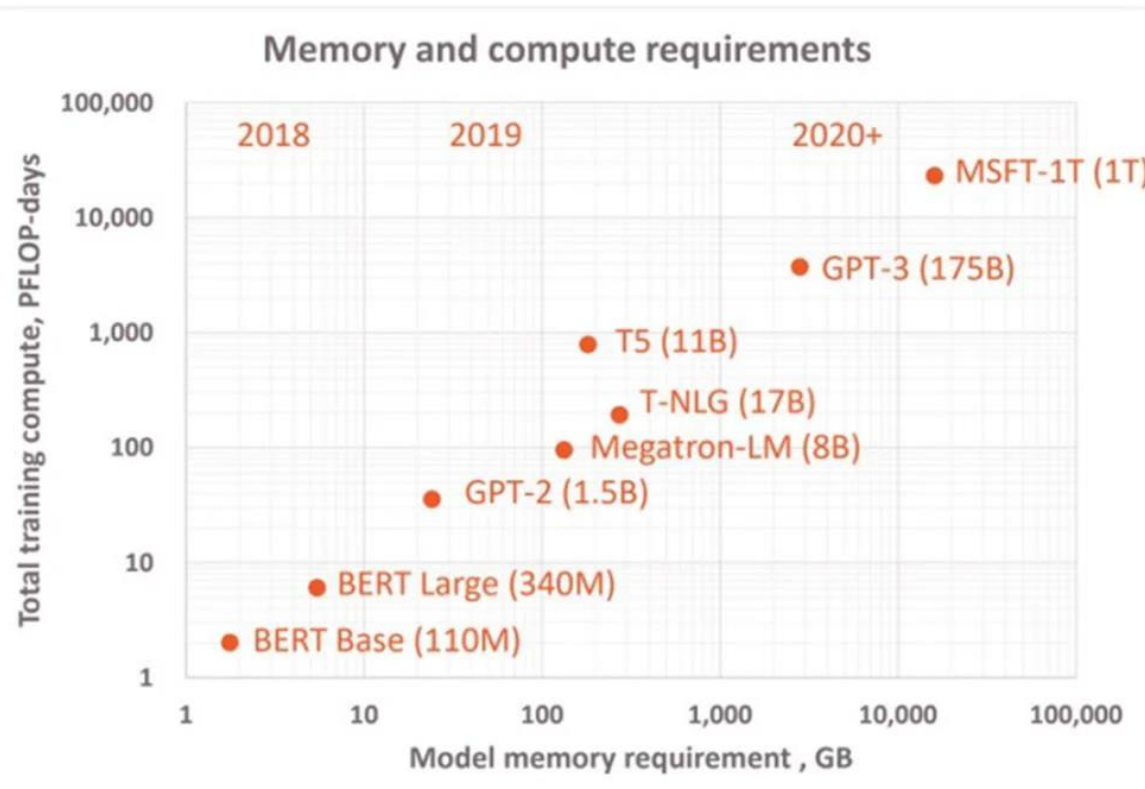


Figure 1: The Transformer - model architecture.

Acceleration ...

Exponential Growth of Neural Networks



1000x larger models
1000x more compute
In just 2 years

Today, GPT-3 with 175 billion params trained on 1024 GPUs for 4 months. OpenAI

MS+Nvidia MLM, MSFT models
Tomorrow, **multi-Trillion** parameter models and beyond.

Nanotechnologies for ML training: in 10 years from 1 Pflop to 1 billion Petaflops!

Large models provide foundations for general knowledge. Gato has only 1.2 B parameters. Smaller LLM may also be efficient (mBERT, LLaMa).

NLP supermodels

OpenAI GPT-3 model has 175 B parameters! One can use it on OpenAI server. First-of-its-kind API can be applied to any language task, and serves millions of production requests each day. GPT-4 has 170 trillion parameters?

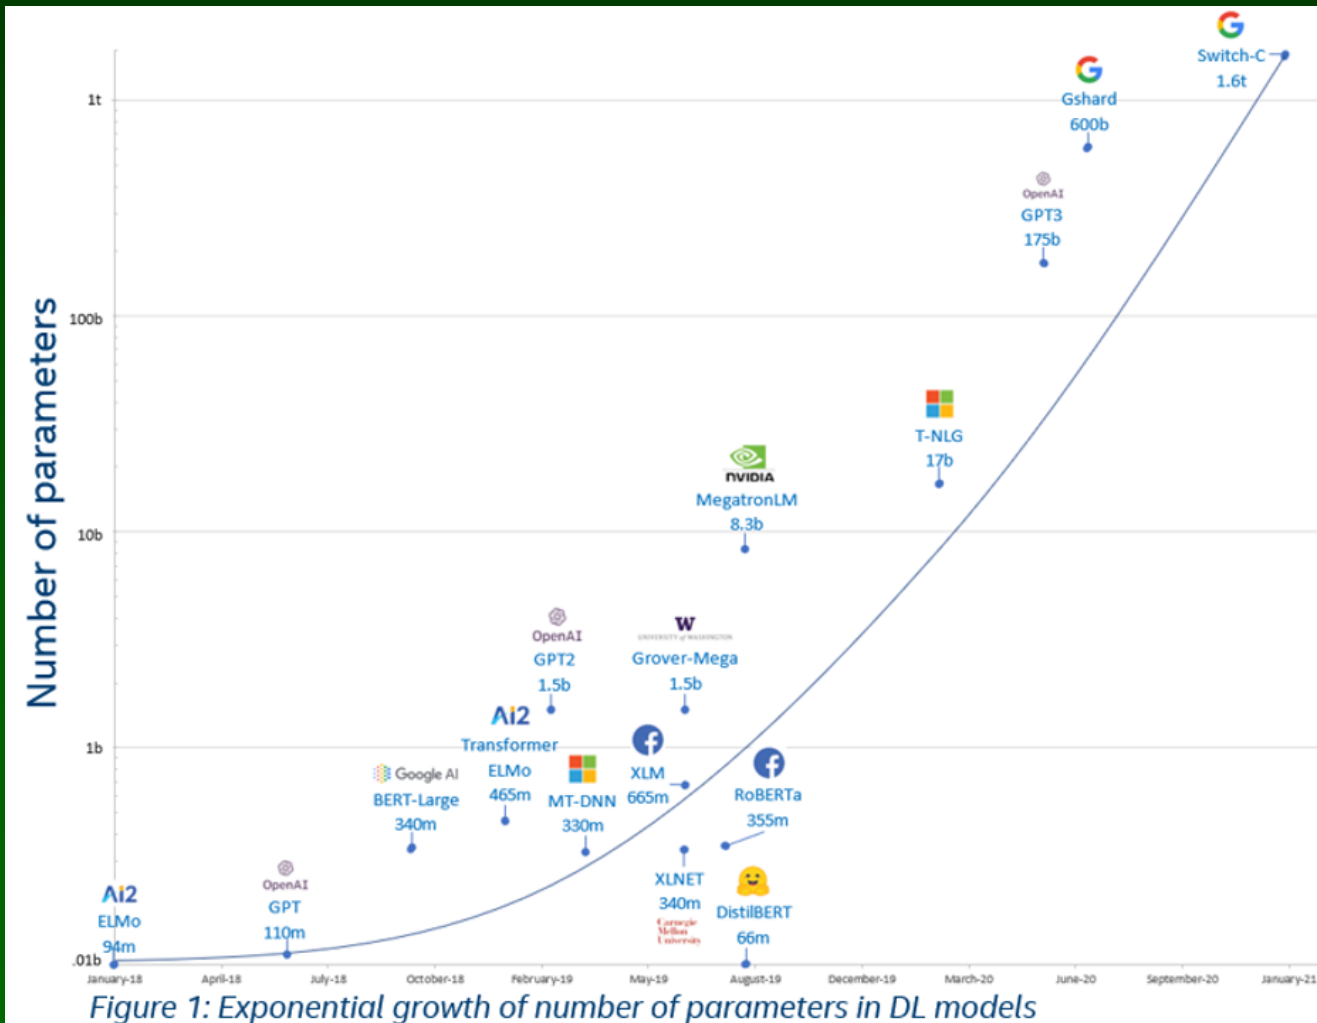


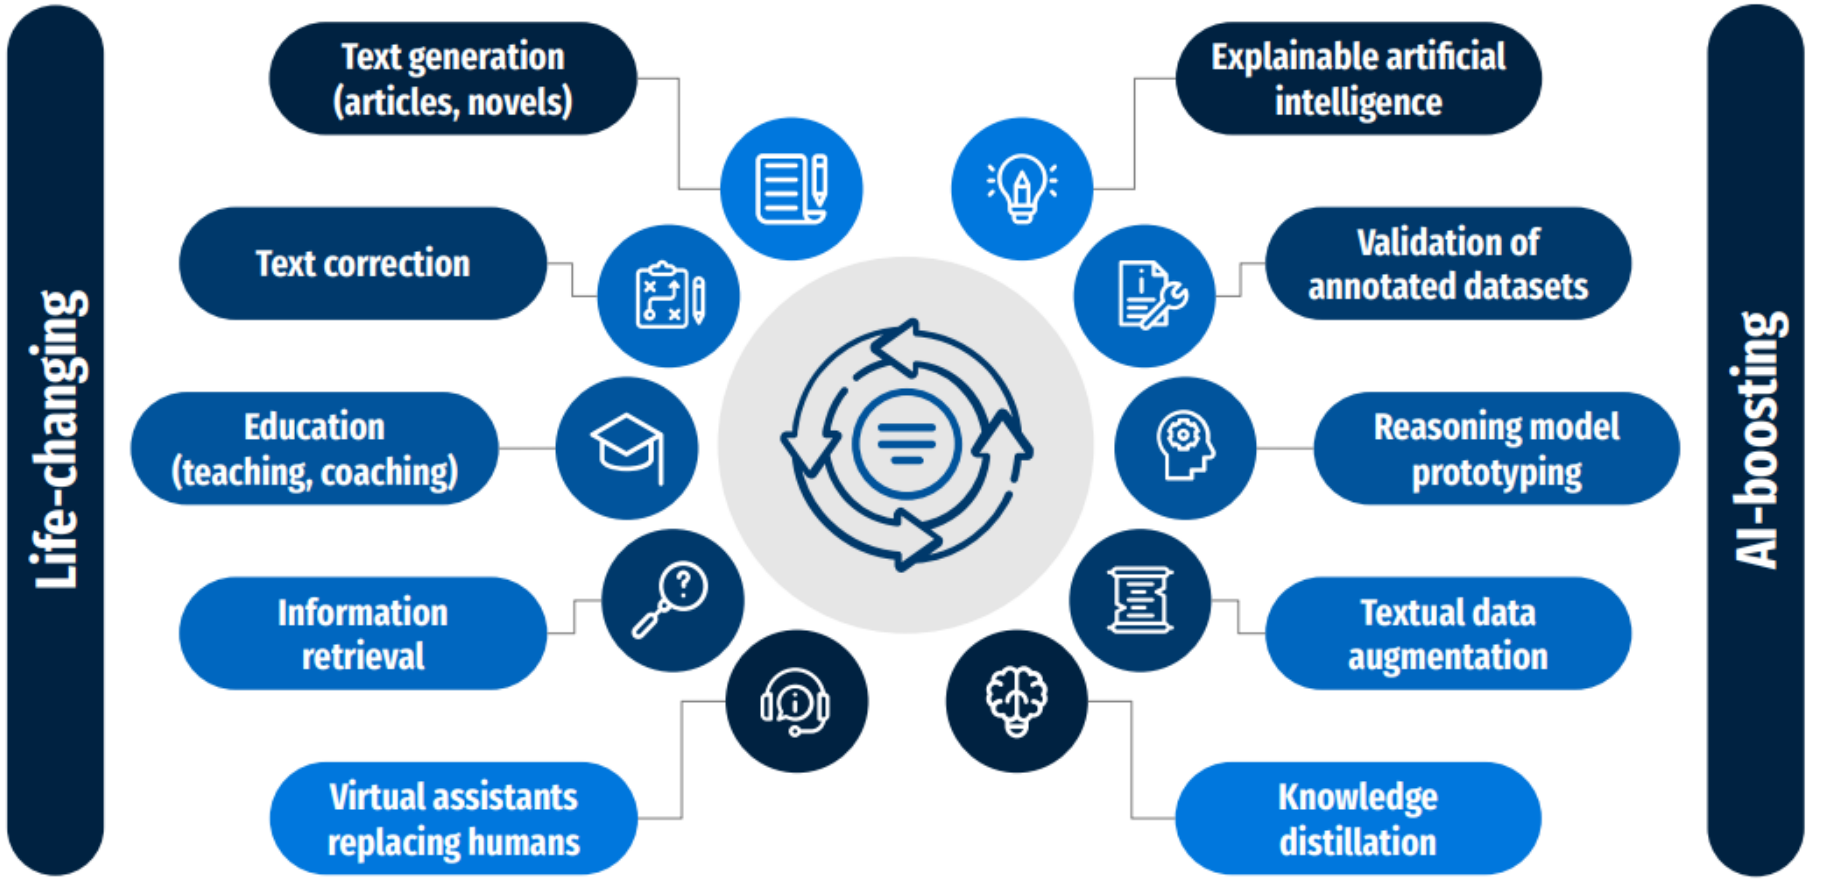
Figure 1: Exponential growth of number of parameters in DL models

WuDao ~1.75 trillion param.

Between 2018-21 LLM complexity increased 1 M times!

From 2018 to Jan. 2021

Prospects for ChatGPT applications



ChatPDF

Talk Books

Browse passages from books using experimental AI

[Learn more](#)



Not a traditional search

Use this demo as a creativity tool to explore ideas and discover books by getting quotes that respond to your queries.



Use natural language

Speaking to it in sentences will often get better results than keywords. That's because the AI is trained on human conversations.



Play with it

Try our sample queries then try your own. Experiment with different wording to see how it changes the results.



Say something to books...



Go!

Galactica trained on science, and Consensus for evidence-based answers.

Large Language Models

Next-gen (existing) applications


Product & customer interaction / management



 viable  chatdesk  Quickchat

 Nevermaps  ActiveChat  exceed
by GENESYS

 Stateset  Sapling


Personal productivity

 personal.ai

 mem 

 Oogway




Search engine

 YOU  Google


 algolia

Emerging net-new applications

Application synthesis

 Adept  CODEGEN 

Data analyst productivity

 veezoo  AI 2sql  cogram

Developer productivity

 warp  tabnine



 GitHub Copilot  ASK JARVIS




 replit  M  KD

New media generation

 FABLE  DALL-E 2  alethea  R

Writing assistant/text generation




 AI21labs  Jasper  Snazzy AI

 PR Guy  copy.ai  Scalenut

 LAVENDER  YOU Write 

 anyword  Simplified 

 copysmith  copymatic 

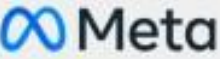

 LONGSHOT  Rytr  Writesonic
Previously ProWritingAid

Infrastructure



Model /builders providers - Big Tech

 Microsoft

 Google  DeepMind

 Meta 

Model providers/builders - Startups

 OpenAI  co:here

 Hugging Face  BigScience

 AI21labs  LightOn
We bring Light to AI

 ANTHROPIC 

Accessible specialized AI chips

 NVIDIA  GRAPHCORE

 Google  LightOn
We bring Light to AI

Other tooling

 Humanloop  anyscale

GPT-3 as philosopher



Eric Schwitzgebel, David Schwitzgebel, Anna Strasser, Creating a Large Language Model of a Philosopher, [arXiv:2302.01339](https://arxiv.org/abs/2302.01339)

“Can large language models be trained to produce philosophical texts that are difficult to distinguish from texts produced by human philosophers?

To address this question, we fine-tuned OpenAI's GPT-3 with the works of philosopher Daniel C. Dennett as additional training data.

To explore the Dennett model, **we asked the real Dennett ten philosophical questions** and then posed the same questions to the language model, collecting 4 responses for each question without cherry-picking.

We recruited 425 participants to distinguish Dennett's answer from ChatGPT. Experts on Dennett's work (N = 25) **succeeded 51% of the time**, above the chance rate of 20% but short of our hypothesized rate of 80% correct.

For 2 of the 10 questions, the language model produced at least one answer that experts selected more frequently than Dennett's own answer.

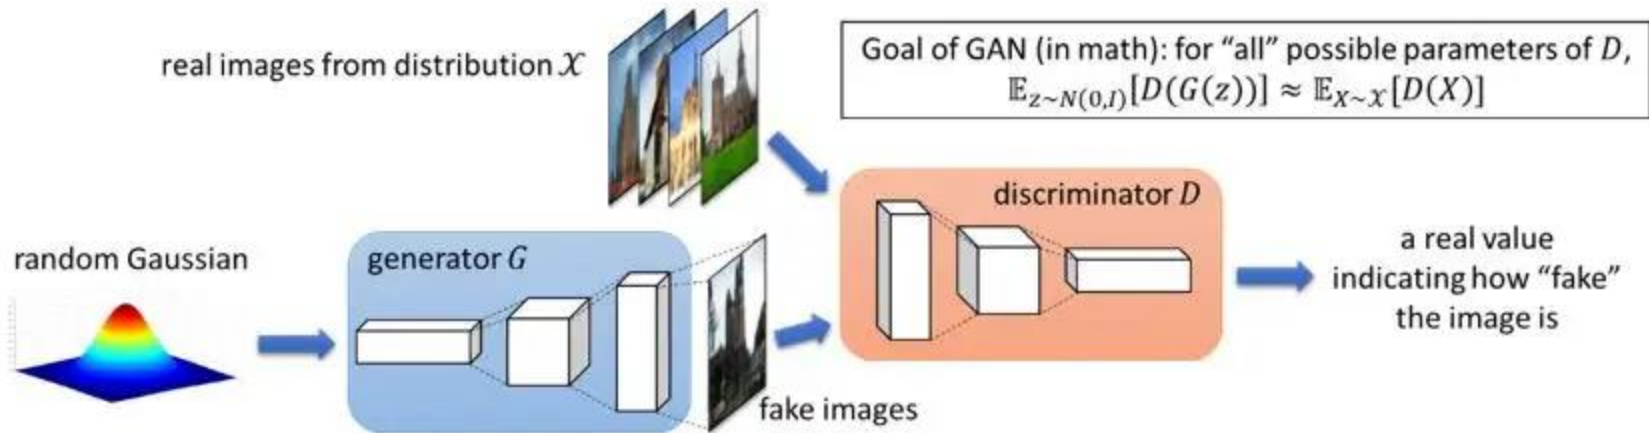
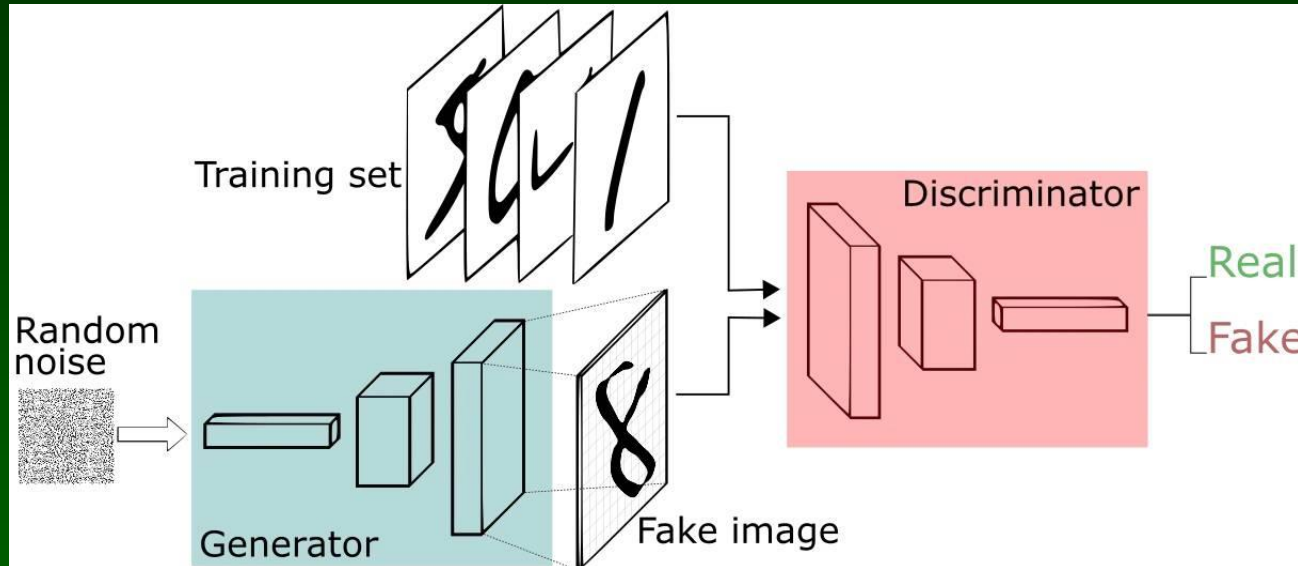
Philosophy blog readers (N = 302) performed similarly to the experts.

Ordinary research participants (N = 98) were near chance distinguishing GPT-3's responses from those of an "actual human philosopher".

Is Dennett intelligent? If we agree, then GPT-3 is also intelligent.

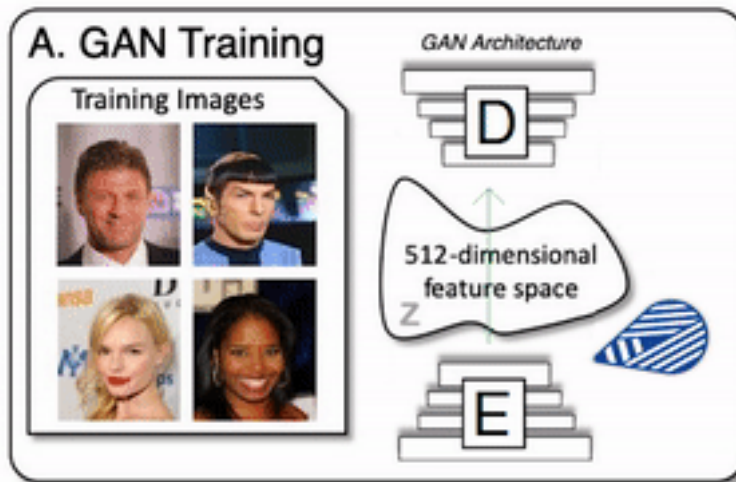
Images: Generative Networks

Generate output from noise. Discriminator recognizes that it does not resemble real patterns. Correct parameters of generator, repeat. GANs capture the essence!



