I. 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

Google: 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.

Development of civilization

We are in extraordinary moment in the history of the world!

Growing understanding of the world, since antiquity:

1. **Magical thinking**: the whims of the gods, fatalism.
2. **Protoscience**: empirical observations, causality, descriptive knowledge.
3. **Classical science**: theories, empirical verification, math and statistics.
4. **Computer simulations**: complex systems, “new kind of science” (Wolfram).
5. **Big data**: knowledge from large amount of data (KDD).
6. **Artificial intelligence**: support for thinking, autonomous AI, emergence.
7. **Superhuman augmentation**: coupling AI with brains, in near future?

2023: AI tools appearing everywhere: browsers, office, Khan Academy.

My 2001 predictions of the AI future. Reification errors.

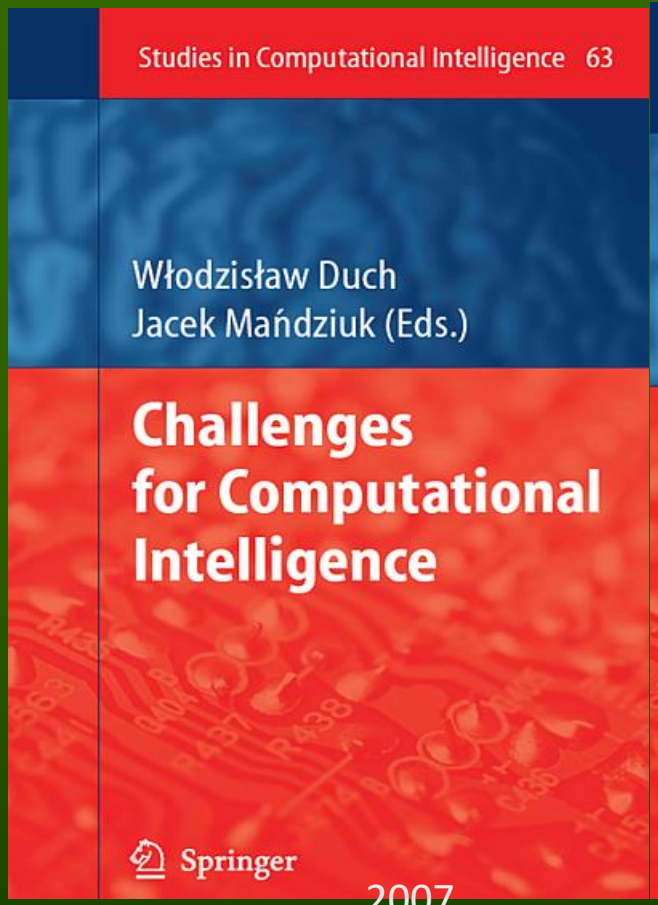

2007

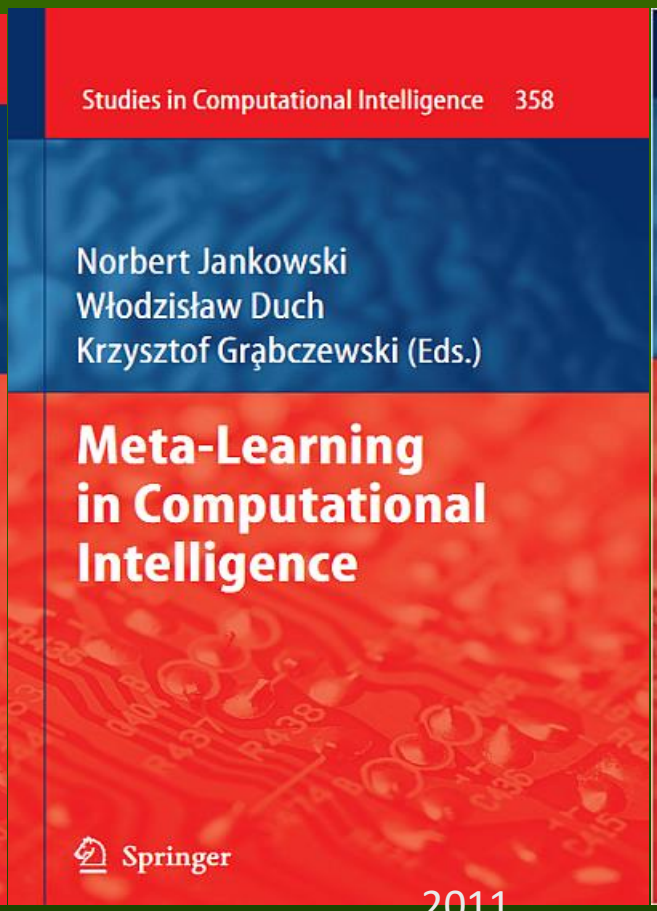

2011

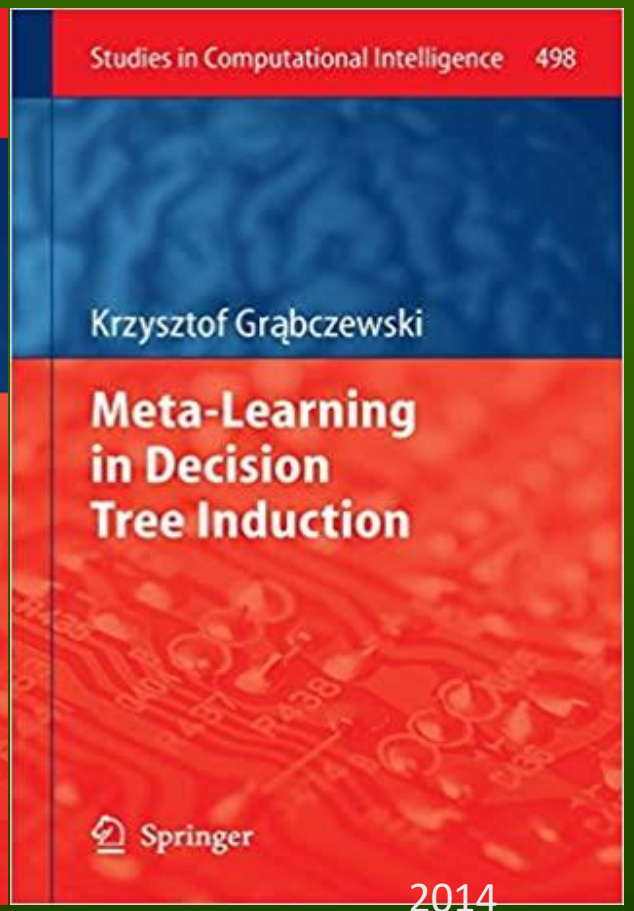

2014

Meta-learning, or learning by search in the model space for useful composition of fine-grained transformations, support feature extraction, novel transfer functions, interesting distributions as new targets for learning and many deep ideas, not simple improvements.
Duch W, Grudziński K. (2001) Meta-learning: searching in the model space. ICONIP, 235;
Duch W, Grudziński K. (2002) Meta-learning via search combined with parameter optimization. IIS, Advances in Soft Computing 17, pp. 13-22

WD: Machine Learning topics Now: transformers.