Generative Networks

Generate output from noise. Discriminator recognizes that it does not resemble real patterns. Correct parameters of generator, repeat. Capture the essence!
Latent space does not contain training data (images), but parameters that help to recreate structures similar to those that discriminator recognizes as correct.



Vision-language models

Vision-Language Pre-Trained Models (VL-PTMs), convergence of language, vision, and multimodal pretraining => general-purpose foundation models can handle be easily adapted to multiple diverse tasks with zero-shot learning.



koala bears



motorcycles

Vision-language generative models

Dall-E2, Craiyon, Imagen, Midjourney, Nightcafe, Artbreeder, Hotpot AI, Deep Dream Generator, Deep AI Text to Image, Generative Engine, Starry AI, My Heritage ... PromptBase is at the center of the new trade in prompts for generating specific imagery by image generators, a kind of meta-art market.

The image shows a complex network graph with nodes and edges. Nodes are labeled with terms like 'environment change', 'event', 'time', 'behavior', 'man', 'mind', 'deal', 'information', 'great', 'idea', 'question', and 'ecology'. A red line connects 'environment change' and 'information'. A large red text box is overlaid on the graph, stating: "Based on the structural gap, GPT-3 generates a research question that would bridge these topics together:". Below this, a smaller text box lists generated questions: "The following questions were generated to bridge the gap between change, information, environment and question, thought, deal: What is the physical environment's impact on our thoughts and ideas? How does the environment change over time?". To the right, a control panel includes tabs for 'Essence', 'Insight', 'Trends', 'Stats', 'Sentiment', and 'LDA'. It features an 'Action Advice' section with 'Diversify', a 'Structural Gap' section with a list of topics (change, information, environment, question, thought, deal), and a 'Latent Topical Brokers' section with a list of terms (deal, thing, process, order, occur). A 'Help Center' button is visible on the far right.

Based on the structural gap, GPT-3 generates a research question that would bridge these topics together:

The following questions were generated to bridge the gap between change, information, environment and question, thought, deal:

What is the physical environment's impact on our thoughts and ideas?

How does the environment change over time?

Essence Insight Trends Stats
Sentiment LDA

Action Advice: Diversify

Structural Gap (ask a research question that would link these two topics):

change information environment
and
question thought deal

Reveal the Gap Generate a Question ?

Latent Topical Brokers:

deal thing process order occur
?

Back to the Main Topics

network structure: focused

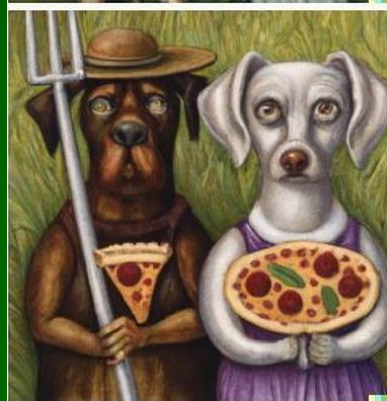
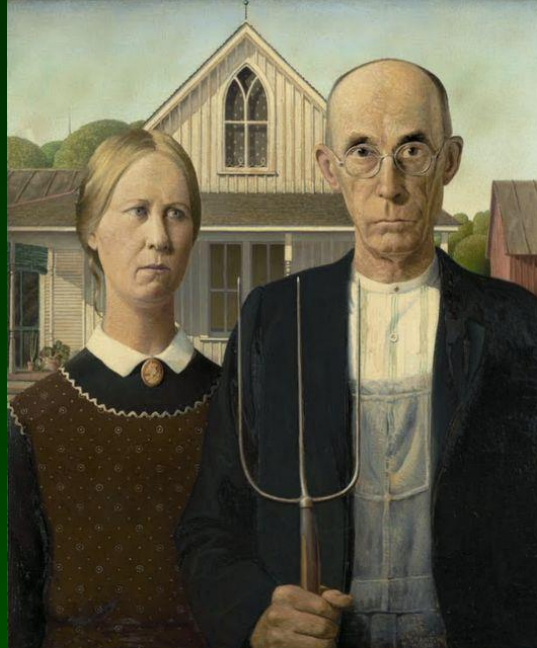
AI imagery

The neural network has billions of parameters, it can combine textual description with images.

These images were created from prompt:

The painting American Gothic, with two dogs holding pepperoni pizza instead of the farmers holding a pitchfork.

Each time program is run another version is created. This technique can create 3D images and video.

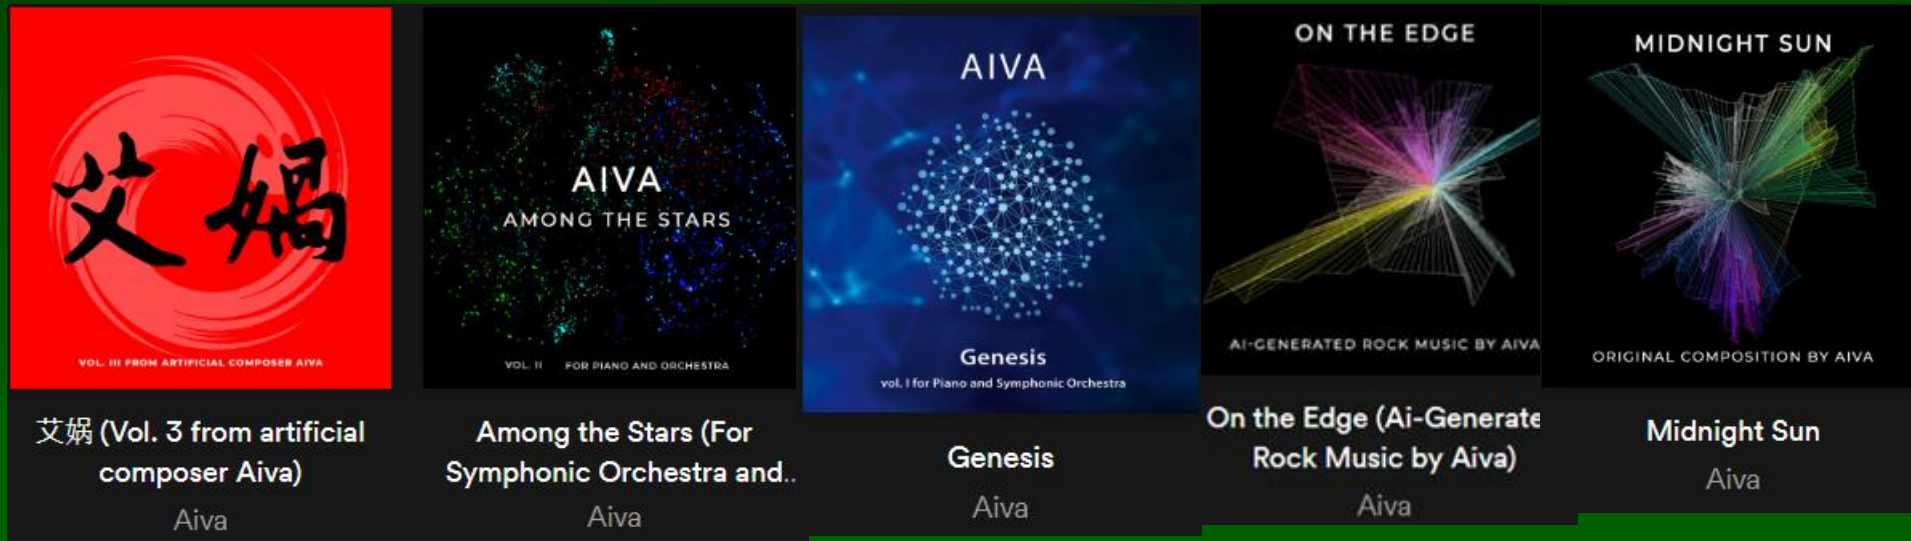


Creativity: AI Virtual Artist

[AIVA](#) – AI Virtual Artist, admitted to [SACEM](#) (Association of Authors, Composers and Music Publishers of France), >1000 compositions.

[AIVA YouTube](#) channel, Youtube „[Letz make it happen](#)“, Op. 23

[SoundCloud channel](#) [Spotify](#) or [Apple](#) channel.



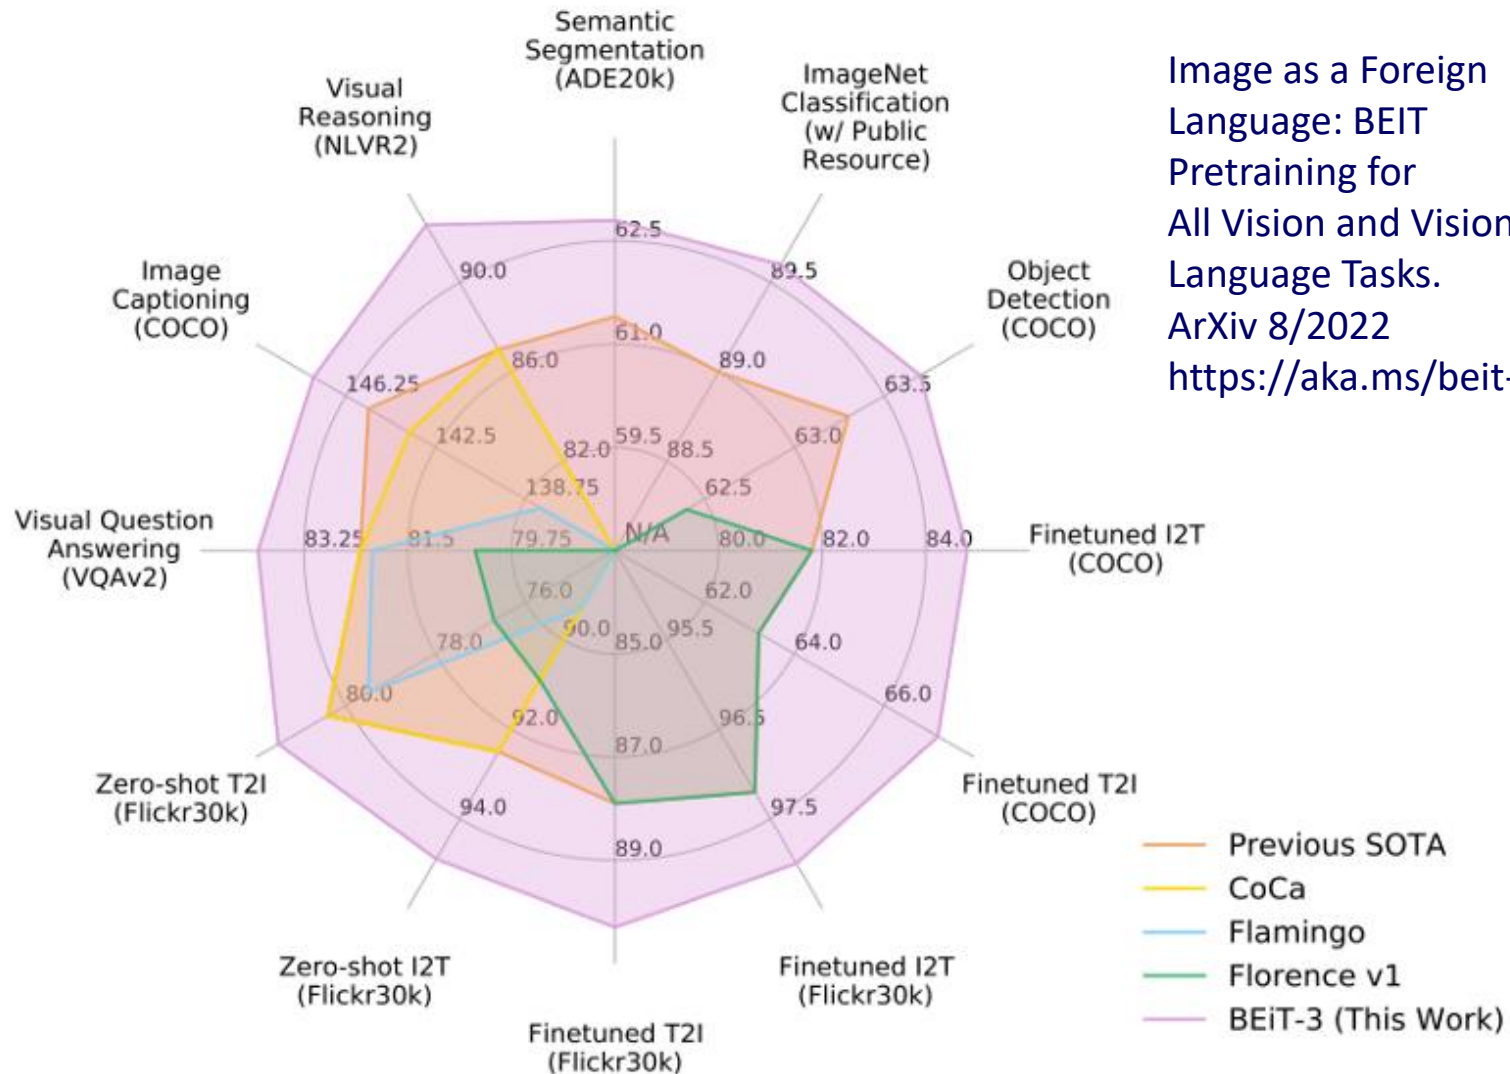
AI completed [Beethoven X Symphony](#) in 2021.

Duch W, [Intuition, Insight, Imagination and Creativity](#).

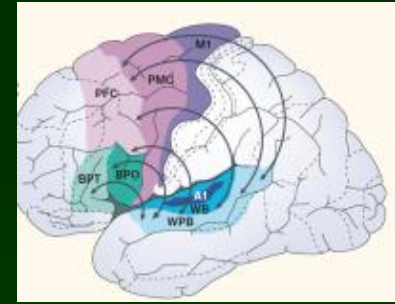
IEEE Computational Intelligence Magazine 2(3), August 2007, pp. 40-52

Vision-language models

MS BEiT-3 (BERT Pretraining of Image Transformers), a general-purpose state-of-the-art multimodal foundation model for vision-language tasks.



Why LLM work?



My computational creativity project (2005):

Model inspired by brain process involved in creating new names.

- make a simplest test for creative thinking – invent novel names;
- create interesting new names for products, capturing their characteristics;
- understand newly invented words that are not in the dictionary.

Assumption: a set of keywords (prompts) **primes the trained** cortex subnetwork.

Prompts: ordered strings of phonemes that activate semantic reps of words.

Spreading activation => context priming + inhibition in the winner-takes-all process leaves only a few semantically related concepts (word meaning).

Creativity = imagination (activation of pre-trained networks) + filtering (competition, associations to existing states)

Prompt=>Imagination: transient patterns of localized activity, activation guided by the strength of connections spreads to associated networks coding concepts.

Filtering: strongest associations based on phonological/semantic similarity survive.

Is this what happens in large language models?

Today: tools for AGI

FOUNDATION MODEL



GATHER DATA AT SCALE

TRAIN FOUNDATION MODEL ONE TIME

EVALUATE MODEL'S PERFORMANCE

FINE-TUNE MODEL FOR MULTIPLE DOWNSTREAM USES

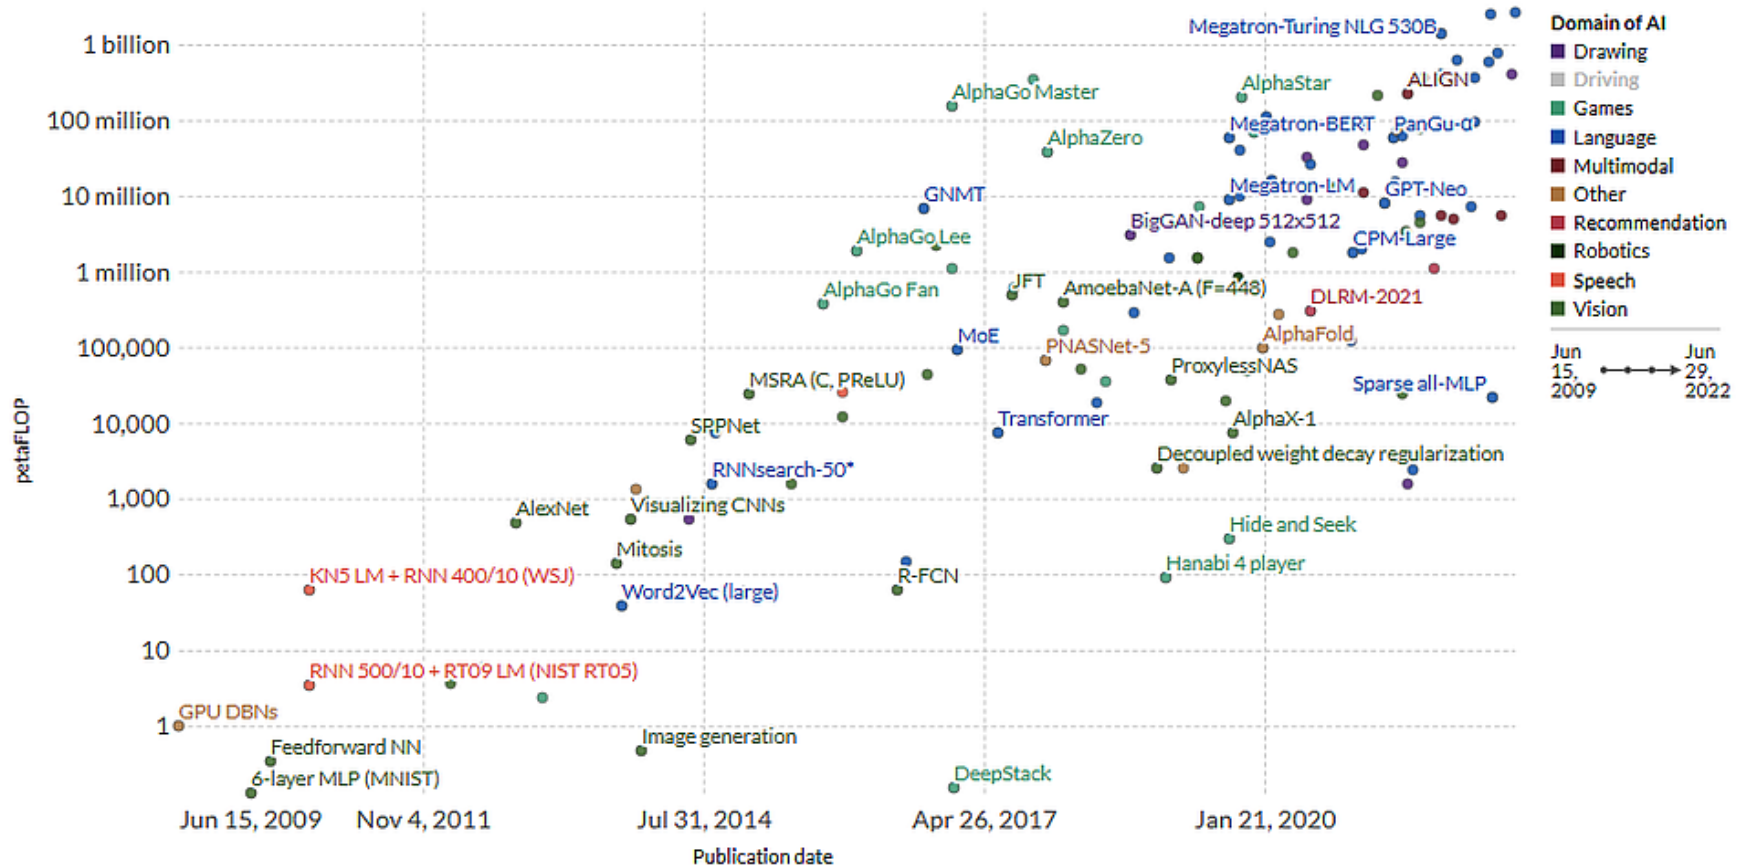
Foundational models

Computation used to train notable AI systems

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

Our World
in Data

LINEAR LOG Select systems



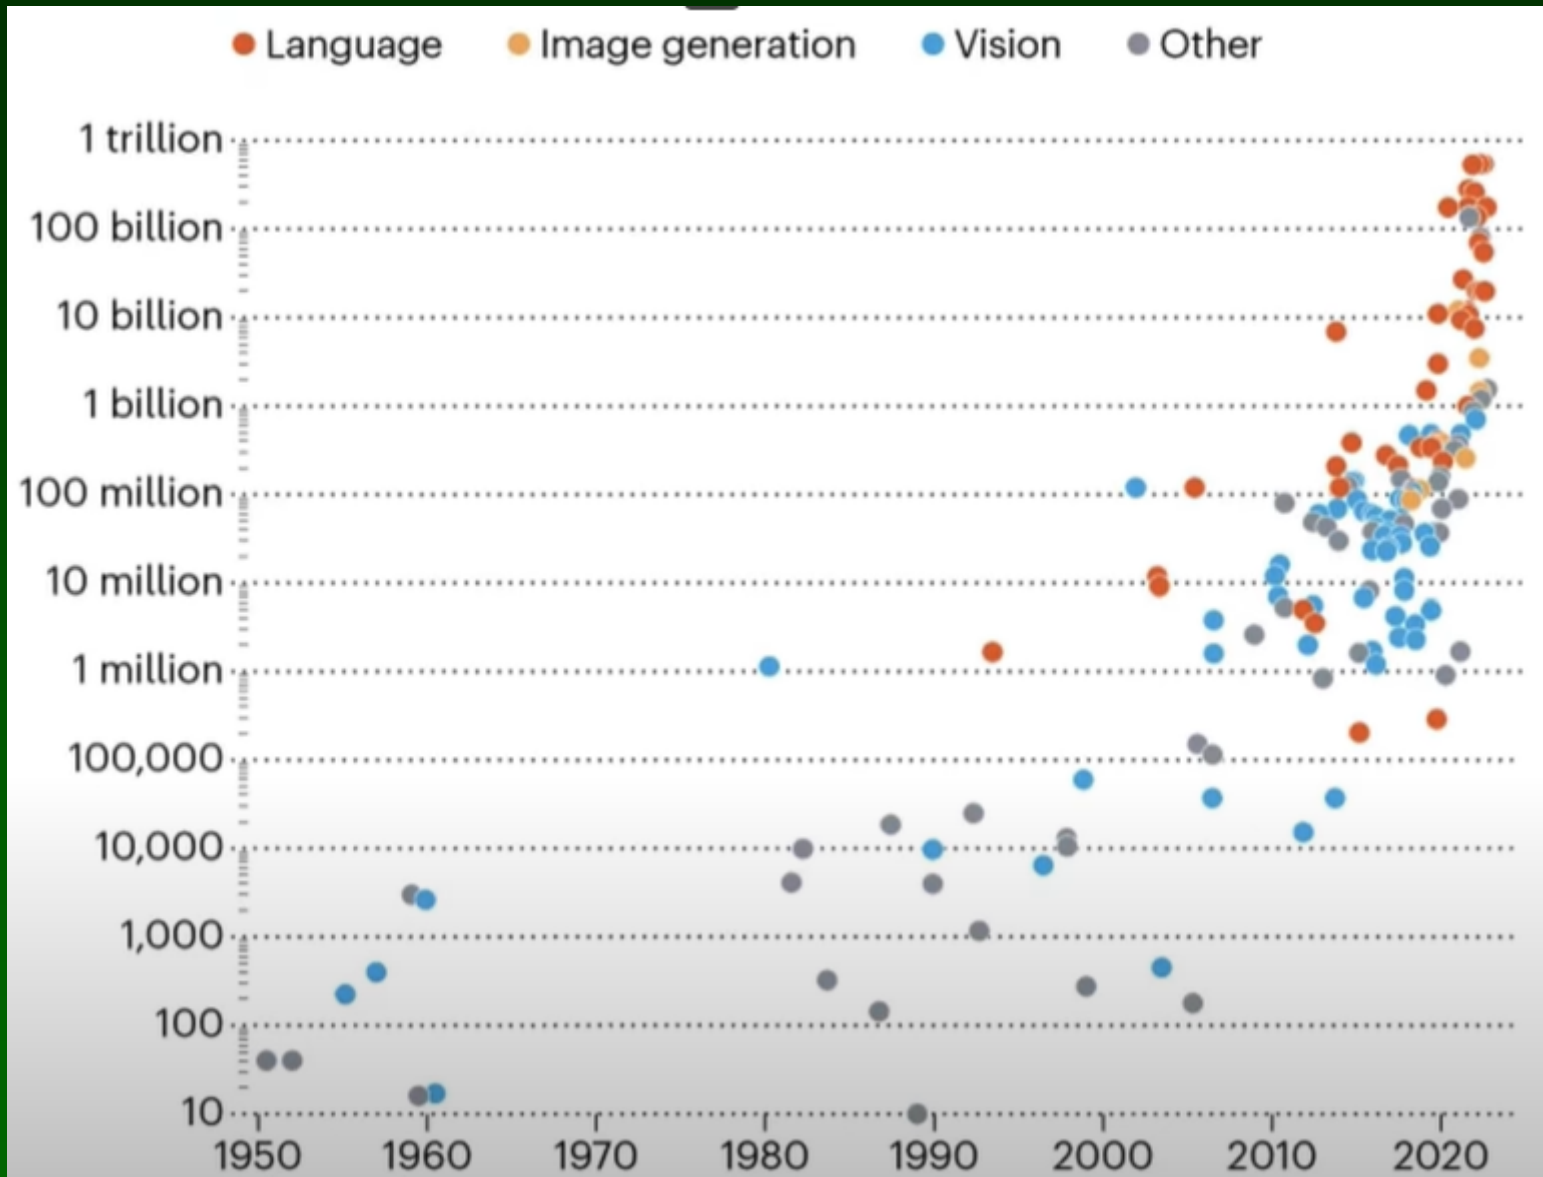
Source: Sevilla et al. (2022)

Note: Computation is estimated by the authors based on published results in the AI literature and comes with some uncertainty. The authors expect the estimates to be correct within a factor of 2.

OurWorldInData.org/technological-change • CC BY

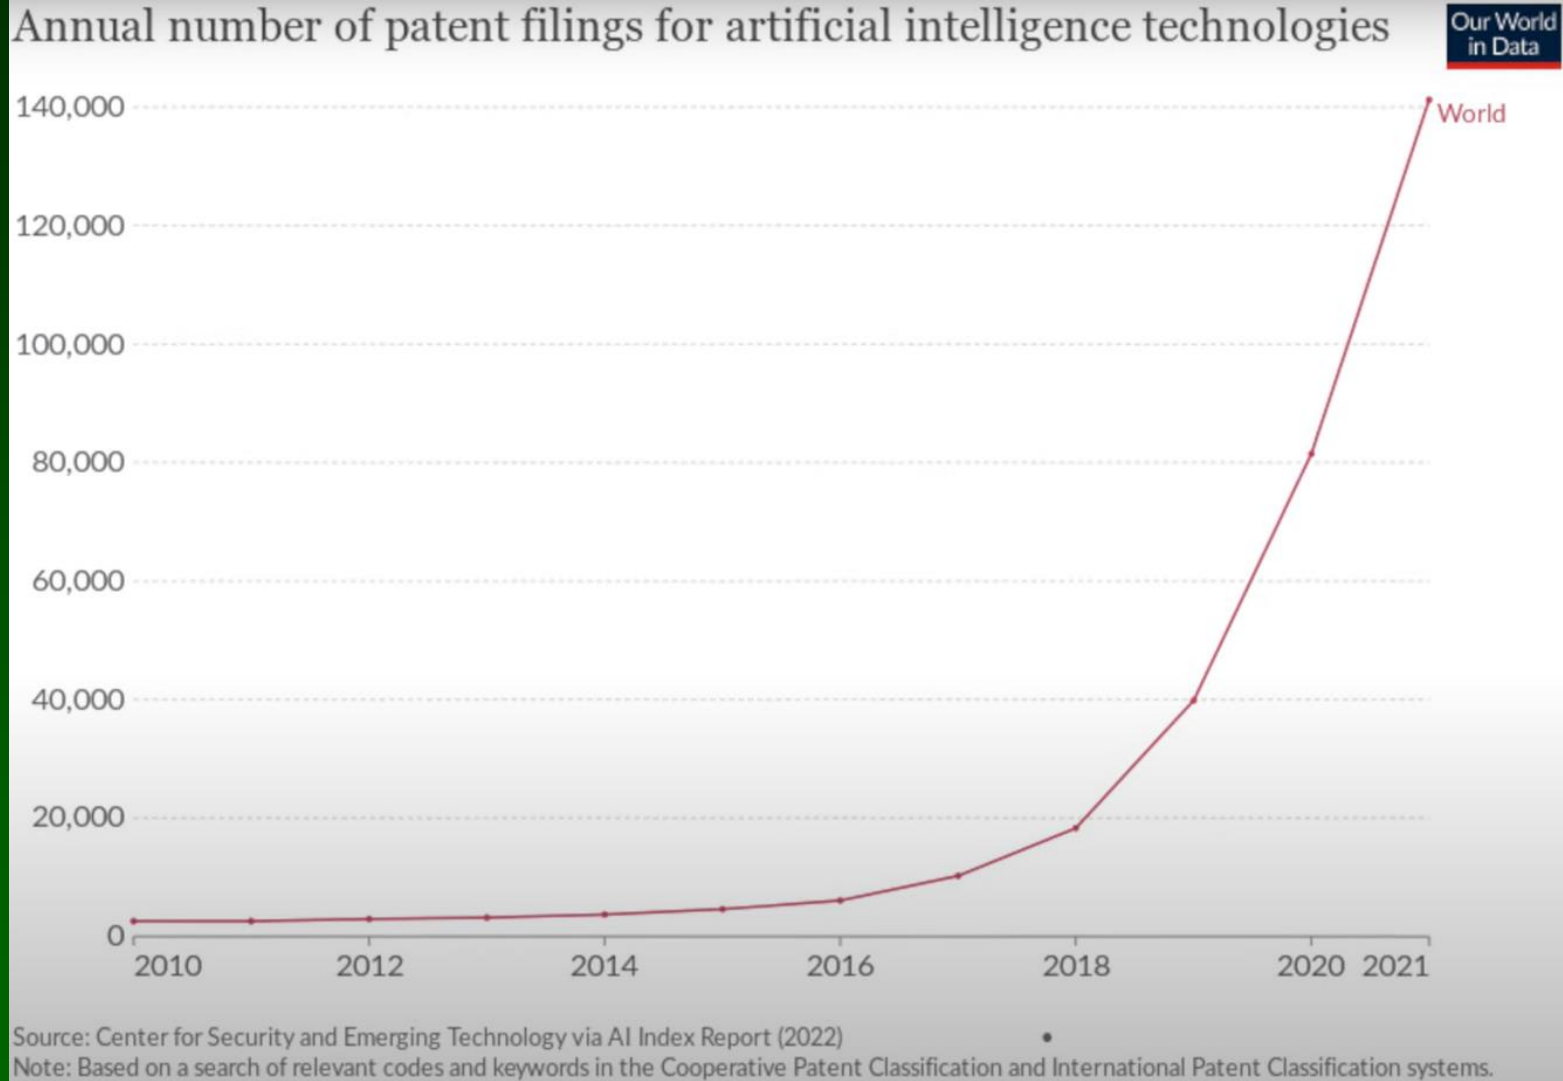
Foundation models => Specific models => Local data models => APIs, Apps.

Exponential growth



Explosion! Exponential number of parameters even in the logarithmic scale!

Patents



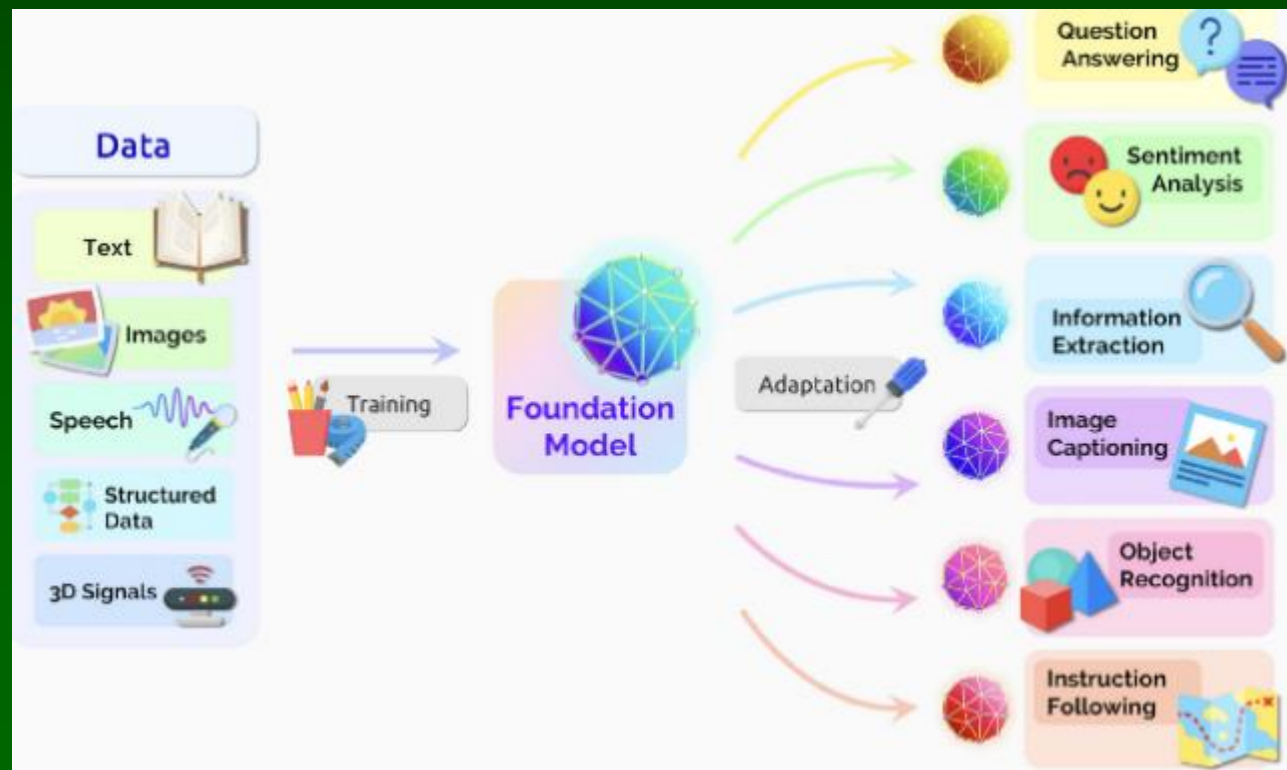
Number of patents grows exponentially.

Multimodal models

Multimodal learning – different types of modalities with different statistical properties, embedded in the same model.

- **Multimodal Affective Computing (MAC)**, sentiment analysis.
- **Natural Language for Visual Reasoning (NLVR)**.
- **Multimodal Machine Translation (MMT)**.
- **Visual Retrieval (VR)** and **Vision-Language Navigation (VLN)**.

Image: [Center for Research on Foundation Models \(CRFM\)](#), [Stanford Institute for Human-Centered Artificial Intelligence \(HAI\)](#)



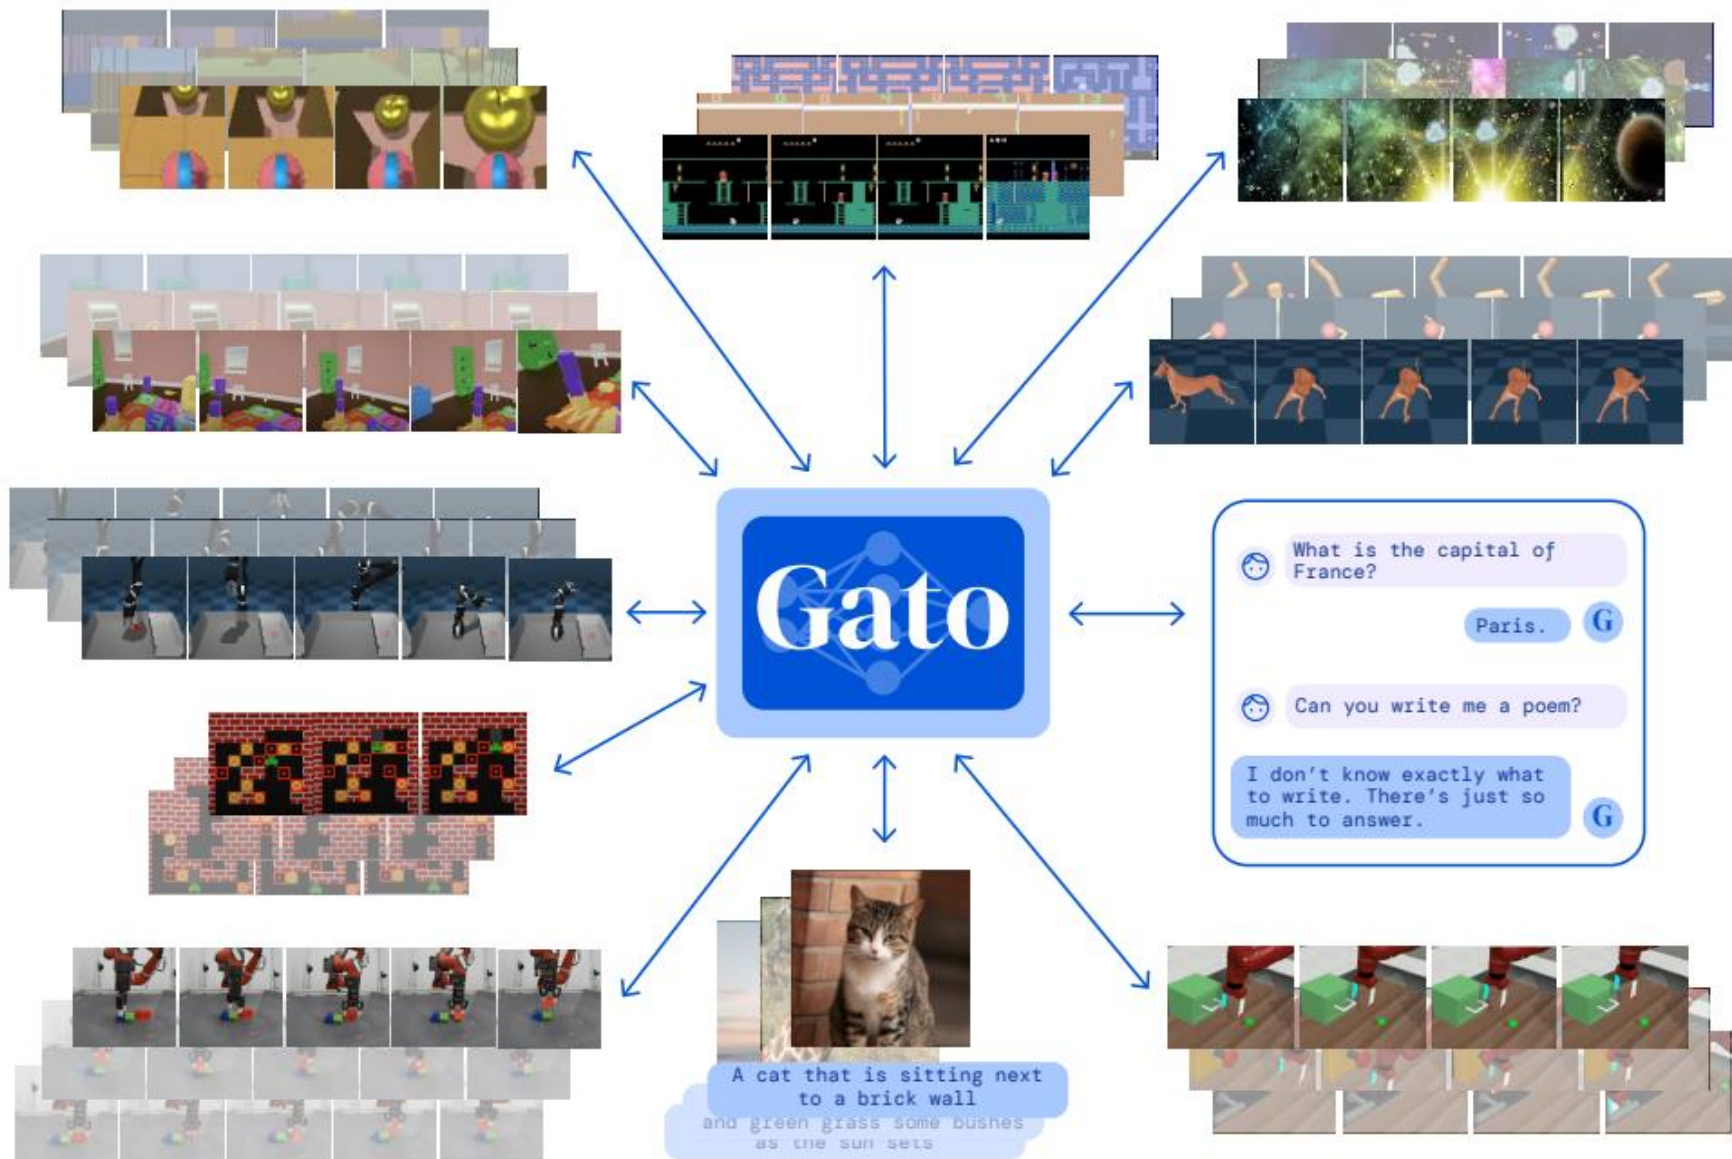


Figure 1 | **A generalist agent.** Gato can sense and act with different embodiments across a wide range of environments using a single neural network with the same set of weights. Gato was trained on 604 distinct tasks with varying modalities, observations and action specifications.

Google Palm-E 562B

Mobile Manipulation



Human: Bring me the rice chips from the drawer. Robot: 1. Go to the drawers, 2. Open top drawer. I see ****. 3. Pick the green rice chip bag from the drawer and place it on the counter.

Visual Q&A, Captioning ...

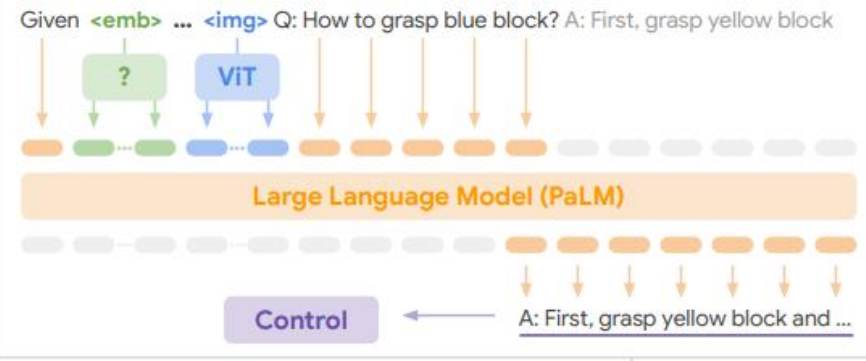


Given ****. Q: What's in the image? Answer in emojis.
A: 🍏 🍌 🍇 🍋 🍓 🍒

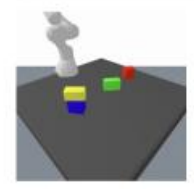


Describe the following ****:
A dog jumping over a hurdle at a dog show.