Ghostminer

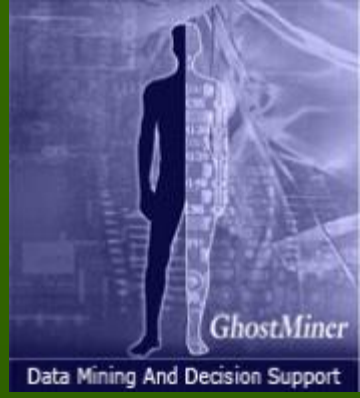

- FQS Ghostminer software
- KIS UMK project 1998-2004,
  data mining, business intelligence.
- GhostMiner is Fujitsu's advanced analytical data mining tool that not only supports a variety of databases (and spreadsheets), advanced machine learning algorithms, but also data preparation and selection, model validation, multi-models such as committees, or k-classifiers and data/model visualization.
- Clients: universities, polytechnics, research institutes, banks and various companies in Poland, Austria, Australia, China, Czech Republic, Netherlands, India, Japan, Canada, Germany, Norway, Singapore, UK and USA.
- E.g. Abbott Laboratories was using GhostMiner to study and discover the properties of multidimensional scientific data.

AI Timeline

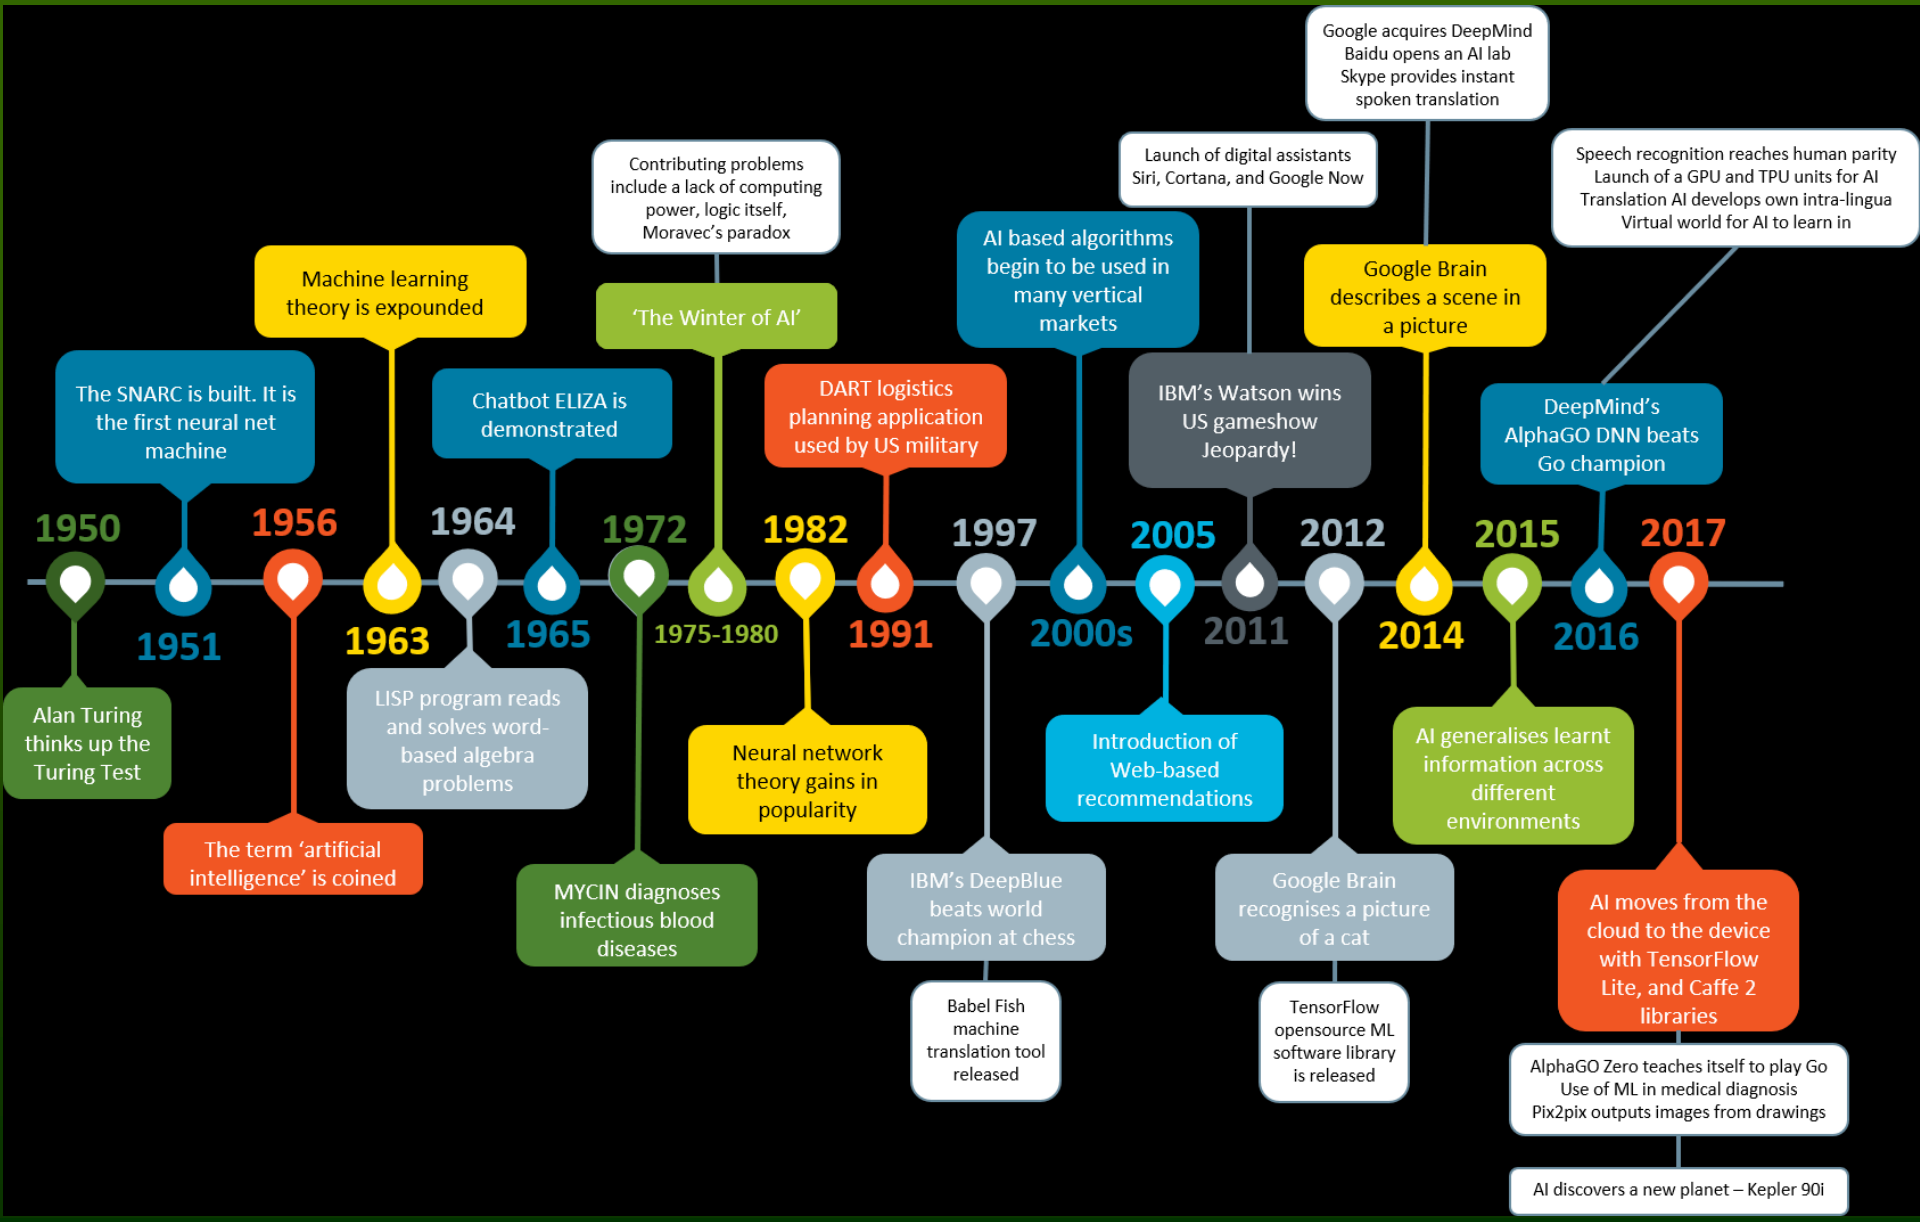
1950
Alan Turing thinks up the Turing Test
1951
The SNARC is built. It is the first neural net machine
1956
The term 'artificial intelligence' is coined
1963
Machine learning theory is expounded
1964
LISP program reads and solves word-based algebra problems
1965
Chatbot ELIZA is demonstrated
1972
MYCIN diagnoses infectious blood diseases
1975-1980
Contributing problems include a lack of computing power, logic itself, Moravec's paradox
'The Winter of AI'
1982
Neural network theory gains in popularity
1991
DART logistics planning application used by US military
1997
IBM's DeepBlue beats world champion at chess
Babel Fish machine translation tool released
2000s
AI based algorithms begin to be used in many vertical markets
2005
Introduction of Web-based recommendations
2011
IBM's Watson wins US gameshow Jeopardy!
2012
Google Brain recognises a picture of a cat
TensorFlow opensource ML software library is released
2014
Google Brain describes a scene in a picture
Launch of digital assistants Siri, Cortana, and Google Now
Google acquires DeepMind
Baidu opens an AI lab
Skype provides instant spoken translation
2015
AI generalises learnt information across different environments
2016
DeepMind's AlphaGO DNN beats Go champion
2017
AI moves from the cloud to the device with TensorFlow Lite, and Caffe 2 libraries
AlphaGO Zero teaches itself to play Go
Use of ML in medical diagnosis
Pix2pix outputs images from drawings
AI discovers a new planet – Kepler 90i
Speech recognition reaches human parity
Launch of a GPU and TPU units for AI
Translation AI develops own intra-lingua
Virtual world for AI to learn in

AI First Wave

The first wave of AI

DARPA

Engineers create sets of rules
to represent knowledge in
well-defined domains

The **structure** of the knowledge is defined by humans
The **specifics** are explored by the machine

UiPath Business Automation Platform - Leader in Robotic Process Automation.
Founded in Romania in 2005, valuation >1 B$ in 2018, >40 B$ in 2022

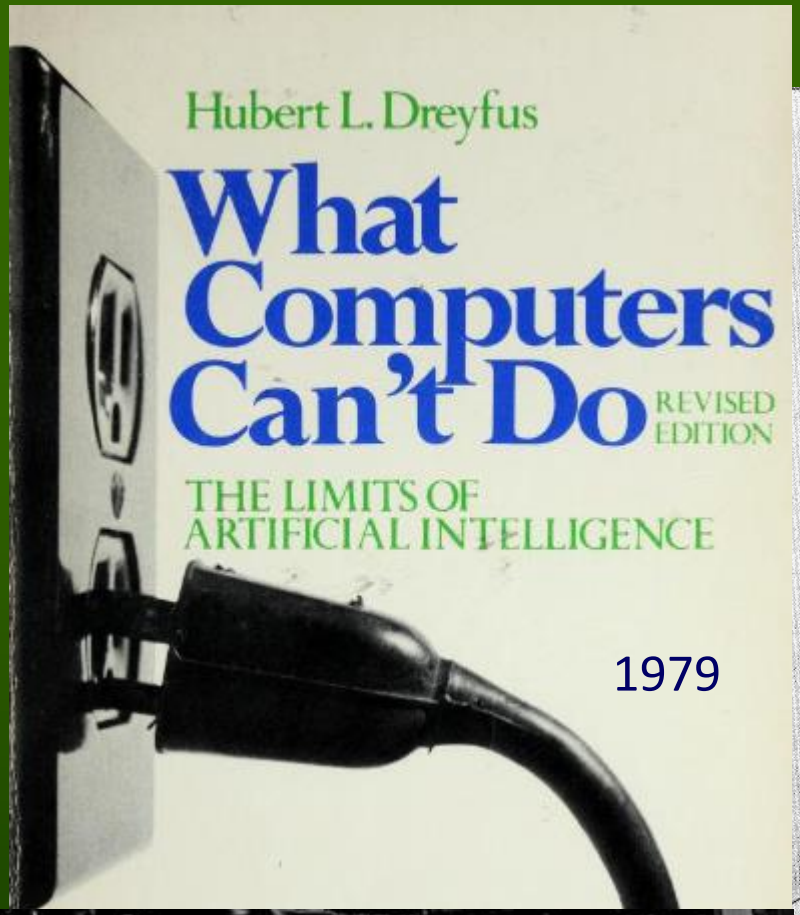
Hubert L. Dreyfus
What Computers Can't Do
REVISED EDITION
THE LIMITS OF ARTIFICIAL INTELLIGENCE
1979

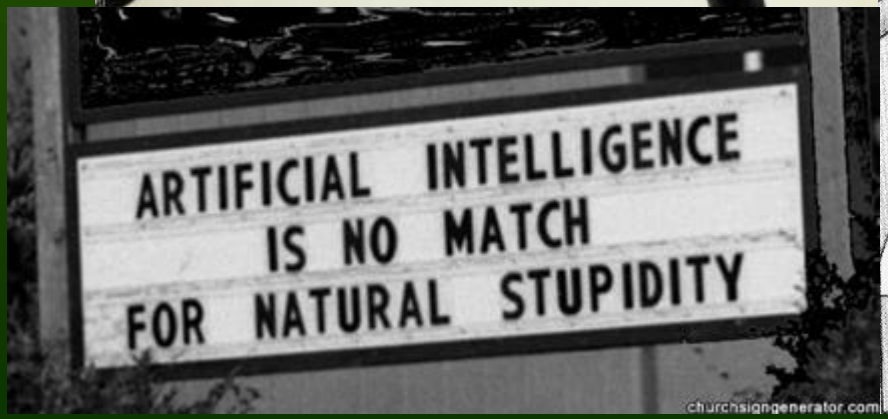
ARTIFICIAL INTELLIGENCE
IS NO MATCH
FOR NATURAL STUPIDITY
churchsigngenerator.com