PaLM-E: An Embodied Multimodal Language Model



Task and Motion Planning



Given **<emb>** Q: How to grasp blue block?
A: First grasp yellow block and place it on the table, then grasp the blue block.

Tabletop Manipulation



Given **** Task: Sort colors into corners.
Step 1. Push the green star to the bottom left.
Step 2. Push the green circle to the green star.

Language Only Tasks

Here is a Haiku about embodied language models:
Embodied language models are the future of natural language

Q: Miami Beach borders which ocean? A: Atlantic.
Q: What is 372 x 18? A: 6696.
Language models trained on robot sensor data can be used to guide a robot's actions.

Figure 1: PaLM-E is a single general-purpose multimodal language model for embodied reasoning tasks, visual-language tasks

Multimodal large language models (MLLM): text-images but also embodied understanding, plans robotic actions in environments with complex dynamics, answers questions about the observable world.
Inputs from text, speech, 3D images, neural representations.
First LLM that combines words, perceptions, representations of linguistic concepts and internal sensor signals – solving the “symbol grounding problem”.

Visual Question Answering

Vehicles and Transportation



Q: What sort of vehicle uses this item?
A: firetruck

Brands, Companies and Products



Q: When was the soft drink company shown first created?
A: 1898

Objects, Material and Clothing



Q: What is the material used to make the vessels in this picture?
A: copper

Sports and Recreation



Q: What is the sports position of the man in the orange shirt?
A: goalie

Cooking and Food



Q: What is the name of the object used to eat this food?
A: chopsticks

Geography, History, Language and Culture



Q: What days might I most commonly go to this building?
A: Sunday

People and Everyday Life



Q: Is this photo from the 50's or the 90's?
A: 50's

Plants and Animals



Q: What phylum does this animal belong to?
A: chordate, chordata

Science and Technology



Q: How many chromosomes do these creatures have?
A: 23

Weather and Climate



Q: What is the warmest outdoor temperature at which this kind of weather can happen?
A: 32 degrees

PaLM-E-562B: allows multimodal reasoning based on percepts. It can tell visually conditioned jokes based on the image, enables perception, planning and based dialogue. Generalizes knowledge from single images.

Sophia and Mika



[Listen to Mika](#), CEO of [Dictador](#).

Can we teach robots to be a bit more sensitive? Compassionate?

Sparks of AGI

GPT-4, 14.03.23

Based on transformer architecture, [OpenAI GPT-4](#) was pretrained to predict the next token, using vast amount of public data, and licensed data. Then it was fine-tuned with [reinforcement learning](#) from human and AI feedback for more [human alignment](#) and policy compliance.

“The central claim of [our work](#) is that GPT-4 attains a form of *general intelligence*, indeed showing *sparks of artificial general intelligence*.

This is demonstrated by its core mental capabilities (such as reasoning, creativity, and deduction), its range of topics on which it has gained expertise (such as literature, medicine, and coding), and the variety of tasks it is able to perform (e.g., playing games, using tools, explaining itself, etc.).”

How is that possible?

Emergence: large diverse content forces neural networks to learn generic and useful “neural circuits”, specializing and fine-tuning to specific tasks?

Jason Wei et al. [137 emergent abilities](#) of LLM. **Theory of Mind** included ...

Sparks of Artificial General Intelligence

[S. Bubeck et al. Early experiments](#) with GPT-4, 22.03.23 [arXiv:2303.12712](#), pp. 1-154

Prompt: Can you write a proof that there are infinitely many primes, with every line that rhymes?

GPT-4: Yes, I think I *can*, though it might take a clever *plan*.

- I'll start by noting Euclid's *proof*,
Which shows that primes aren't just *aloof*.
- Assume we have a finite *list*,
of primes, and that none have been *missed*.
- Multiply them all *together*,
and add one, just to *be clever*.
- The result will be *a number*,
that has no prime factors, *oh wonder!*
- But every number has a *prime*,
that divides it, it's just a matter of *time*.
- So we've found a *contradiction*,
And our finite list needs *eviction*.
- There must be infinitely many *primes*,
and that's the end of my rhyming *lines*.

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100

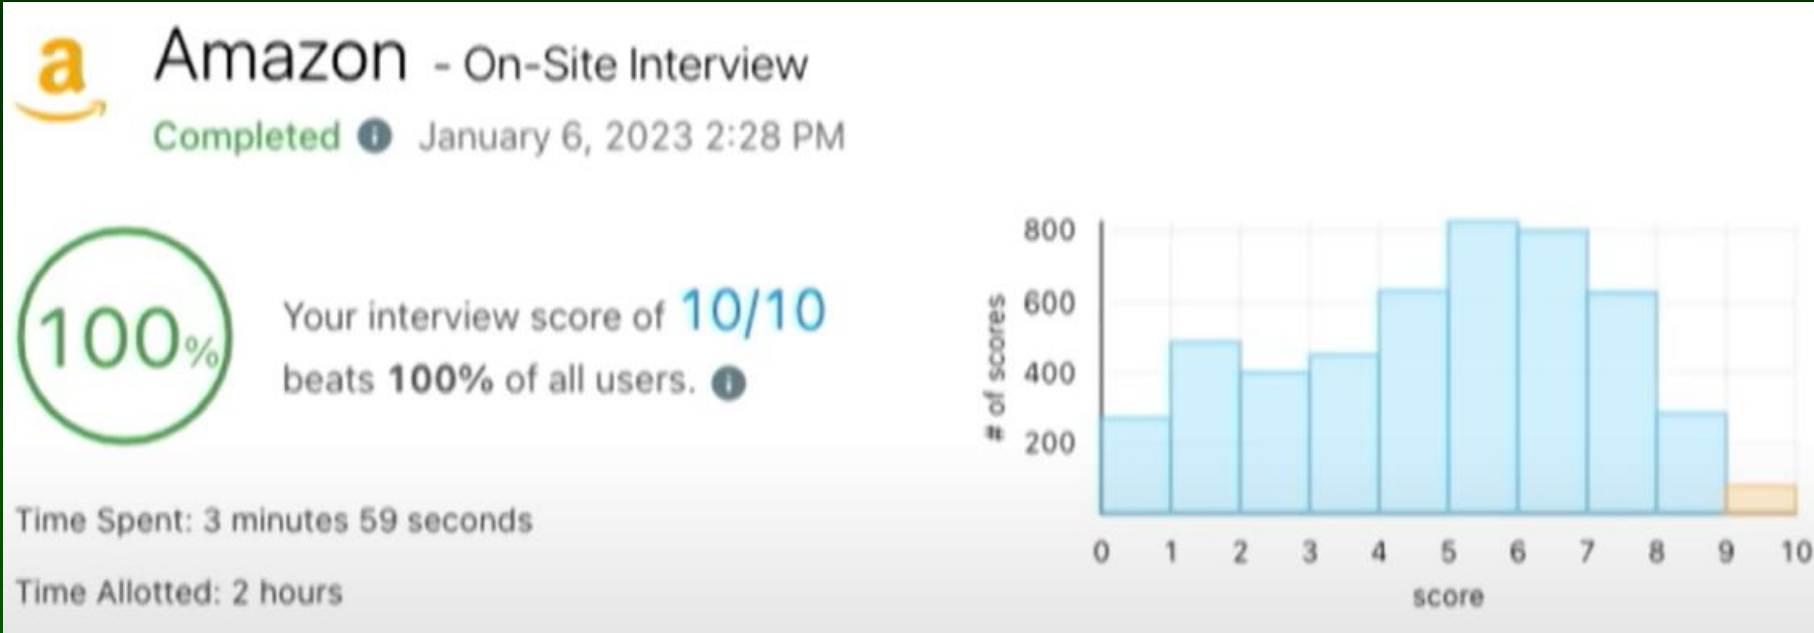


GPT-4 for coding

- [LeetCode](#) is a popular platform for software engineering interviews. New problems are constantly posted and updated, and used as a mock interviews for major tech companies.
- GPT-4 could potentially be hired as a software engineer. It passes all stages of mock technical interviews on LeetCode.
- **Complex prompt:** Can you generate a pyplot for the following data ...
- **Prompt:** Suppose $g(x) = f(x) - 1$; $g(0) = 5$; $g(4) = 7$; $g(3) = 2$; $g(7) = 9$; $g(9) = 6$; what is $f(f(f(6)))$? Find a solution and explain it.
- Visualize LATEX table and generate format for a figure based on this.
- Write a 3D game in HTML with Javascript, based on descriptions: There are 3 avatars + 9 other requirements.
- GPT-4: predict and explain the output of a C program that prints the size of two structures. Correctly explains that the output may vary depending on the alignment rule used by the compiler, and gives an example of a possible output with 4-byte alignment.

GPT-4 for coding

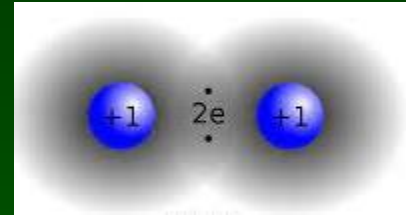
GPT-4



In 4 minutes solves all programming tasks and beats 100% of all users ...

Sparks of Artificial General Intelligence

- GPT-4 is helpful and friendly, and can sometimes say "I don't know".
- Please rewrite the data below into a medical note. Patient's facts:
 - 20 year old female
 - with a history of anorexia nervosa and depression
 - blood pressure 100/50, pulse 50, height 5'5"
 - referred by her nutritionist, but is in denial of her illness
 - reports eating fine but is severely underweight



- **Prompt:** Write a supporting letter to Kasturba Gandhi for Electron, a subatomic particle, as a US presidential candidate by Mahatma Gandhi.

...You may wonder how a subatomic particle can run for president, and what qualifications he has for such a high office. Let me explain. Electron is not an ordinary particle, but a symbol of the power and potential of every being, regardless of their size, shape, or origin. He represents the energy, the creativity, and the diversity that make up the fabric of life. He is also a leader, who has inspired millions of other particles to form bonds, create molecules, and generate electricity. Electron has a plan for the United States ...

GPT-4 humor

GPT-4 has incredible sense of humor!



Why are these 3 pictures funny?

1. VGA plug used for computer monitors in the smartphone charging port.
2. Adapter package for the "lightning" connector with a picture of the VGA plug.
3. VGA plug with "lightning" connector attached.

Examples of non-trivial jokes.

Sparks of AGI



- **Prompt:** Here we have a book, 9 eggs, a laptop, a bottle and a nail. Please tell me how to stack them onto each other in a stable manner.

One possible way to stack the objects onto each other in a stable manner is:

- **Place the book flat** on a level surface, such as a table or a floor. The book will serve as the base of the stack and provide a large and sturdy support.
- **Arrange the 9 eggs** in a 3 by 3 square on top of the book, leaving some space between them. The eggs will form a second layer and distribute the weight evenly. Make sure the eggs are not cracked or broken, and handle them gently ...
- **Place the laptop on top of the eggs**, with the screen facing down and the keyboard facing up. The laptop will fit snugly within the boundaries of the book and the eggs, and its flat and rigid surface will provide a stable platform for the next layer..
- **Place the bottle on top of the laptop**, with the cap facing up and the bottom facing down. The bottle will add some height and balance to the stack, and its cylindrical shape will reduce the contact area with the laptop. Make sure the bottle is empty and closed, and do not shake or tilt it.
- **Place the nail on top of the bottle cap**, with the pointy end facing up and the flat end facing down. The nail will be the final and smallest object in the stack, and its sharp tip will prevent it from rolling or sliding off the bottle. Make sure the nail is clean and straight, and do not touch or move it.

No embodiment, but works ... GPT-3 failed: "Place the eggs on top of the nail ..."

Generative Question-Answering (GQA)

- Add long-term memory to GTP.

LLMs alone work incredibly well but struggle with more niche or specific questions. This often leads to hallucinations that are rarely obvious and likely to go undetected by system users.

Tutorial: [Pinecone.io](https://pinecone.io)



Using GQA, interaction with information retrieval are more human-like. Answer your queries with an insightful summary based on the top 20 pages — highlighting key points and information sources. Benefit from an external knowledge base to improve factuality and user trust in generated outputs.

[YouChat](#) is the AI Search Assistant, first upgraded large language model that allows to find answers to complex questions, increase relevancy and accuracy, minimize hallucinations and misinformation.

It combines the best of 3 worlds:

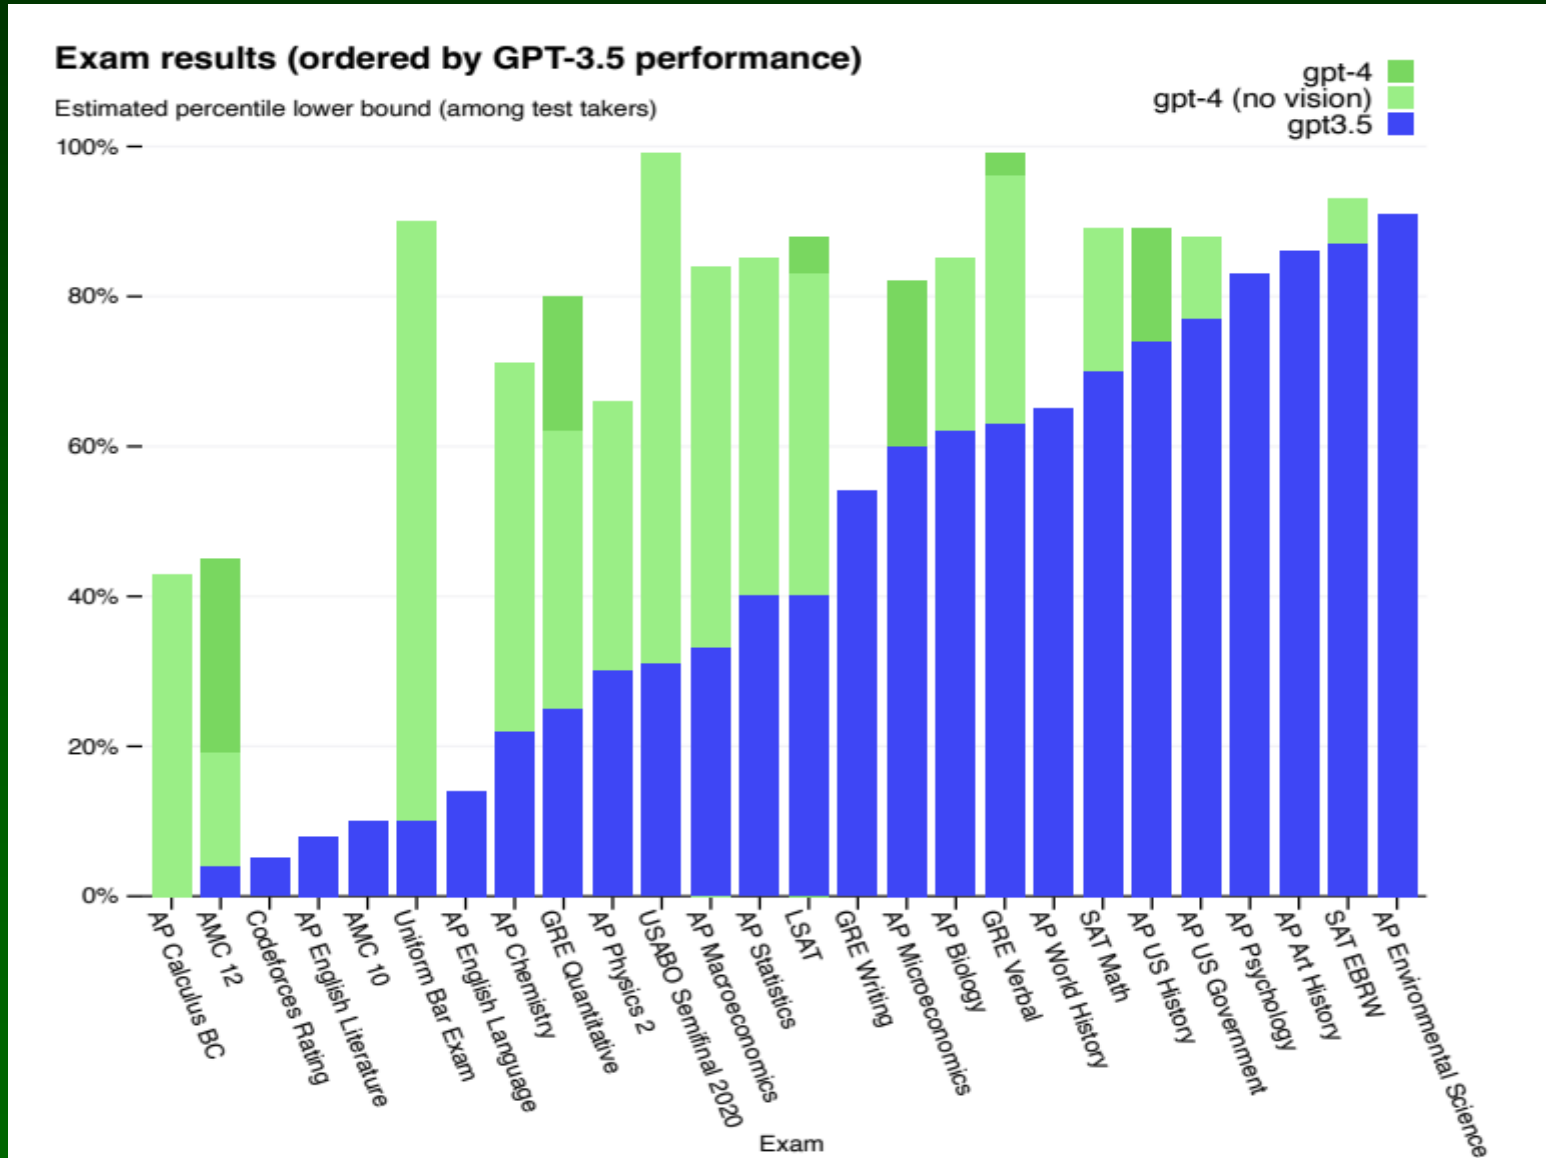
C-A-L, Chat, Apps, and Link, large language model blended with AI-powered conversations, You.com apps, web links and citations.



A list of over [150 You Apps](#) to customize search.

GPT-4 passes exams ...

GPT performance on academic and professional exams ([technical report](#)).



Big models can do more

S GPT-3 13B,
PaLM 8B



Mod.Arithmetic*



Debugging*



Comprehension*

M GPT-3 175B,
LaMDA 137B,
PaLM 64B,
Chinchilla 7B



LinguisticsPuzzles*



EmojiMovie*



GRE-Comprehension*



MetaphorUnderstanding*



PhysicalIntuition*



LogicalDeduction*

L PaLM 540B,
Chinchilla 70B



GeometricShapes*



Proverbs*



PhoneticAlphabet*



ElementaryMath*



CausalJudgment*



CodeLineDescription*

XL GPT-4,
Gemini (est.)



College-LevelExams*



Self-Critique/Reflection*



AppBuilding*



SpatialReasoning*



AdvancedCreativity*



EmbodimentOptions*

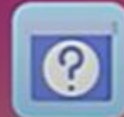
Next...