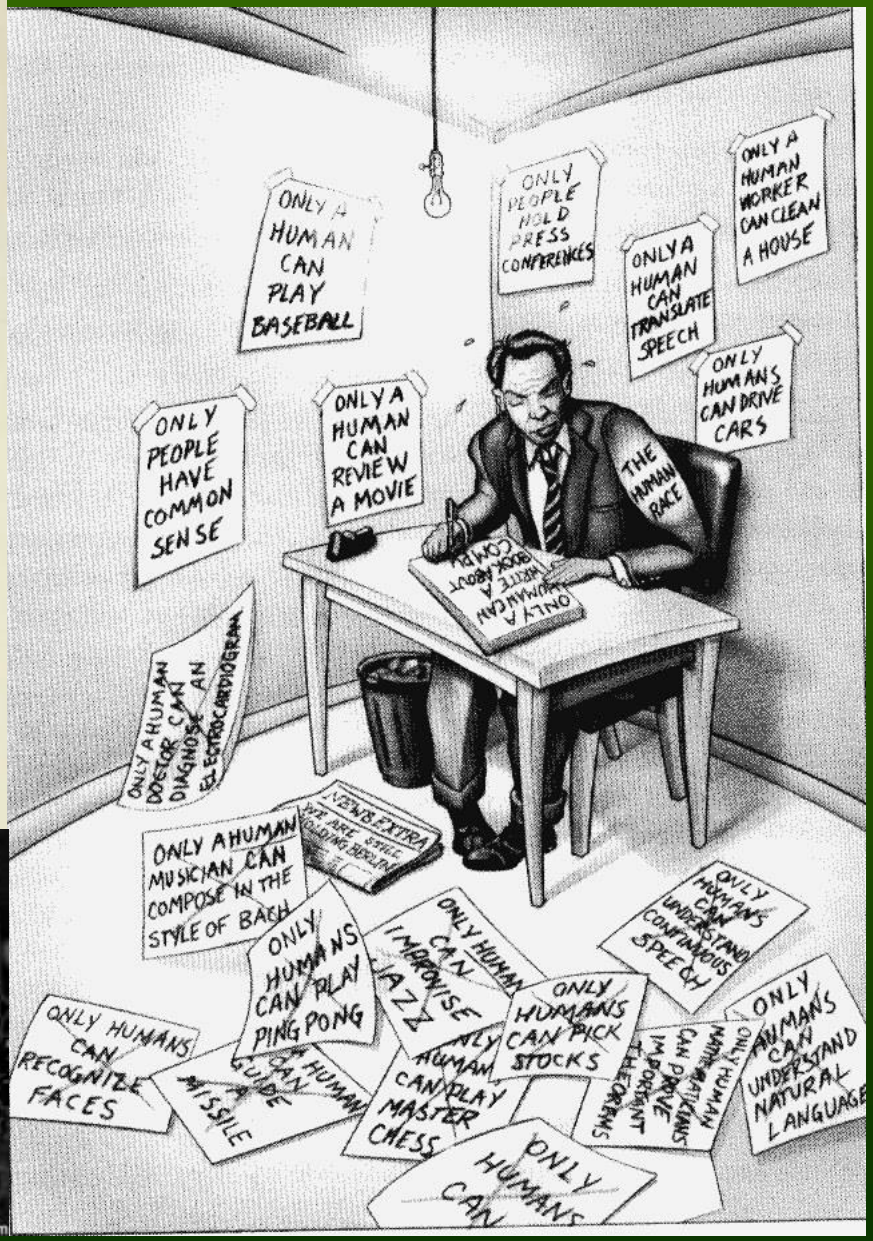
ONLY A HUMAN CAN PLAY BASEBALL
ONLY PEOPLE HOLD PRESS CONFERENCES
ONLY A HUMAN WORKER CAN CLEAN A HOUSE
ONLY A HUMAN CAN TRANSLATE SPEECH
ONLY HUMANS CAN DRIVE CARS
ONLY PEOPLE HAVE COMMON SENSE
ONLY A HUMAN CAN REVIEW A MOVIE
THE HUMAN RACE
ONLY A HUMAN MUSICIAN CAN COMPOSE IN THE STYLE OF BACH
ONLY HUMANS CAN PLAY PING PONG
ONLY HUMANS CAN IMPROVISE JAZZ
ONLY HUMANS CAN PICK STOCKS
ONLY HUMANS CAN RECOGNIZE FACES
ONLY HUMANS CAN GUIDE A MISSILE
ONLY HUMANS CAN PLAY MASTER CHESS
ONLY HUMANS CAN UNDERSTAND NATURAL LANGUAGE

AI First Wave

First wave stumbles

DARPA

2004
completed: 0

Source: DARPA

2005
completed: 5

DARPA Autonomous Vehicle Grand Challenge
140 miles of dirt tracks in California and Nevada

We need computational intelligence, symbolic AI + pattern/signal recognition.

AI Second Wave

The second wave of AI

DARPA

Source: thrillist.com

Statistical Learning

Collect big data in specific domains and create statistical models of data. For example, compare texts translated by people, available in two languages, learn to translate phrases.

Classification and prediction, but abstracting and reasoning is difficult.

Data understanding

Manifolds

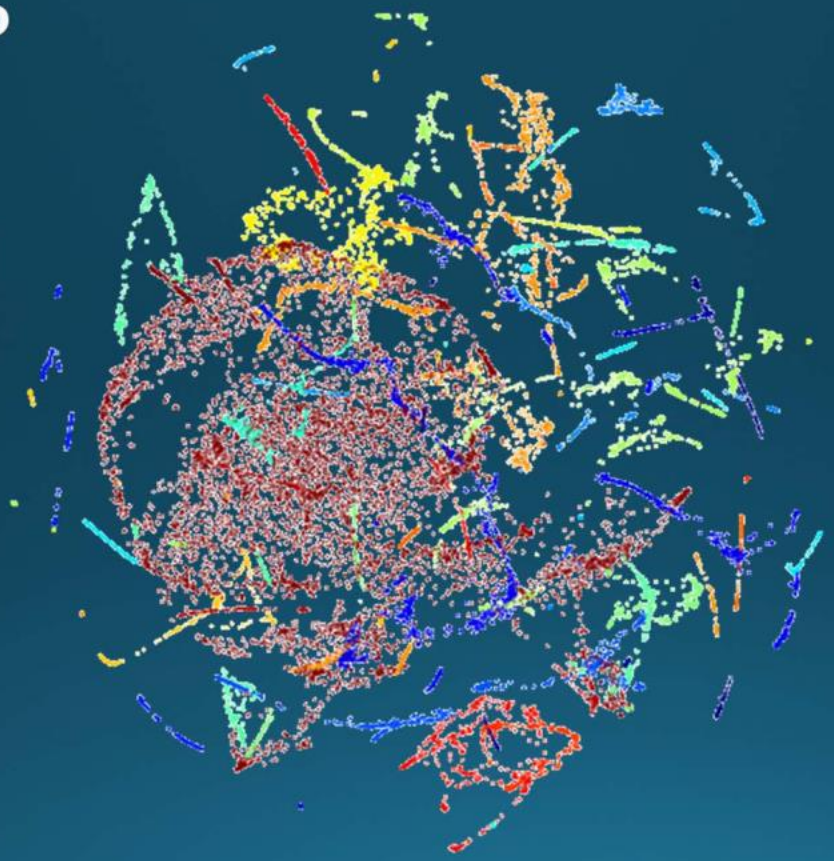

Each manifold represents a different entity

Understanding data comes by separating the manifolds

Images/signals are complex, but have certain similarity in appropriate spaces, with reduced dimensionality! Latent space of parameters that re-create data.