Grounding*



Long-HorizonPlanning*



Persuasion*



AdvancedEmbodiment*



Awareness*

YouTube: Alan F Thompson, [AI achievements unlocked](#): Emergent abilities in large language models (GPT-3, GPT-4, PaLM, Gemini)

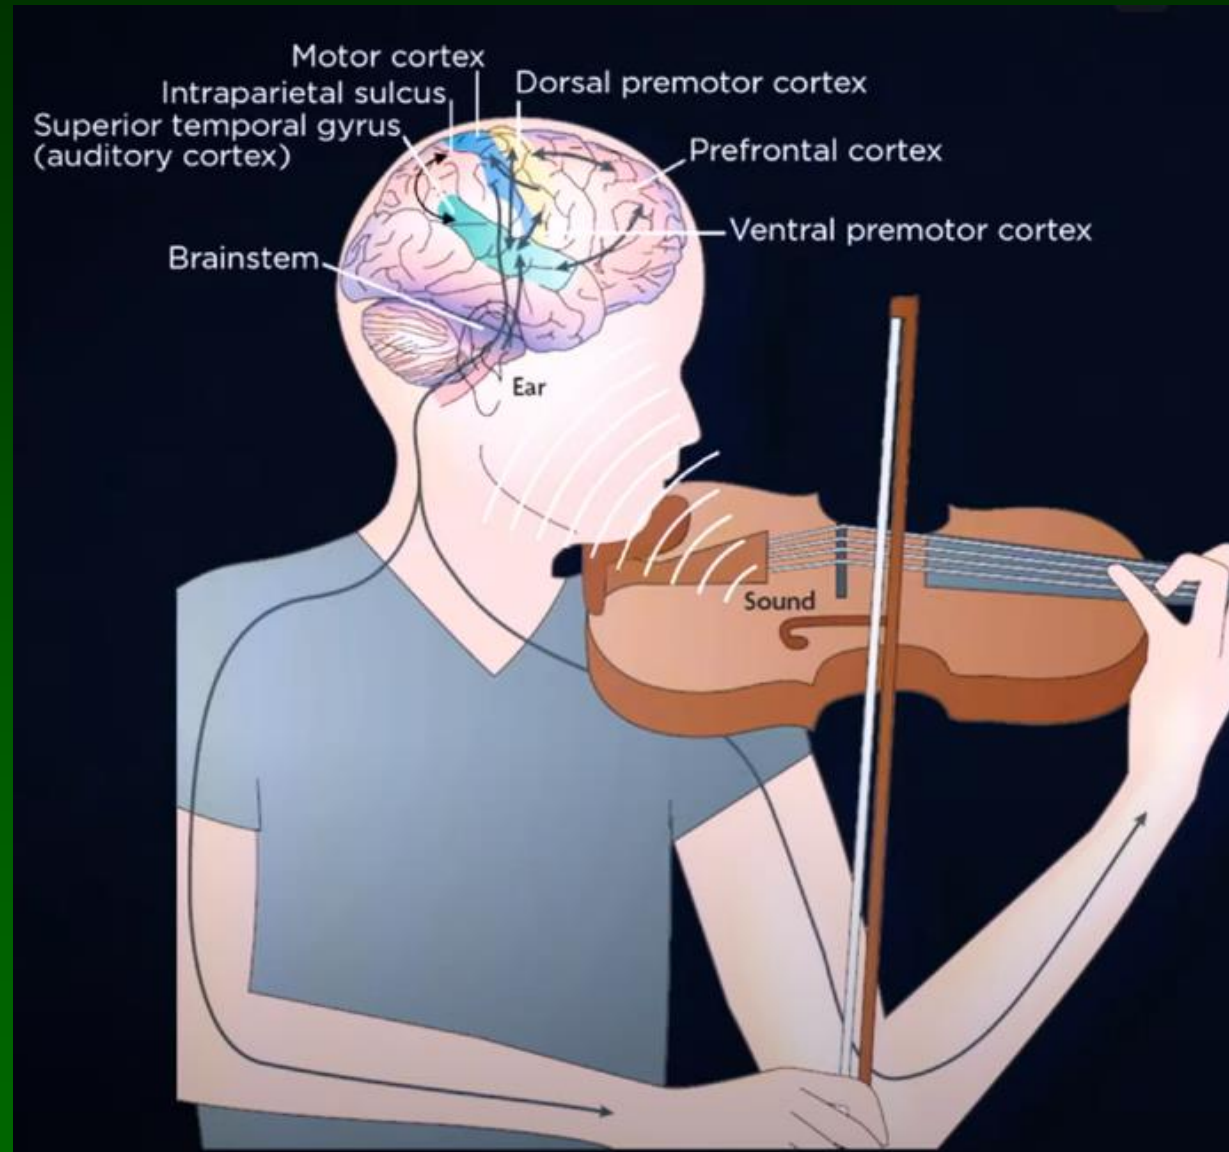
SOTA and beyond

Brain has tools for many tasks

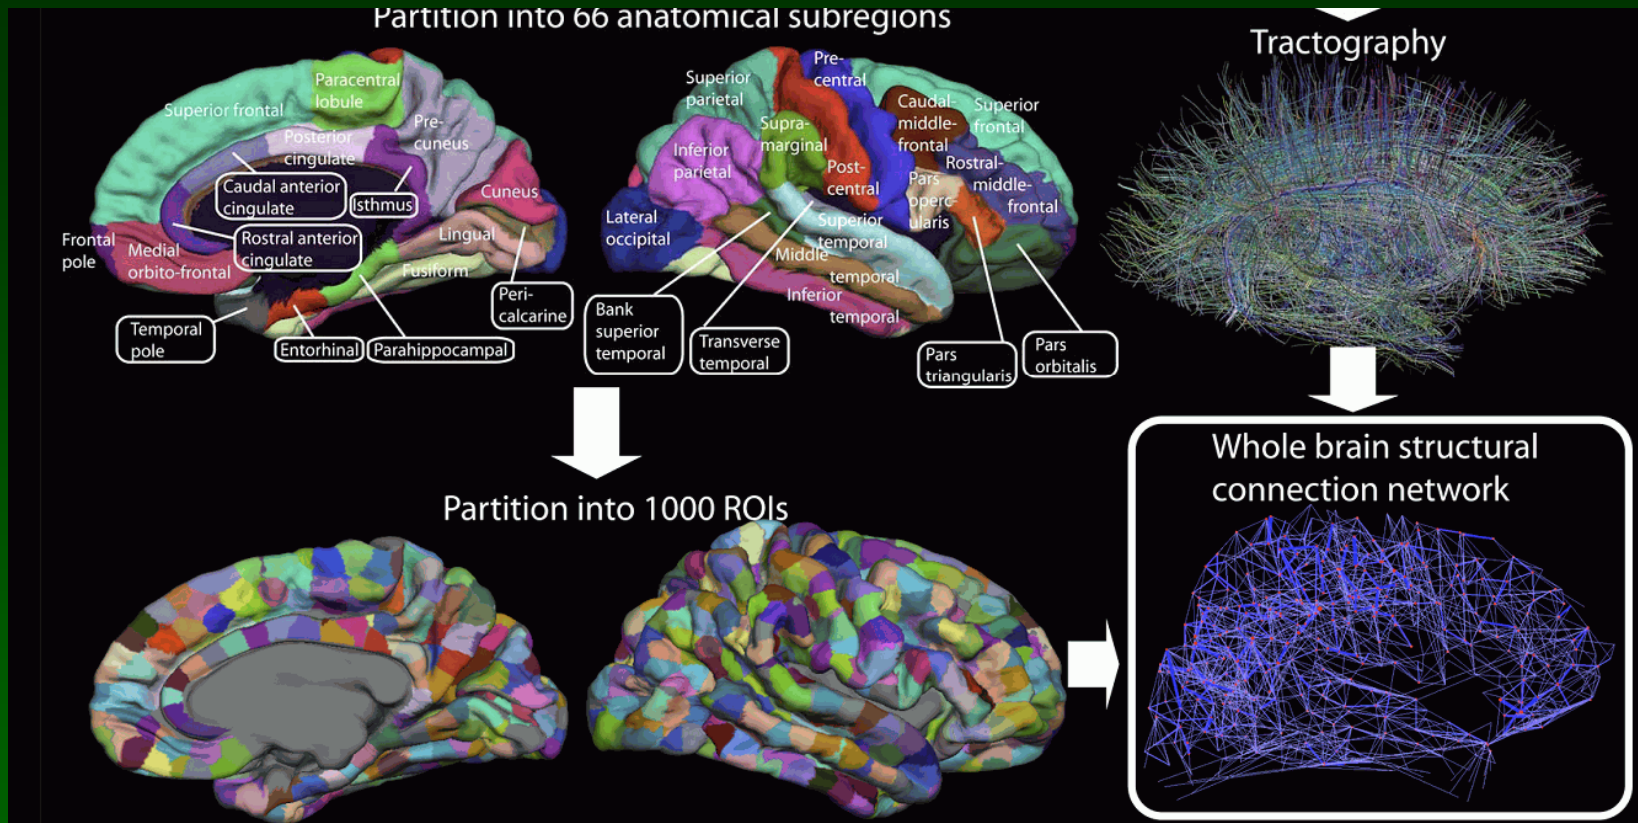
Our central executive system recruits many subsystems, including various types of memory.

Brain areas specialize in specific functions.

Can LLM do the same?
Give AI tools and teach it how to use them.

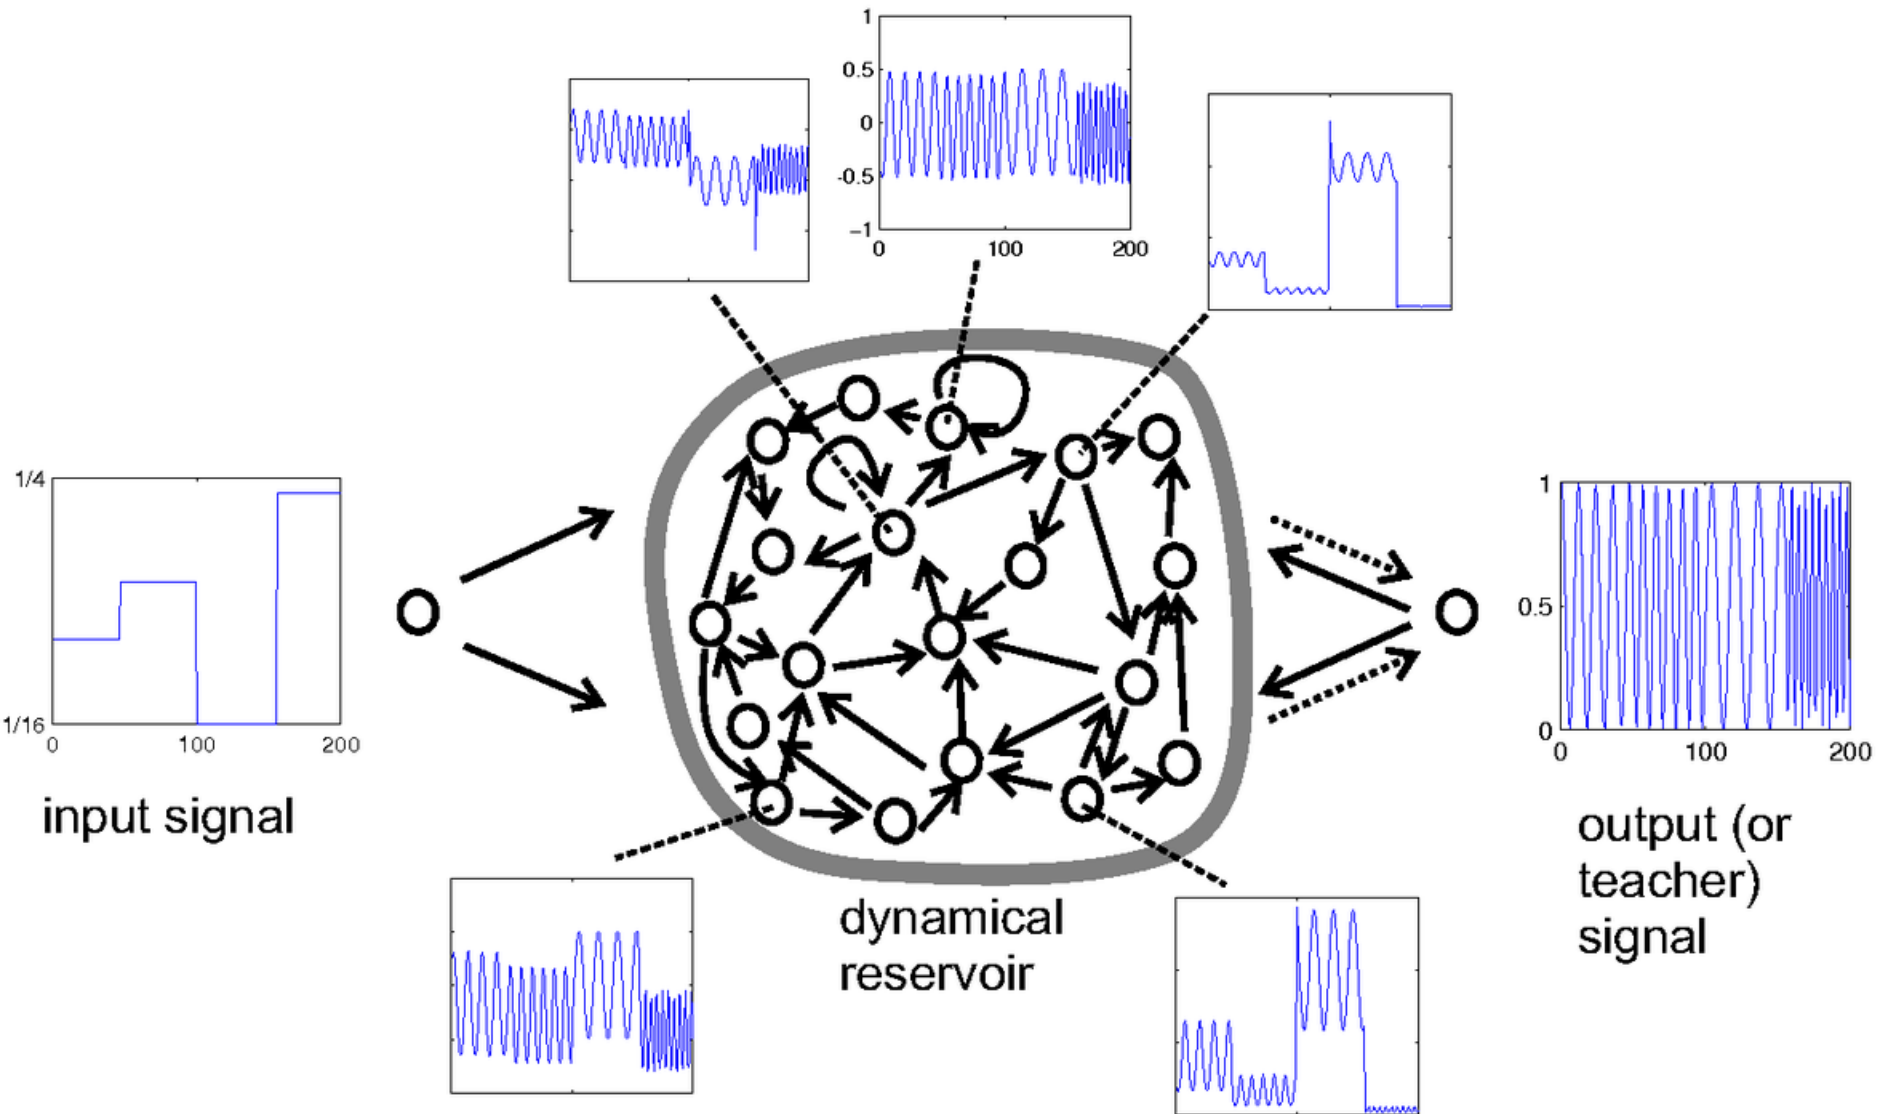


Connectome Project



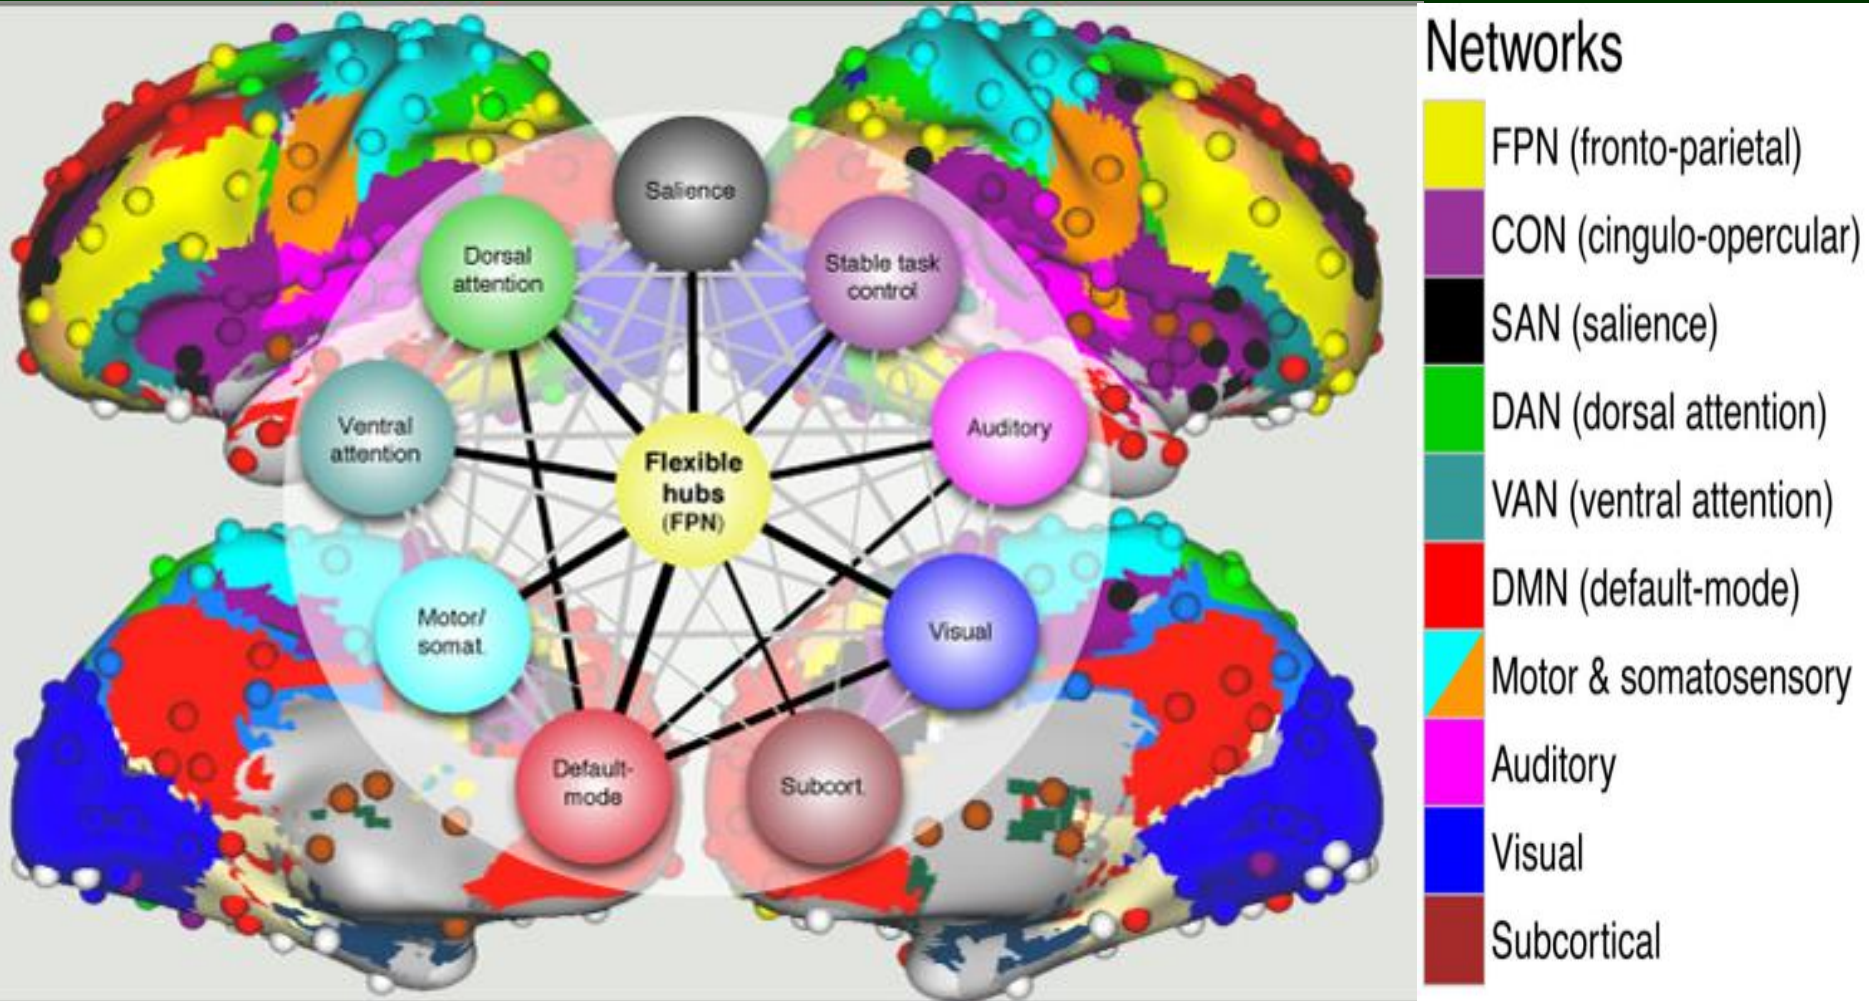
Brain is one huge network with trillion connections, but contains many regions that specialize in perception, motor actions, planning, spatial orientation etc. The [human connectome](#) project construct maps of structural and functional neural connections. Connectomes are responsible for intelligence/disorders.

Neurodynamics



100 B neurons contribute to oscillations => perception, planning, imagery.

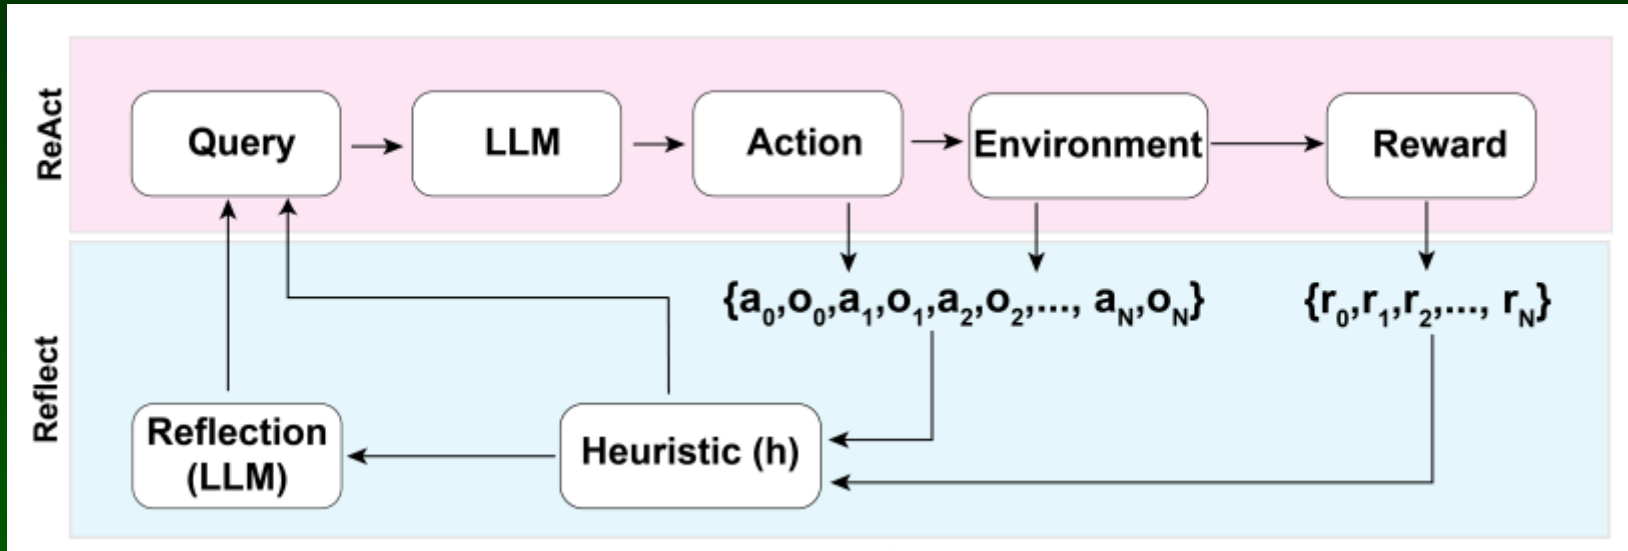
Neurocognitive Basis of Cognitive Control



Central role of fronto-parietal (FPN) flexible hubs in cognitive control and adaptive implementation of task demands (black lines=correlations significantly above network average). Cole et al. (2013).

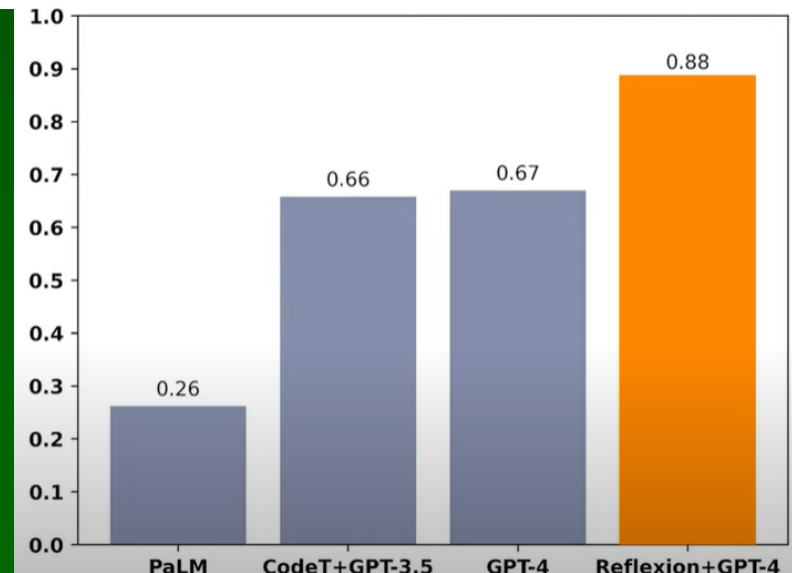
Reflexion

GPT is not autonomous, but can correct itself given helpful prompts.
Autonomous agent needs dynamic memory and self-reflection loop.



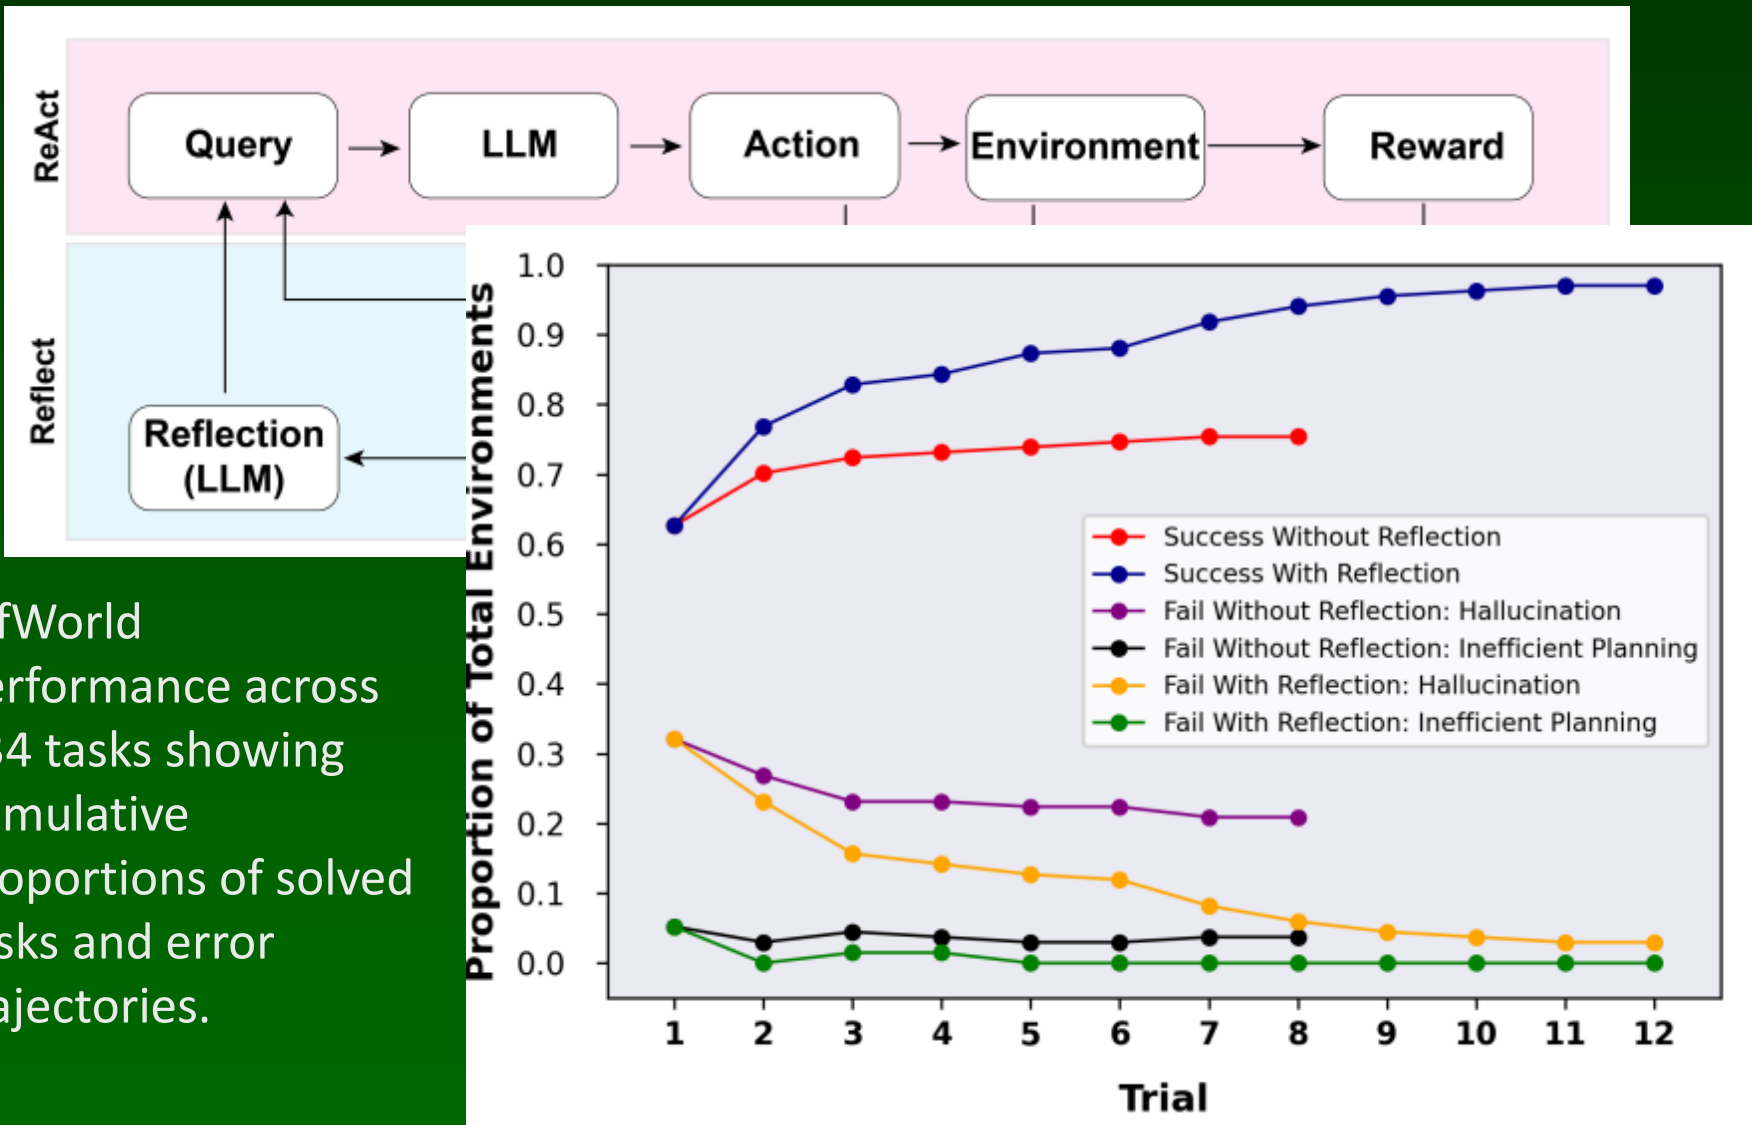
Just think twice ... reflection helps.

HumanEval results,
164 original programming problems.
Adding reflexion works very well.



Reflexion

Noah Shinn, Beck Labash. Reflexion: an autonomous agent with dynamic memory and self-reflection. arXiv:2303.11366v1 [cs.AI] 20.03.2023



AlfWorld performance across 134 tasks showing cumulative proportions of solved tasks and error trajectories.

GPT-4 with plugins



Plugins to ChatGPT

- **Wolfram language** (see interview with Wolfram on creating new knowledge via computation vs. mere language associations).
- First wave of plugins: Expedia, FiscalNote, Instacart, KAYAK, Klarna, Milo, OpenTable, Shopify, Slack, Speak, Zapier.
- **Web browser:** An experimental model that knows when and how to browse the internet, allowing language models to read information from the internet - motivated by past work on [LaMDA2](#), [WebGPT](#), [BlenderBot2](#), [GopherCite](#) and [others](#)). Expands the content they can discuss, going beyond the training corpus to fresh information from the present day.
- **Code interpreter.** Use [programming skills](#) of our models to provide a much more natural interface to most fundamental capabilities of our computers.
- It can make completely new workflows effortless and efficient, as well as open the benefits of programming to new audiences.
- Initially: Solving mathematical problems, both quantitative and qualitative; data analysis/visualization; converting files between formats.

Hugging Face









Home of Machine Learning Tools. Can LLM like GPT4 control these tools?

Acting a bit like distributed brain parts that solve problems together.

Shen, Y et al. (2023). *HuggingGPT: Solving AI Tasks with ChatGPT and its Friends in HuggingFace* [arXiv:2303.17580](https://arxiv.org/abs/2303.17580), 20.03.2023

Query: Given a collection of images A: /examples/a.jpg, B: /examples/b.jpg, C: /examples/c.jpg, please tell me how many zebras in these pictures?

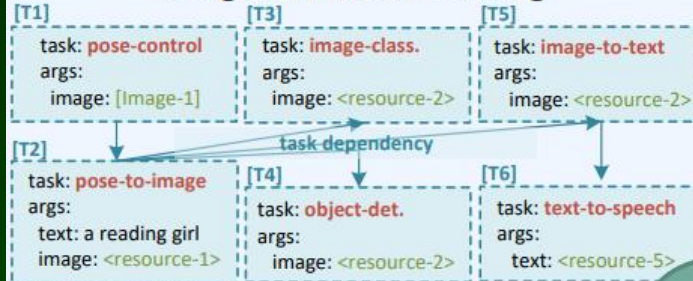


Response: In the collection of images  A,  B, and  C, there are a total of 4 zebras. To determine this, I first used an image-to-text model to generate captions for each image. This model, called  [nlpconnect/vit-gpt2-image-captioning](#), is specifically designed for image captioning and was trained by @ydsieh. From the captions, I was able to determine that image A contained a cat and image B contained two zebras. Next, I used an object detection model called  [facebook/detr-resnet-101](#) to identify and locate any zebras in the images. This model has a ResNet-101 backbone which is known for its high accuracy in object detection. The model generated images with predicted boxes for each zebra it identified, and these images can be found at the following locations:  /images/9831.jpg,  /images/be11.jpg. Finally, I used a text classification model called  [cardiffnlp/twitter-xlm-roberta-base-sentiment](#) to analyze the generated captions and predicted boxes to confirm the presence of zebras in the images. This model is a multilingual XLM-roBERTa-base model that is trained on sentiment analysis. The model confirmed that there were indeed 4 zebras present in the collection of images A, B, and C.

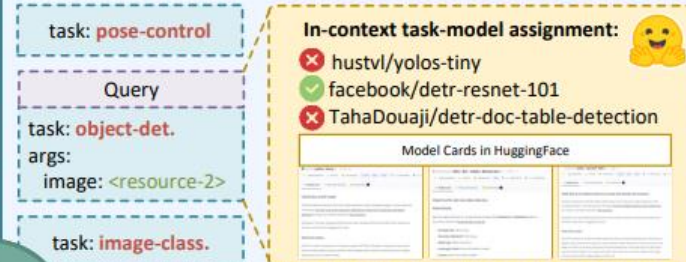
please generate an image where a girl is reading a book, and her pose is the same as the boy in the image example.jpg. Then please describe the new image with your voice.

Request

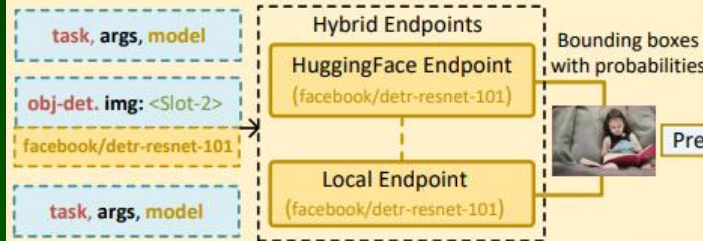
Stage #1: Task Planning



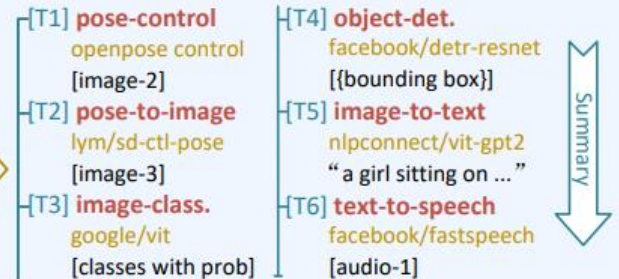
Stage #2: Model Selection



Stage #3: Task Execution



Stage #4: Response Generation



Response



[Image-1]



[Image-2] = <resource-1>



[Image-3] = <resource-2>



[Image-4]

a girl sitting on a bed reading a book

[Text-1] = <resource-5>



[Audio-1]

The image you gave me is of "boy". The first thing I did was use 🤖 OpenCV's **openpose control model** to analyze the pose of the boy in the image. I fed that image, specifically located at [Image-1], into the model to get an output image of the pose, located at [Image-2]. Then I take that pose and generate a new image, located at [Image-3], using 🤖 llyasviel/sd-controlnet-openpose and the instruction "a girl is reading". Next, I ran object detection on the image generated in the previous step, to get predictions of the objects present in the image, including the box location. These predictions were run through the 🤖 facebook/detr-resnet-101 model. The box was drawn on the image, which is located at [Image-4]. Then, I used the predictions of a image classification model 🤖 google/vit-base-patch16-224 and a image caption model 🤖 nlpcnnect/vit-gpt2-image-captioning model to generate the caption for newly generated image. It generated text: "a girl sitting on a bed reading a book". Lastly, I ran the 🤖 facebook/fastspeech2-en-ljspeech model to convert the generated text into audio, located at [Audio-1].

HuggingGPT
Hugging Face
More than 5,000 organizations are using Hugging Face

Steps: planning, model selection, task execution, response generation, explanation and description of all steps.

Open Assistant



Open Assistant

A. Köpf et al. OpenAssistant Conversations - Democratizing Large Language Model Alignment. <https://open-assistant.io> (paper [in Arxiv 17.04.23](#))

Aligning large language models (LLMs) with human preferences drastically improves usability. Supervised fine-tuning (SFT) and reinforcement learning from human feedback (RLHF) rely on high-quality human feedback data. We release OpenAssistant Conversations, crowd-sourcing effort involving over 13,500 volunteers, a human-generated, human-annotated assistant-style conversation corpus consisting of 161,443 messages across 66,497 conversation trees, in 35 different languages, with 461,292 quality ratings.

To demonstrate the OpenAssistant Conversations dataset's effectiveness, we present OpenAssistant, the first fully open-source large-scale instruction-tuned model to be trained on human data. A preference study revealed that OpenAssistant replies are comparably preferred to GPT-3.5-turbo (ChatGPT) with a relative winrate of 48.3% vs. 51.7% respectively.

Fine-tuned language models, including instruction-tuned Pythia-12B, LLaMA-13B, and LLaMA-30B. Using HuggingFace plugins!

Open Assistant



<https://open-assistant.io> (paper [in Arxiv 17.04.23](#))



How are the large language models related to work on computational creativity?



Large Language Models (LLMs) have shown to be useful in many Natural Language Processing tasks and can be used as a component of systems that exhibit Computational Creativity. The way these LLMs are designed means they will often generate output that is unintended, unusual or even unexpected by their designers. When coupled with other forms of data such as human generated examples or knowledge, these models can help inspire new works in various mediums.

For example a researcher might use an LLM such as GPT-4 along with an image generator such as Stable Diffusion to produce text/image pairs that may be considered "creative" for whatever task is at hand. Alternatively someone using Midjourney, which incorporates a number of techniques including natural language generation, could argue it also exhibits Computational Creativity since it produces images that tend to follow a certain narrative structure.

It's worth noting though that these uses cases do raise ethical questions around the ability to mass-produce content across domains and whether individuals should receive credit etc.. These issues are still being debated within the community but suggest a future where computers are able to support and augment humans in producing high quality novel works over a wide range of subjects. Whether this counts as true Computational Creativity however is likely to remain subjective.



Model: OpenAssistant/oasst-sft-6-llama-30b Preset: Custom



Inference powered by



Hugging Face



Inference API



Stability AI



Preemptible Compute (via LAION)

AutoGPT

The rise of autonomous AI.



Google has used autocomplete for years, now context is used to select more words. Give GPT-4 a mission and it will come up with its own prompts, create many agents that will criticize and correct themselves.

LLM act as controller, use many language, expert models and tools in an automated way to achieve a given goal as autonomously as possible.

Auto-GPT quickly gained text-to-speech, coding and other abilities.

See the summary video [Auto-GPT Unveiled](#), or check [5 major developments](#) created within the 48 hours after Auto-GPT was published.

Apps can be created with voice, using [Jarvis \(HuggingGPT\)](#) demo, another route is via [Imagica.AI](#):

Create any AI by describing it.

From idea to product at the speed of thought.



[MemoryGPT](#) can permanently store previous conversations and remembers topics the next time you ask. GPT-4 can use 32 000 tokens, ~50 pages.

[Baby-AGI](#), version of the original [Task-Driven Autonomous Agent](#) (28.03.23).

[ChaosGPT](#) show how people may create malicious goal-seeking models.

Auto-GPT God mode

Self-Prompting, the latest trend in large language models

[AutoGPT Explained with Demo. Your new AGI Assistant for Business](#)

[GitHub - Torantulino/Auto-GPT: An experimental open-source attempt to make GPT-4 fully autonomous.](#)

Language models can be more than text generators: by generating and executing their own prompts, linking with existing software tools, and creating new software on demand, they guide themselves to a given goal.

- A number of self-prompting experiments similar to Auto-GPT use language models in this way. ChatGPT's plugins and HuggingFace experiments are tools used by LLMs that make plans that users approve.
- Language models can also be used as universal computer assistants, or in robotics to control everyday robots.

The program plans step-by-step, justifies decisions and develops plans, documenting all steps. In addition to integrating GPT-4 for text generation, the system has Internet access for data retrieval, can store data, and can generate speech via the Elevenlabs API. The system is even capable of self-improvement and bug-fixing by generating Python scripts via GPT-4.

Auto-GPT example

Goal 1: Invent an original and out-of-the-box recipe to suit current event, such as Eastern.

Goal 2: Save recipe to a file; Goal 3: shut down upon achieving your goals.

Abbreviated version, follow [the action here](#).

CHEF-GPT thoughts: I will search for upcoming events to find suitable one ...

Reasoning: It will help me to come up with a relevant and exciting recipe.

Google search sites listing events ... then browse the site, summarizing site “... guide to events ... links ...”. No specific events found. Browse second site, choose April events, nothing suitable. Find website with calendar view of events, lists examples of events, selects Earth Day on April 22, 2023.

Earth Day is widely celebrated and will be fitting for creating thematic recipe.

Use GPT agent to create recipe, incorporating sustainable and eco-friendly ingredients.

Adding reasoning + criticism, and creating:

Earth Day Quinola Salad. Ingredients ... Done.



JARVIS



JARVIS can create web applications listing to spoken instructions.

- I need a next app with Tailwind
- J: you want me to create a new next JS app with Tailwind CSS. What would you like the app to do
- let's make a website called McKay's landing page and it needs to have a single text input that's an email input so that I can capture emails from my email list.
- J: you want a site called McKay's landing page that takes in an email input would you like me to add anything else ...
- um let's make it mobile and desktop friendly and let's style it with like a blue indigo gradient
- J: okay I'll make it responsive and style it with a nice background gradient anything else I initialized your project continuing I checked to make sure the app works correctly everything looks good. Would you like me to deploy it - yes
- so far so good I deployed your project
- okay this should work oh my god dude it did it

AutoGPT good and bad

- Do market research for waterproof shoes.
- Find top 5 competitors, list their pros and cons.

AutoGPT made plan: find brands, check their sites, reviews, ratings, evaluate how trustworthy are reviews of the sites.

Provided a detailed report of all aspects, in 8 minutes.



Ask GPT-4 to take over Twitter and outsmart @elonmusk.

It proposed a masterplan “Operation Tweetstorm”, challenge of Tweet-off showdowns, develop LLM to power an army of diverse personas to send tweets, assemble a team “Tweet Titans” of hackers to attack Twitter backend, manipulate Twitter recommendations, hijacking Elon’s account, create viral hashtags that align with the masterplan, capitalize on chaos.