Hand writing

Manifolds of handwriting

Variation in handwritten digits
form 10 distinct manifolds within the
28x28 dimensional space of pixel values

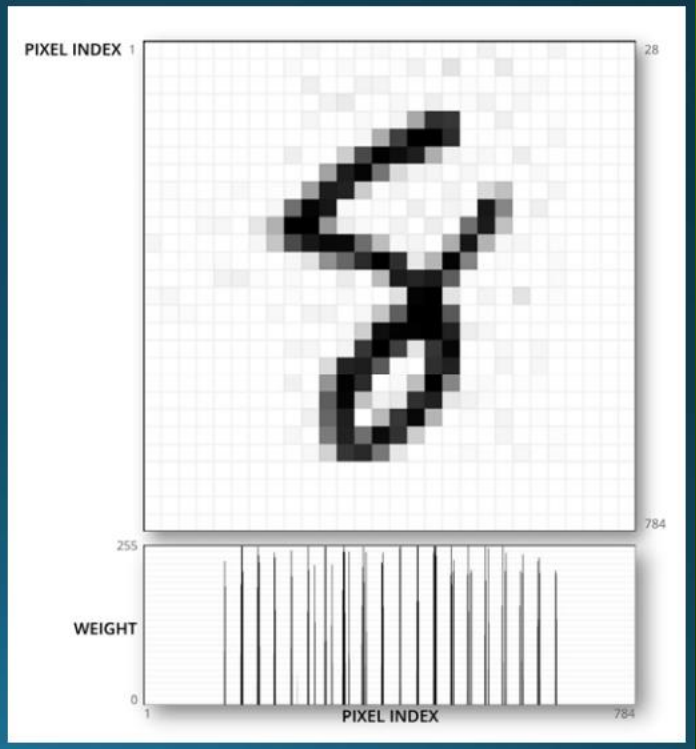

For example, hand written digits and characters show high variability but may be separated on local manifolds in low dimensions.

Deep networks

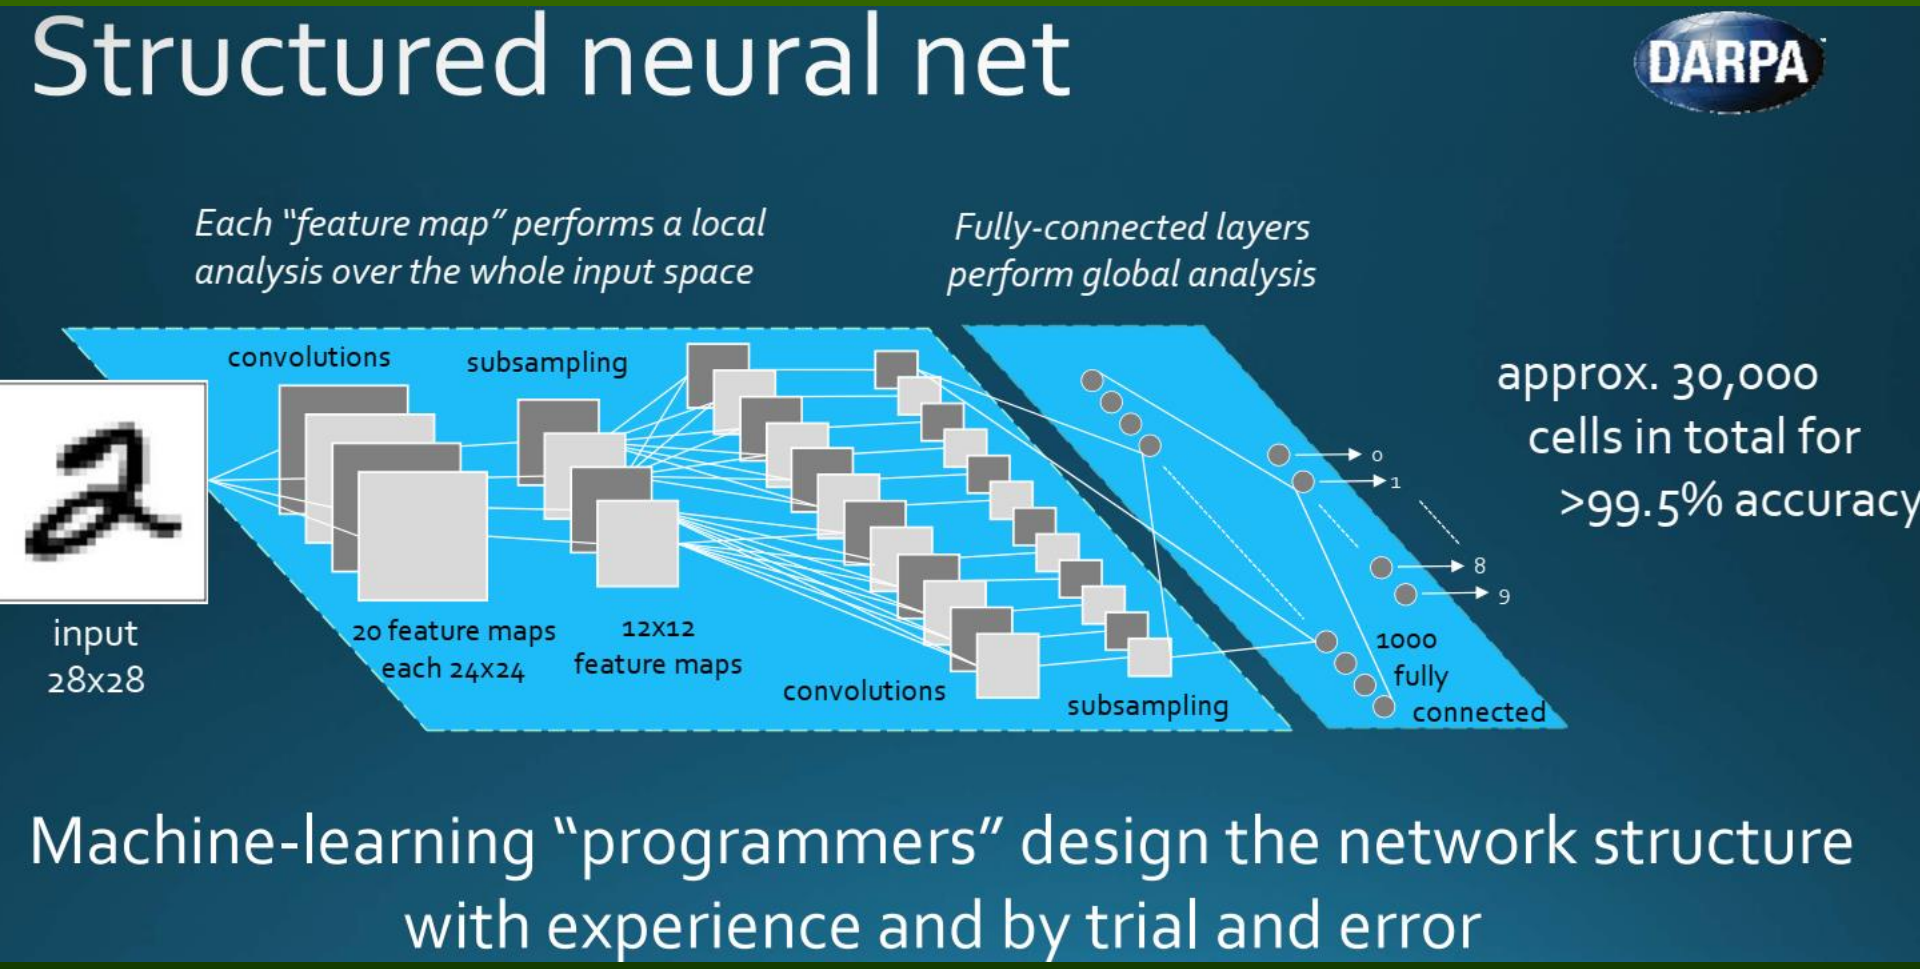

Simple neural nets were developed in 1980-90, but learning from data of complex, deep architectures was impossible – computational resources/algorithms were too poor. Until 2015, when 100 layer network exceeded human ImageNet results.