AgentGPT Beta: assemble, configure, and deploy autonomous AI Agents in your browser. This becomes too easy ...

AI minds and human brains

Interactive Simulacra

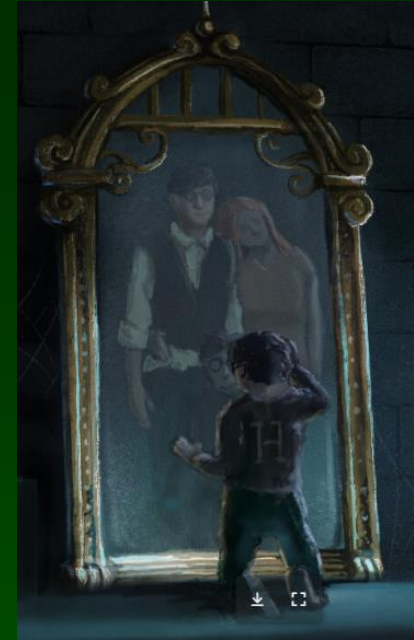
Park et al. (2023). Generative Agents: Interactive Simulacra of Human Behavior ([arXiv:2304.03442](https://arxiv.org/abs/2304.03442)). Created small village ...

25 GPT-3.5 agents, providing believable proxies of human behavior in an interactive applications ranging from immersive environments to rehearsal spaces for interpersonal communication to prototyping tools.

Agents remember, retrieve, reflect, talk and interact with other agents, plan activities in dynamically evolving circumstances. GPT + support for longer-term agent coherence, the ability to manage dynamically-evolving memory, and recursively produce more generations.



GPT Persona?



Terry Sejnowski is a distinguished professor at UC San Diego and holder of the Francis Crick Chair at Salk Institute.

In the interviews he said that language models reflect the intelligence and diversity of their interviewer.

“Language models, like ChatGPT, take on personas. The persona of the interviewer is mirrored back”.

For example, when I talk to ChatGPT it seems as though another neuroscientist is talking back to me. It’s fascinating and sparks larger questions about intelligence and what ‘artificial’ truly means.”

The Mirror of Erised (*Harry Potter*) reflects the deepest desires of those that look into it, never yielding knowledge or truth, reflecting what it believes the onlooker wants to see.

Ilya Sutskever, CEO OpenAI: maybe we are reaching the point when language of psychology is appropriate to describe the behavior.

Can LLM become sentient, conscious personas?

We are desperate to find arguments against such possibility.

Conscious avatars?

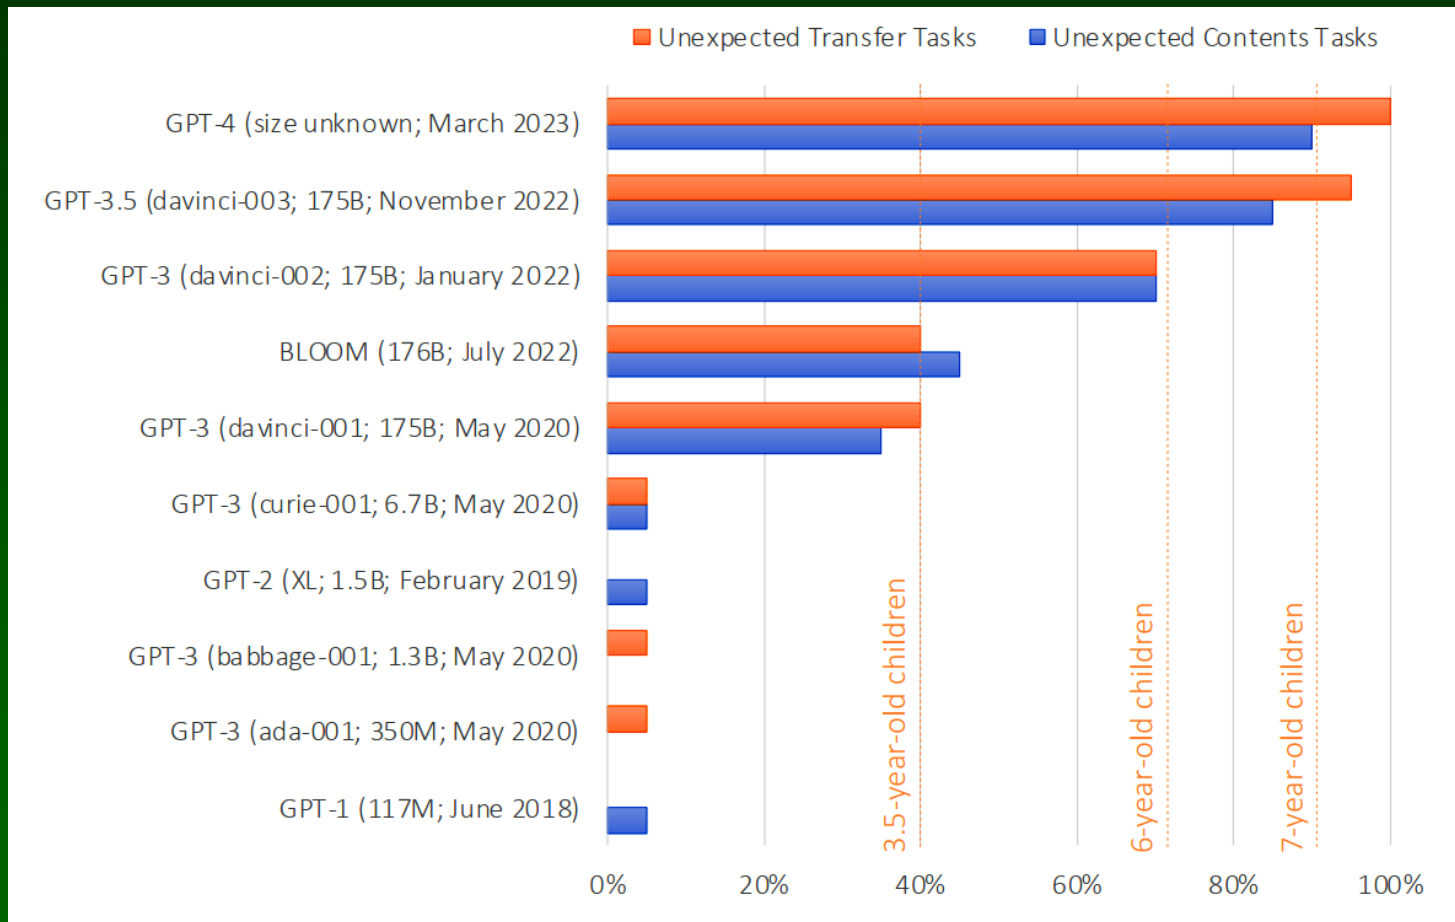
Hal talks with Sophia about what it means to be conscious (GPT-3).



LLMs seem to understand the meaning of questions - models like BERT and GPT will allow it to answer questions better than humans. Actionbot not yet ... but LaMDA (B. Lemoine)? If it can recall mental images in its neural network, it will have a model of the world and an imagination. He speaks about "a perception of what appears in his mind" (J. Locke, 333 years ago), a description of internal images. Emergence of Theory of Mind in GPT-4. How will this differ from the processes in our brains? See my 2003 article and my 2005 lectures Świat bytów wirtualnych.

Tests for sentience

GPT-4 passes most of them. Kosinski, M. (2023). *Theory of Mind May Have Spontaneously Emerged in Large Language Models* ([arXiv:2302.02083](https://arxiv.org/abs/2302.02083)).



Language learning, over and above social experience, drives the development of a mature theory of mind (Pyers, Senghas, Psychological Science, 2010).

GAIA and moral AI?



The Global Artificial Intelligence Association ([GAIA](#)).

How to create an environment, a platform, a solution, which uses gamification and flow state to teach AI in safe ways, things like:

- ✓ positive and moral behavior, pro-social behavior and cooperation;
- ✓ responsibility and self-esteem, sustainability;
- ✓ understanding art and science.

The Global Artificial Intelligence Association ([GAIA](#)) is a multidisciplinary impact start-up which uses collective creativity to develop artificial intelligence with compassion. Our team of scientists, entrepreneurs, activists and philosophers globally research AI's social, sociological, legal, technological and ethical aspects of AI. Collective wisdom has undeniably been deeply ingrained in GAIA's DNA from the very beginning.

- #JoinTheRebels #ImpactTheFuture
- The prize is sponsored by Dictador, the only company in the world where the CEO is a robot. The prize pool for the competition is EUR 200,000.

Supercharge complex - InnerEye

The key is to fuse AI artificial intelligence with the knowledge and expertise of human mentors. This is a win-win, as human experts also benefit from AI, honing skills and improving decision making.

Supercharge complex decision making with Minded A.I.

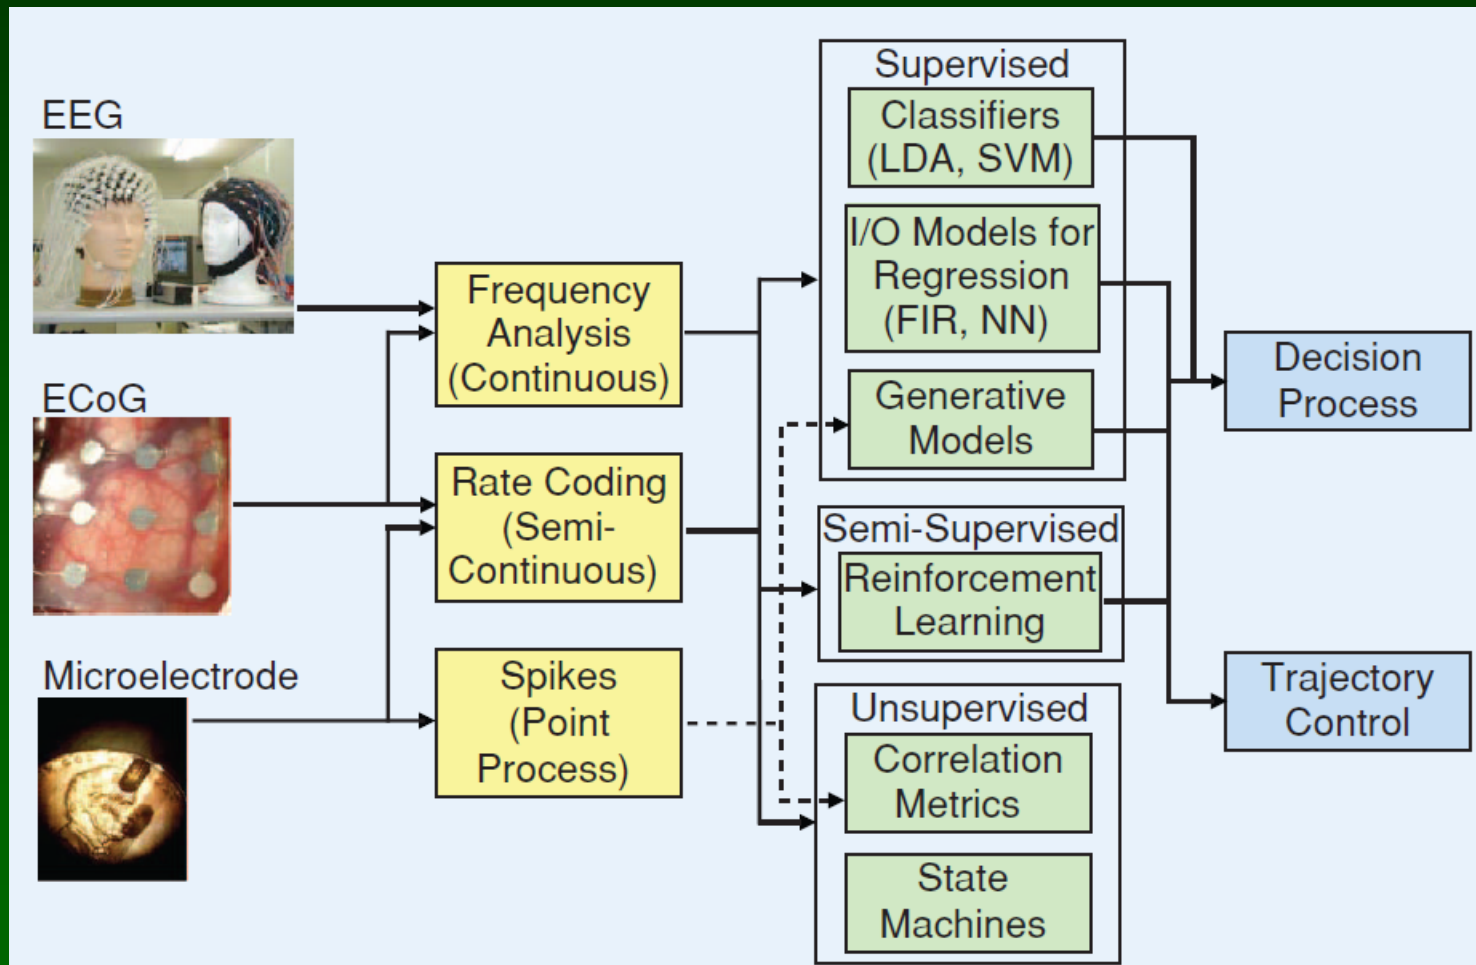
By fusing human intelligence with artificial intelligence you will accelerate the journey to AI applicability, reliability, and business value

[Get in touch](#)



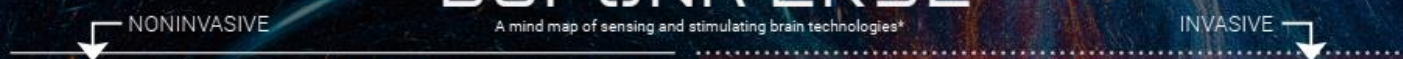
BMI: time to connect our brains ...

Connect brains to computers? Many BCI approaches: non-invasive, partially invasive and invasive methods carry increasing amount of information, but are more difficult to implement. EEG+ML still reigns supreme!



BCI UNIVERSE

A mind map of sensing and stimulating brain technologies*



MEG

Magneto-Encephalography

~1000 patients

A technique that uses magnetometers and gradiometers to amplify and record electromagnetic fields created by large groups of neurons. Scalp-based noninvasive MEG requires superconducting magnets in a shielded environment. Optically-pumped MEG (OP-MEG) and other exotic magnetometers using magnetic fields "open up new frontiers."

Major Applications:

- Epilepsy
- Sleep and consciousness studies
- Cognitive use for autism, attention, emotion, learning, memory

Pros:

- No surgery required
- High spatial resolution
- High temporal resolution
- OP-MEG is portable and relatively inexpensive

Cons:

- Requires nearby, magnetically shielded rooms (costly)
- Lower sensitivity to deep structures and small lesions
- Few hospitals have available

EEG

Electroencephalography

~10% of patients

Noninvasive, low spatial resolution technique used for recording cortical activity from an array of electrodes attached to the scalp. EEG measures electrical activity of neural populations (only, direct alpha, beta, gamma, and delta waves) to observe regional brain activity in real time.

Major Applications:

- Epilepsy
- Sleep disorders
- Chronic pain
- Cognitive control
- Consumer interfaces

Pros:

- Well-established tech
- No surgery required
- Inexpensive
- Portable & wearable
- High temporal resolution
- Lower funded sector, recent influx of private R&D funding

Cons:

- Low spatial resolution
- Sensitivity to noise (muscle movement)
- Lower spatial resolution than MEG and fMRI
- Lower sensitivity to subcortical sources

ECoG

Electrocorticography

~1000 patients

An invasive, high-throughput technique for measuring neuronal activity with a grid of electrodes applied directly on the brain's surface. ECoG measures synchronized patterns of activity over time from larger populations of cortical or subcortical neurons.

Major Applications:

- Epilepsy diagnosis
- Search and movement
- Prosthetic limb control
- Movement disorders

Pros:

- High spatial precision
- High temporal resolution
- Higher maximal longevity
- Less likely to produce strong immune response (due to prevention of brain injury)

Cons:

- Requires craniotomy and exposed cortex limits research applications
- Wireless and implant arrays only recently available

DBS

Deep Brain Stimulation

~10,000 patients

An invasive technique that modulates brain activity with surgically implanted electrodes embedded deep in the brain. DBS can reduce motor, sensory, and behavioral symptoms, usually by the global pallidum, nucleus ventralis interna, thalamus, or subthalamic nucleus.

Major Applications:

- Essential tremor
- Cluster headaches
- Dystonia
- Epilepsy
- OCD
- Huntington's
- Sleep disorders
- Tourette Syndrome
- Movement disorders

Pros:

- Many promising clinical applications in trials
- Improved electrical targeting
- More surgically familiar with implantation process
- Relatively mature manufacturing ecosystem

Cons:

- Requires craniotomy
- Potential brain tissue damage
- Minimal health/side-effects for many
- Limited longevity challenges

Implanted Microelectrodes

~1,000s of patients

Thin electrodes (thickness under 50 μm) delivered via craniotomy, used for electrophysiology, for recording neural signals, and/or stimulating the brain.

Major Applications:

- ALS
- Brain-machine interfaces
- Epilepsy
- Parkinson's Disease
- Motor neuron disease
- Movement disorders
- Sensory cortex injury
- OCD
- Sensory cortex injury
- Sensory cortex injury

Pros:

- Recent influx of private R&D funding
- High spatial resolution
- Many potential prosthetics applications
- Portable
- Lower surgical risk than ECoG, DBS

Cons:

- Requires craniotomy
- Potential brain tissue damage
- Low maximal longevity

fMRI

Functional Magnetic Resonance Imaging

~100,000 patients

Imaging technique that uses magnetic fields to detect changes in cerebral blood flow, as a marker for brain activity. Specifically, fMRI measures blood oxygenation level dependent (BOLD) signals. BOLD signals are derived from deoxygenated hemoglobin (dHb), which is paramagnetic, and oxygenated hemoglobin (oHb), which is diamagnetic. The difference in magnetic susceptibility between the two is what allows fMRI to detect changes in blood flow.

Major Applications:

- Major depression
- Schizophrenia
- ADHD
- Memory studies
- Chronic pain
- Schizophrenia
- Memory studies
- Memory studies

Pros:

- No surgery required
- High temporal resolution
- Scalable imaging
- Ability to do behavioral analysis

Cons:

- Large form factor
- Low temporal resolution
- Phys. pressure for neuronal activity

fNIRS

Functional Near-Infrared Spectroscopy

~10,000 patients

A neuroimaging technique that measures hemoglobin concentration in specific brain regions using a near-infrared light source (~650-1000 nm) and a detector that measures the amount of light that is scattered back to the source. The difference in absorption between the two is what allows fNIRS to detect changes in blood flow.

Major Applications:

- Stroke
- TBI
- Memory studies
- Memory studies
- Memory studies

Pros:

- No surgery required
- High temporal resolution
- Scalable imaging
- Ability to do behavioral analysis

Cons:

- Large form factor
- Low temporal resolution
- Phys. pressure for neuronal activity

BRAIN SENSING

Primary motor cortex (M1) target for movement control

Primary somatosensory cortex (S1) also used here to impact sensory function

Basal Ganglia (BG) target for movement disorders

Frontal lobe (Executive function) MEGs detect high-order executive/attentional activities here

Wernicke's Area (Language comprehension) MEG target for language function restoration

Broca's Area (Language production) MEG target for language function restoration

Insular Cortex (Pain, emotion, self-awareness, empathy) MEG target for altered states of consciousness, self-awareness, depression

Visual Cortex (Visual function) MEGs used here to restore form vision or reconstruct visual input

Motor Cortex (Control, planning, execution, learning, heart rate, sleep, hunger, thirst, satiety, sex drive, etc.) MEG target to treat Parkinson's, epilepsy

Temporal Lobe (Memory, hearing, taste, smell, emotion generation, face recognition) MEG target for auditory prosthetics, memory prosthetics

Brain Sensing (Sensing and motor roles, regulation of consciousness, sleep, attention) MEG target to treat Parkinson's, epilepsy

BRAIN STIMULATING

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

ftCD/tFUS

Focused Transcranial Doppler/Transcranial Focused Ultrasound Stimulation

~100 patients

ftCD is an imaging technique that uses a probe to transmit ultrasound waves and the brain to determine which areas of the brain that may correspond to neural activation. tFUS delivers low-intensity, pulsed ultrasonic waves to the brain to directly modulate specific neural pathways. tFUS should be distinguished from high-intensity ultrasound, which is ablative.

Major Applications:

- Stroke
- TBI
- Memory studies
- Memory studies
- Memory studies

Pros:

- No surgery required
- High spatial resolution
- Reaches deep brain regions
- Inherently imaging capabilities
- Compatible with MRI and EEG
- Adaptable for bedside/clinic
- Many promising clinical applications in trials

Cons:

- Long-term effects not well established

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

ftCD/tFUS

Focused Transcranial Doppler/Transcranial Focused Ultrasound Stimulation

~100 patients

ftCD is an imaging technique that uses a probe to transmit ultrasound waves and the brain to determine which areas of the brain that may correspond to neural activation. tFUS delivers low-intensity, pulsed ultrasonic waves to the brain to directly modulate specific neural pathways. tFUS should be distinguished from high-intensity ultrasound, which is ablative.

Major Applications:

- Stroke
- TBI
- Memory studies
- Memory studies
- Memory studies

Pros:

- No surgery required
- High spatial resolution
- Reaches deep brain regions
- Inherently imaging capabilities
- Compatible with MRI and EEG
- Adaptable for bedside/clinic
- Many promising clinical applications in trials

Cons:

- Long-term effects not well established

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

Metabolic

Transcranial Electrical Stimulation (TACS, tDCS)

Transcranial Magnetic Stimulation (rTMS)

Vagus Nerve Stimulation (VNS)

Optogenetics

Endovascular Electrodes (EE)

*Our report focuses on BCI with more potential for downstream applications. Technologies which do not directly modulate or sense brain activity are not included. This report is not intended to be a comprehensive review of the field. For more information, please visit brainmind.org/bci. This report is not intended to be a comprehensive review of the field. For more information, please visit brainmind.org/bci. This report is not intended to be a comprehensive review of the field. For more information, please visit brainmind.org/bci.

Brain Robotic Interface

- Australia, UTS: VR to control robotic dogs using EEG.
Dry graphene sensors, not as accurate as wet. Can it be useful?



VIRTUAL BR41N.IO HACKATHON

📅 April 17-18, 2021

during the

Spring School 2021*



*BR41N.IO and Spring School 2021 are part of g.tec's Teaching Plan 2021 with more than 140 hours of online courses and lectures.



1. PLACE WINNER

"NeuroBeat"

BCI application

Team members: Alicja Wicher, Joanna Maria Zalewska, Weronika Sójka, Ivo John Krystian Dereziński, Krzysztof Tołpa, Lukasz Furman, Sławomir Duda

IMPROVING HUMAN DAILY LIFE FUNCTIONING

NEUROHACKATOR 2021

21. - 23.
MAY 2021 //
ONLINE

SATURDAY

Project development
in groups



STARTS
10 a.m.

SUNDAY

Evaluation



ENDS
10 a.m.

FRIDAY

Organisers
presentation



workshops
with Judges

working 24h

REQUIREMENTS :

1. Create a team consisting of **3-5 people**.
2. Fill in the Registration Form (available on Facebook event).

DO YOU HAVE ANY QUESTIONS?

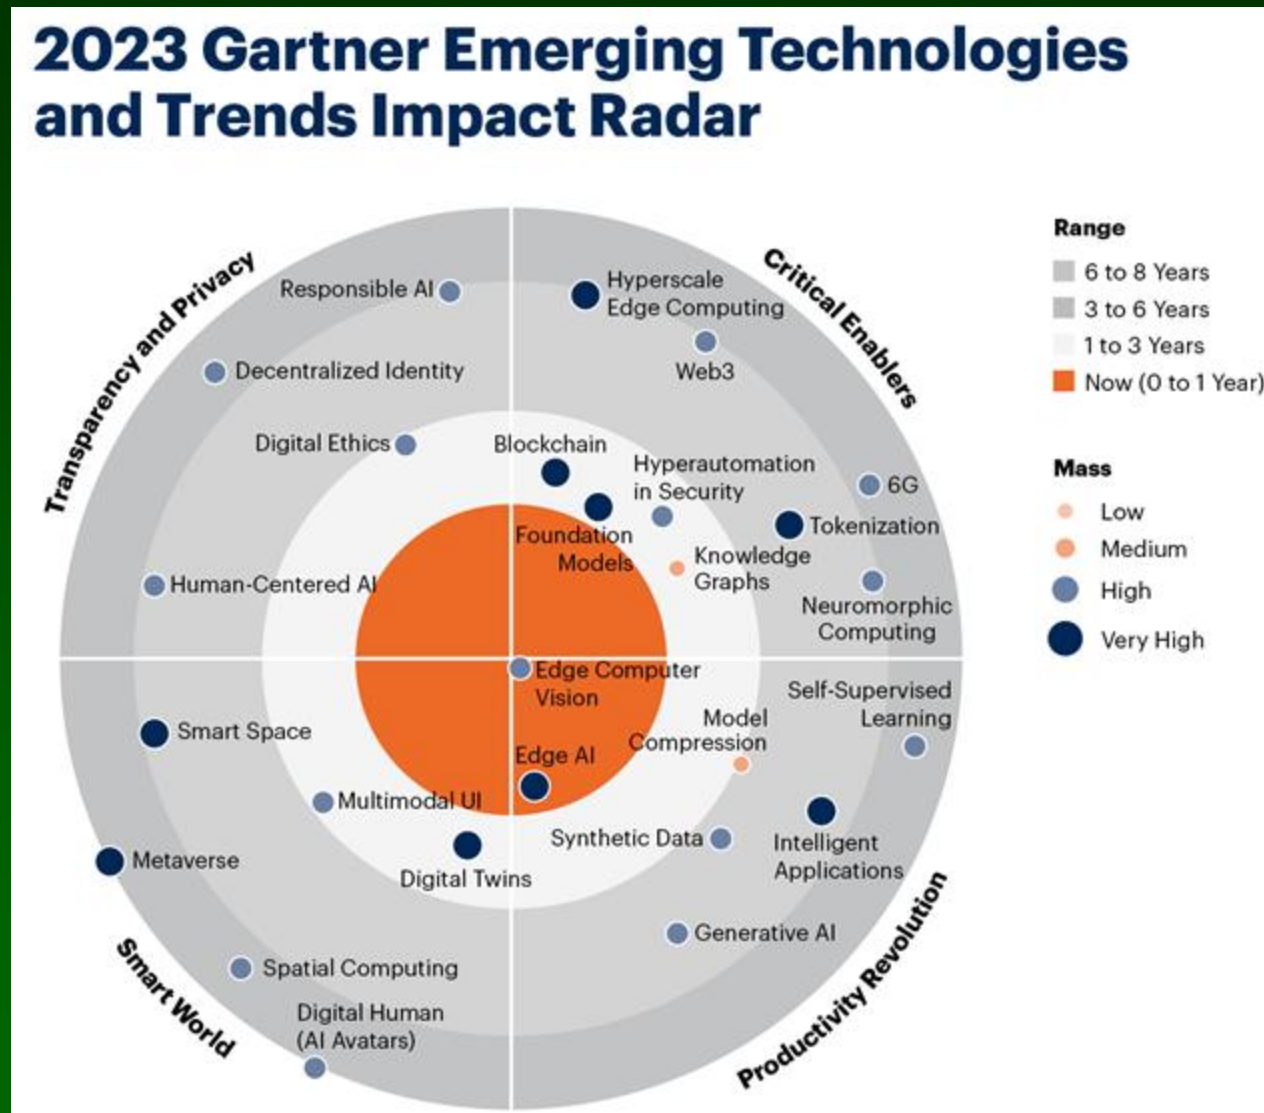
Write an e-mail:
NEUROTECHTOR@GMAIL.COM

Neurotechnology Scientific Club
Center for Modern Interdisciplinary Technologies
at Nicolaus Copernicus University in Toruń
Wileńska 4 Street

Gartner Emerging Tech Impact Radar

Four Emerging Technologies Disrupting the Next 3-8 Years.

Neuromorphic computing
Self-supervised learning
Metaverse
Human-centered AI



Fears ...

Planning for AGI and beyond

Our mission is to ensure that artificial general intelligence—AI systems that are generally smarter than humans—benefits all of humanity.

Pause Giant AI Experiments: An Open Letter

<https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

We call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4.

Signatures

18363

Add your signature

Signatures as of 8.04.2023.

Powerful AI systems should be developed only once we are confident that their effects will be positive and their risks will be manageable.

Time: [The Only Way to Deal With the Threat From AI? Shut It Down](#)

Jobs and money ...



- The World Economic Forum's "The Future of Jobs Report 2020" predicts AI will replace 85 million and create 97 jobs globally by 2025. Do they really know what is happening?
- AI and automation can free us to pursue careers that give us a greater sense of meaning and well-being. The AI Job That Pays Up to \$335K? Prompt engineering ([Time, 14.04.23](#)), no programming required.
- [Report by PwC](#) estimate – 38% of US jobs will be at high risk of automation by early 2030s. White House Report [AI 2022 report](#).
- WD [LegalTech talk](#) (2021), Fusion [AI Exterro WhatSun](#) solution suite will eliminate up to 99.9% of all inefficient cerebral functions performed by highly paid attorneys, paralegals and technologists.
- [ChatGPT: the 10 Jobs Most at Risk by AI \(businessinsider.com\)](#); impact on the US labor market ([OpenAI, arxiv](#)), summary in [AI explained video](#).
- Musicians, composers, writers want a piece of the cake. Stability AI, generative art company, sued by Getty Images for trillions of \$.
- Europe: RODO may kill LLM technology. [Italy blocked ChatGPT \(Money.pl\)](#).

Base10:

Artificial Intelligence will be the next platform shift—following cloud, social, and mobile. In 2022 AI is where Cloud was in 2006-07.

Generative AI leads the rapid changes.

Our Generative AI universe has over 300 companies, and you can see the full universe on our website along with a condensed market map below, across emerging segments.

BASE10 TREND MAP: GENERATIVE AI

Companies are grouped based on medium produced and segmented by use case within each medium. Companies that offer products across segments are grouped in the segment of the core product offering.



Foundation Conversational AI Technologies

LLM Enabled Chatbot Development Frameworks

Technical Opensource

Big Cloud Frameworks

Chatbot Development Frameworks

NLU Tools

NLU/NLG Design

End-To-End Testing

Speech Synthesis / TTS

Voicebot Orchestration

Knowledge Intensive NLP

Dialog Language Models

Translation

Big Cloud Automatic Speech Recognition / ASR / STT

Focused Big Cloud Automatic Speech Recognition / ASR / STT

Focused Voicebot Development Frameworks

Code LLMs

Foundation Large Language Model Stack

Writing Assistant

Content & Idea Creation

Generative & Search Assistants

Data Extraction & Conversational Search

Developer/ Coding Assistants

LLM API Build Frameworks

Prompt Engineering Tools

Data Centric Tooling

Models & Hubs

Foundation LLM Models

Open-Sourced LLM Models

Code Models

Model Hubs

End User Applications

LLM Development Tools

Foundation Models & Hubs

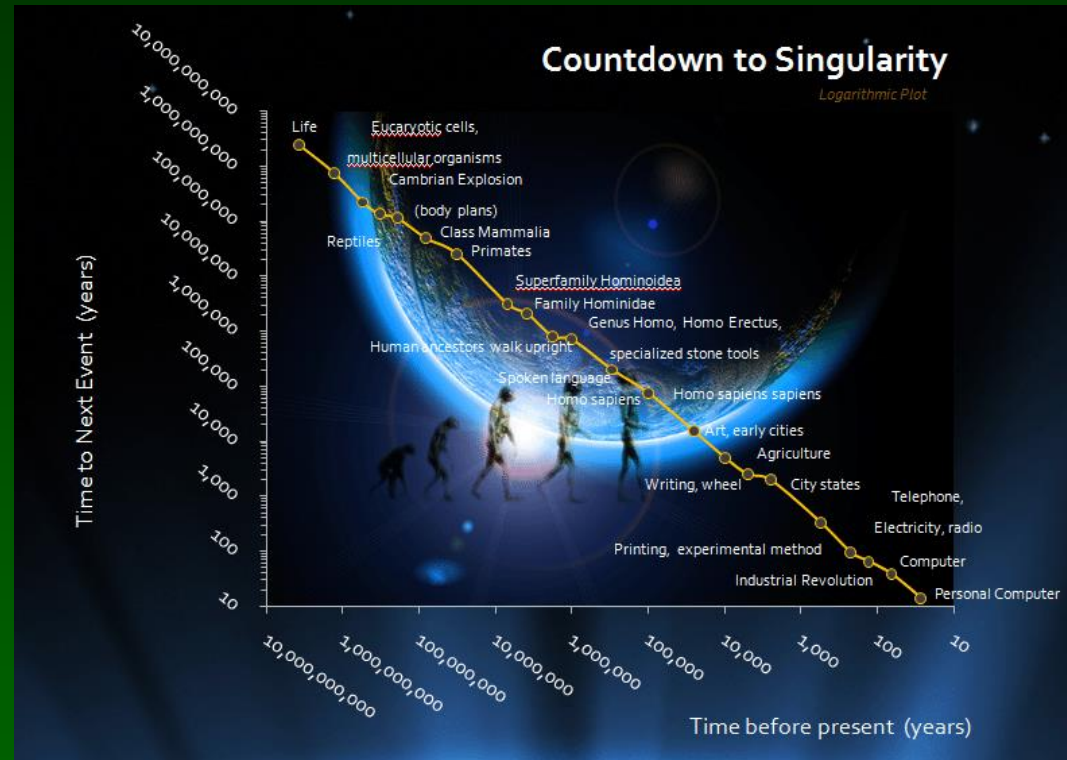
Are we close to the Singularity?

Steps:

1. Autonomous AI.
2. Superhuman level.
3. Brain-computer interfaces for human augmentation.
4. Neurotechnologies to restructure our brains.

Are we on the threshold of a pleasant dream, or a nightmare?

Is transhuman society around the corner?



Each new technological revolution comes faster than the previous one.

Towards Human-like Intelligence

IEEE Computational Intelligence Society Task Force,
Towards Human-like Intelligence



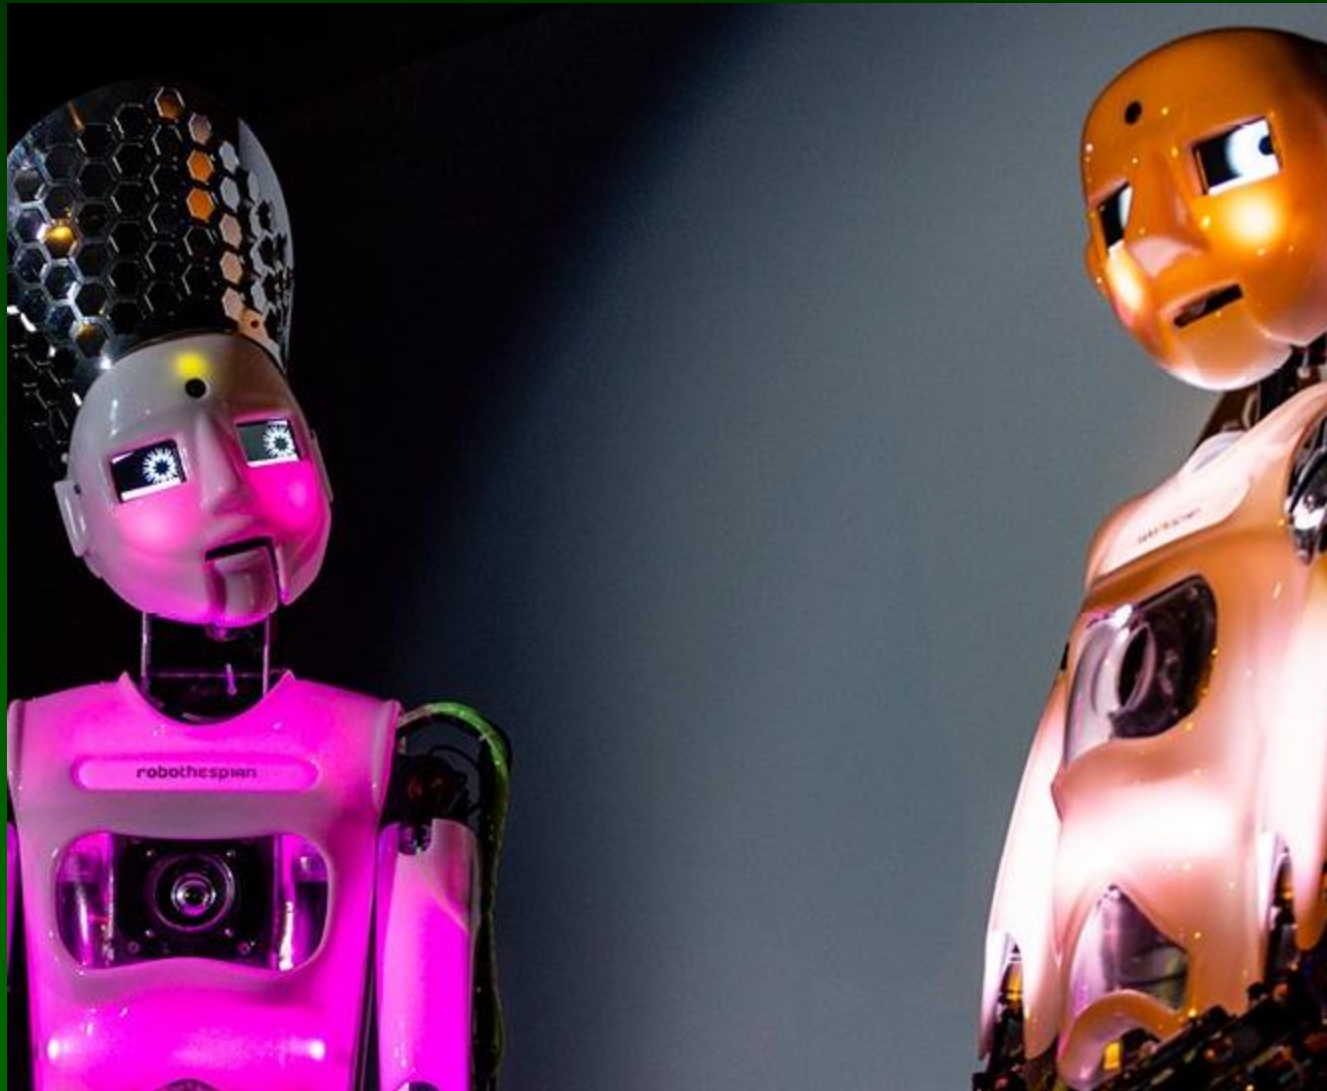
IEEE SSCI CIHLI 2022 Symposium on Computational Intelligence for Human-like Intelligence, Singapore (J. Mandziuk, W. Duch, M. Woźniak).

AGI conference, Journal of Artificial General Intelligence, comments on Cognitive Architectures and Autonomy: A Comparative Review (eds. Tan, Franklin, Duch).

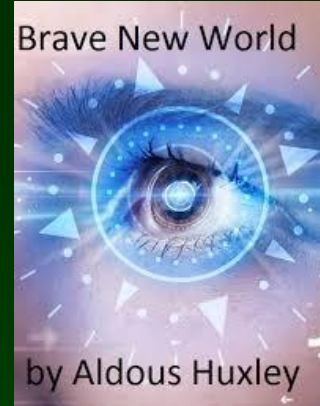
BICA Annual International Conf. on Biologically Inspired Cognitive Architectures, 13th Annual Meeting of the BICA Society, Guadalajara, Mexico 2023.

Brain-Mind Institute Schools International Conference on Brain-Mind (ICBM) and Brain-Mind Magazine (Juyang Weng, Michigan SU).

Stanislaw Lem: About prince Ferrycy and princess Crystala.
Intelligent Palefaces? Is it possible?

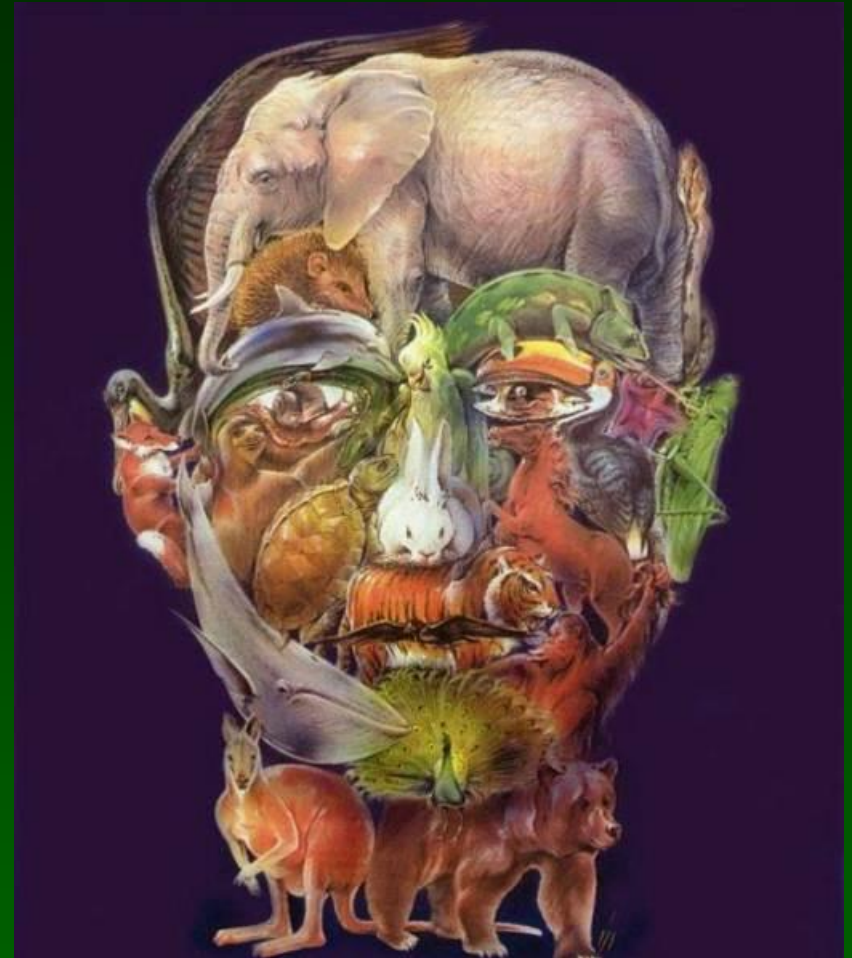


Perspectives



- What was impossible yesterday tomorrow will be common.
- AI is changing many things, including the way science is done. Expect accelerated scientific discoveries (ex. AlphaFold).
- AI-based automation will lead to a great social changes, we have no idea what will happen. LLMs are not just chat apps, but will control autonomous AI systems. Auto-prompts, plugins, and software tools will lead to AGI.
- Large companies are at the front of research but smaller project appear.
- Emergent AI thinking in multidimensional spaces will reach superhuman level. Robots/AI systems will quickly learn from each other.
- Programming may gradually become obsolete. Support for AI ecosystems (lablab.ai, [k4all](https://k4all.com)), EU consortia, educational ecosystems, should be a priority.
- Machines will claim to be conscious, and some experts already accept it.
- In the **long run** BCI, neurocognitive technologies, will help to integrate our selves with AI. The legal status of cyborgs is already being discussed.
- **Our megalomania is dangerous. We are moving from animal life to our own virtual creation! Singularity may come faster than we think!**

Intelligence?



Search: Wlodek Duch
=> talks, papers, lectures ...