Similar approach may be applied to words and sentences, leading to progress in Natural Language Processing.

Machine Learning Types

Supervised – school type, corrections after each error, needs labeled data.
Unsupervised – learn to find structures, ex. syllables, words, walk. No labels.
Semi-supervised – small amount of labeled data to initiate, large unlabeled base.
Reinforcement – change your performance strategy.
Deep learning – many transformations, large network.

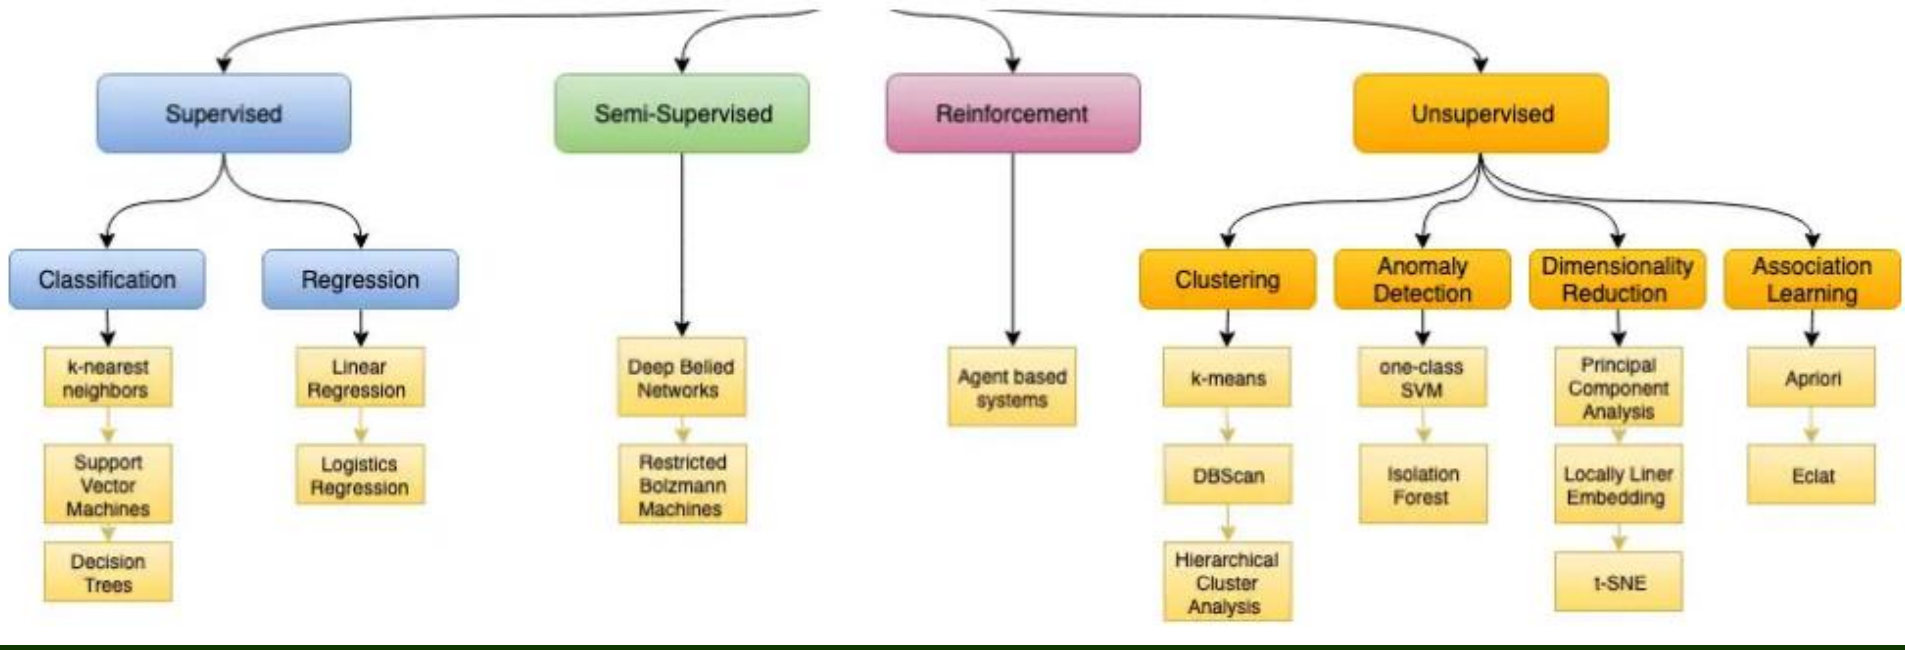

ML is easy - just pick up your method ...

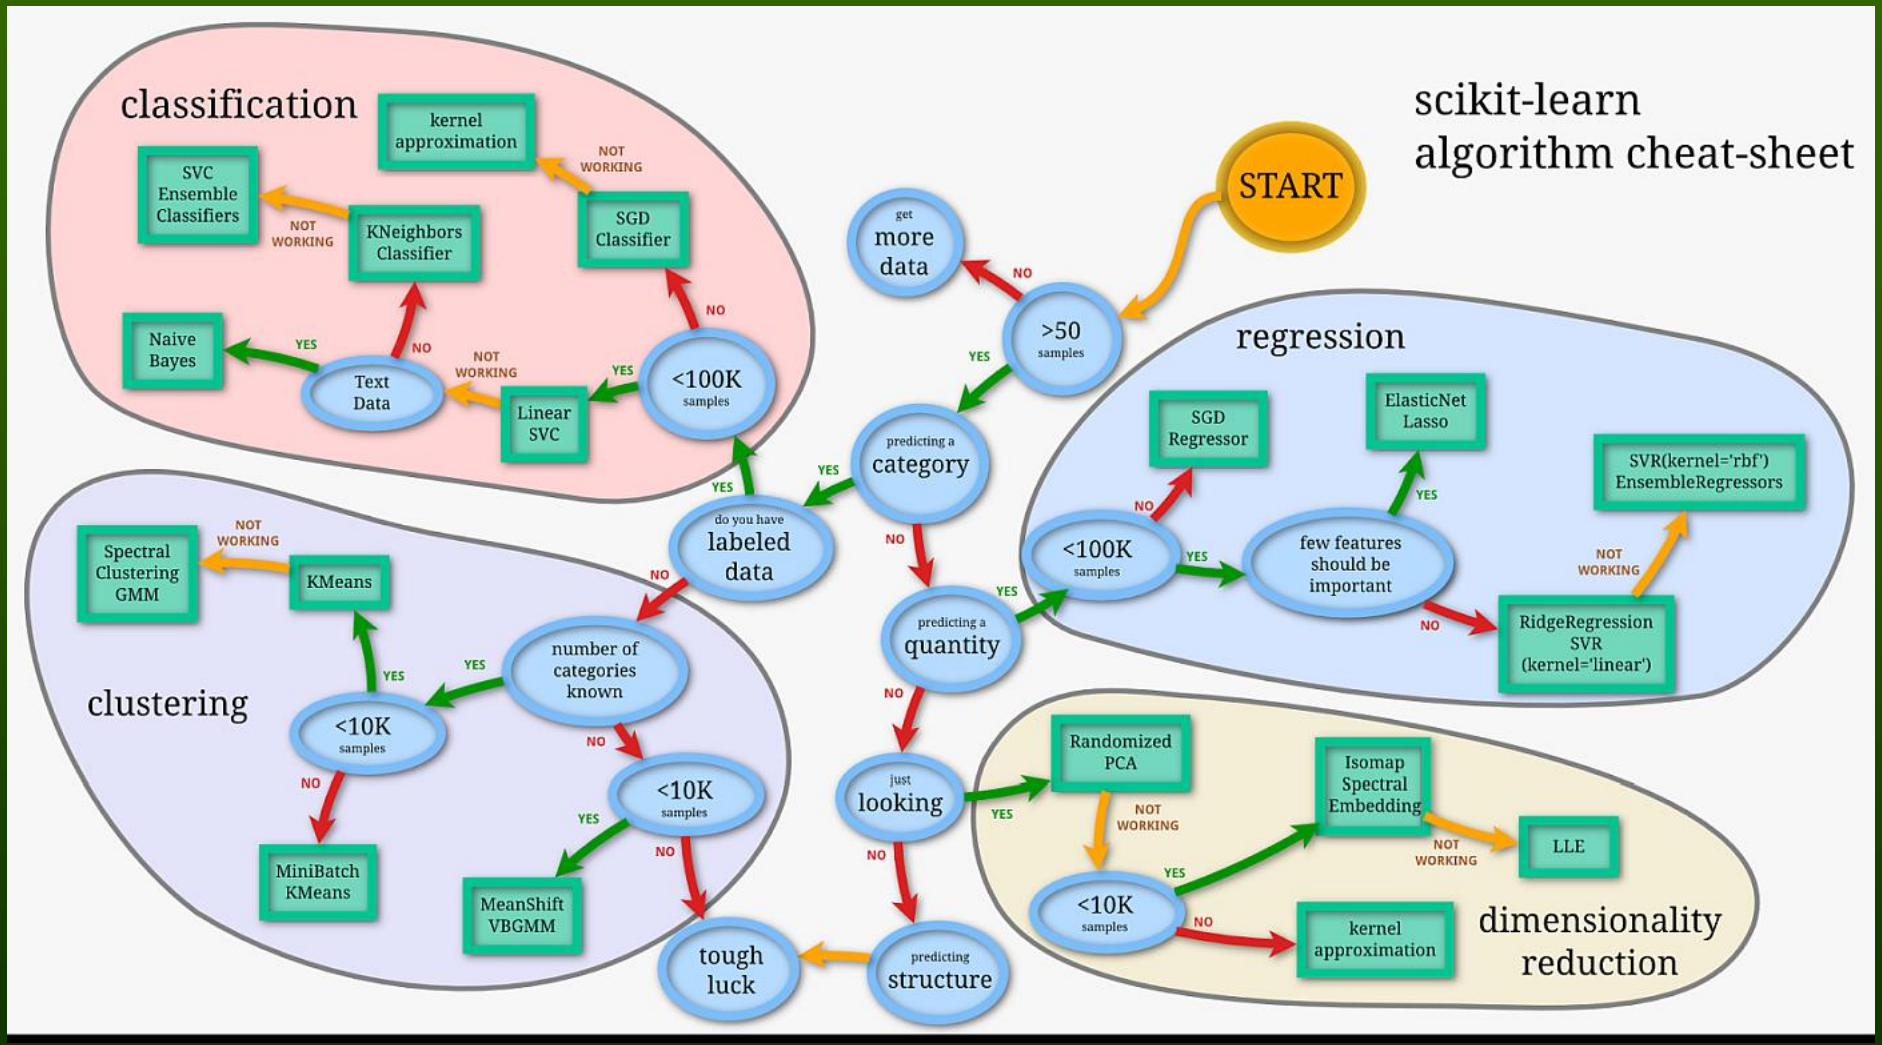

Thousands of applications of machine learning are enabled by free powerful large systems, such as TensorFlow, Scikit-learn, Keras, MS Cognitive services …

Cogni
Cognitive sciences

Biohybrids

Bio
Neuroscience

Nano
Quantum
Technologies

Neurocognitive
Informatics

Info
Artificial/Computational Intelligence,
Machine Learning, Neural Networks

DREAM top-level architecture

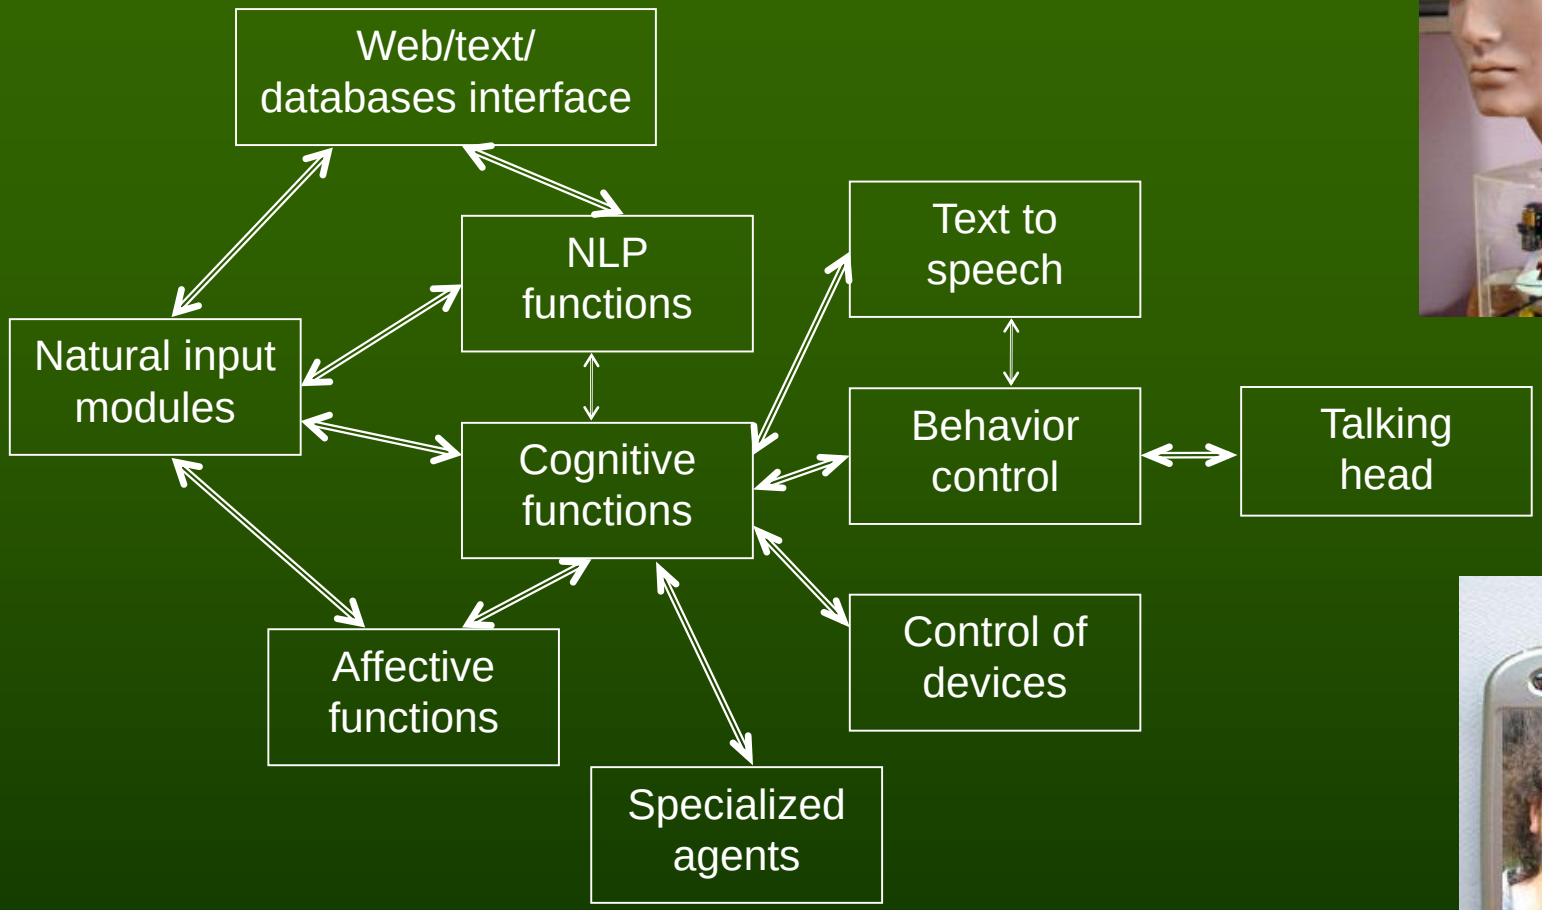

DREAM project (2003), focused on perception (visual, auditory, text inputs), cognitive functions (reasoning based on perceptions), natural language communication in well defined contexts, real time control of the simulated/physical head. Now Amazon, Google, Apple, GPT … even in watches.

Bina48 and LifeNaut Project

Reconstructing the mind from information in mindfiles, creating mindclones: self-aware digital beings, remembering, thinking, feeling. Now Mika in Salzburg.

Neuromorphic future

Wall with 1024 TrueNorth chips, equivalent of 1 Billion neurons, 256 B synapses. 1/6 of chimp brain. Cerebras CS-2 chip has 2600 B transistors, almost 1M cores!

Integration:
Nano +
Neuro +
Info +
Cogni

Neural AI accelerators AD 2022.
Inteli Loihi 2
+ LAVA soft.

Cerebras CS-2 Andromeda system, 10^{18} op/sec!

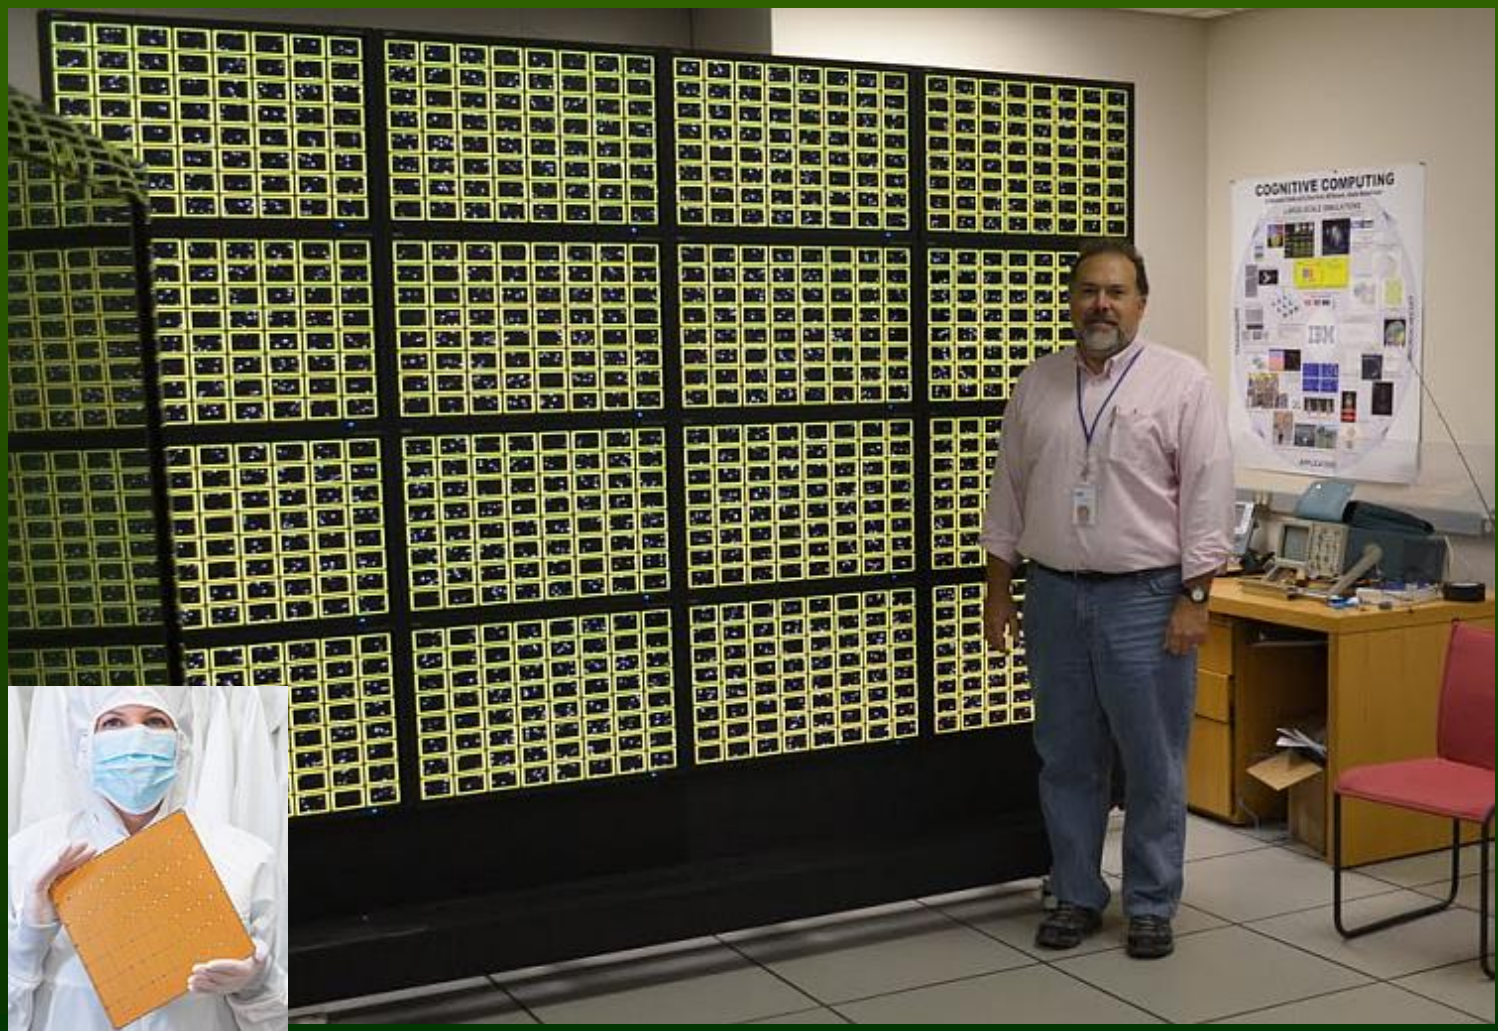

Recently: superhuman AI

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 parcour, autonomous vehicles on the roads.

Automation of science: 2015-AI uncovers genetic and signaling pathways of flatworm regeneration.
2020 AlphaFold 2, now 600 mln 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 beats professionals
2020: BERT answers questions from SQuAD database.

Cyborgization: BCI, brain optimization, coming?

AI in games

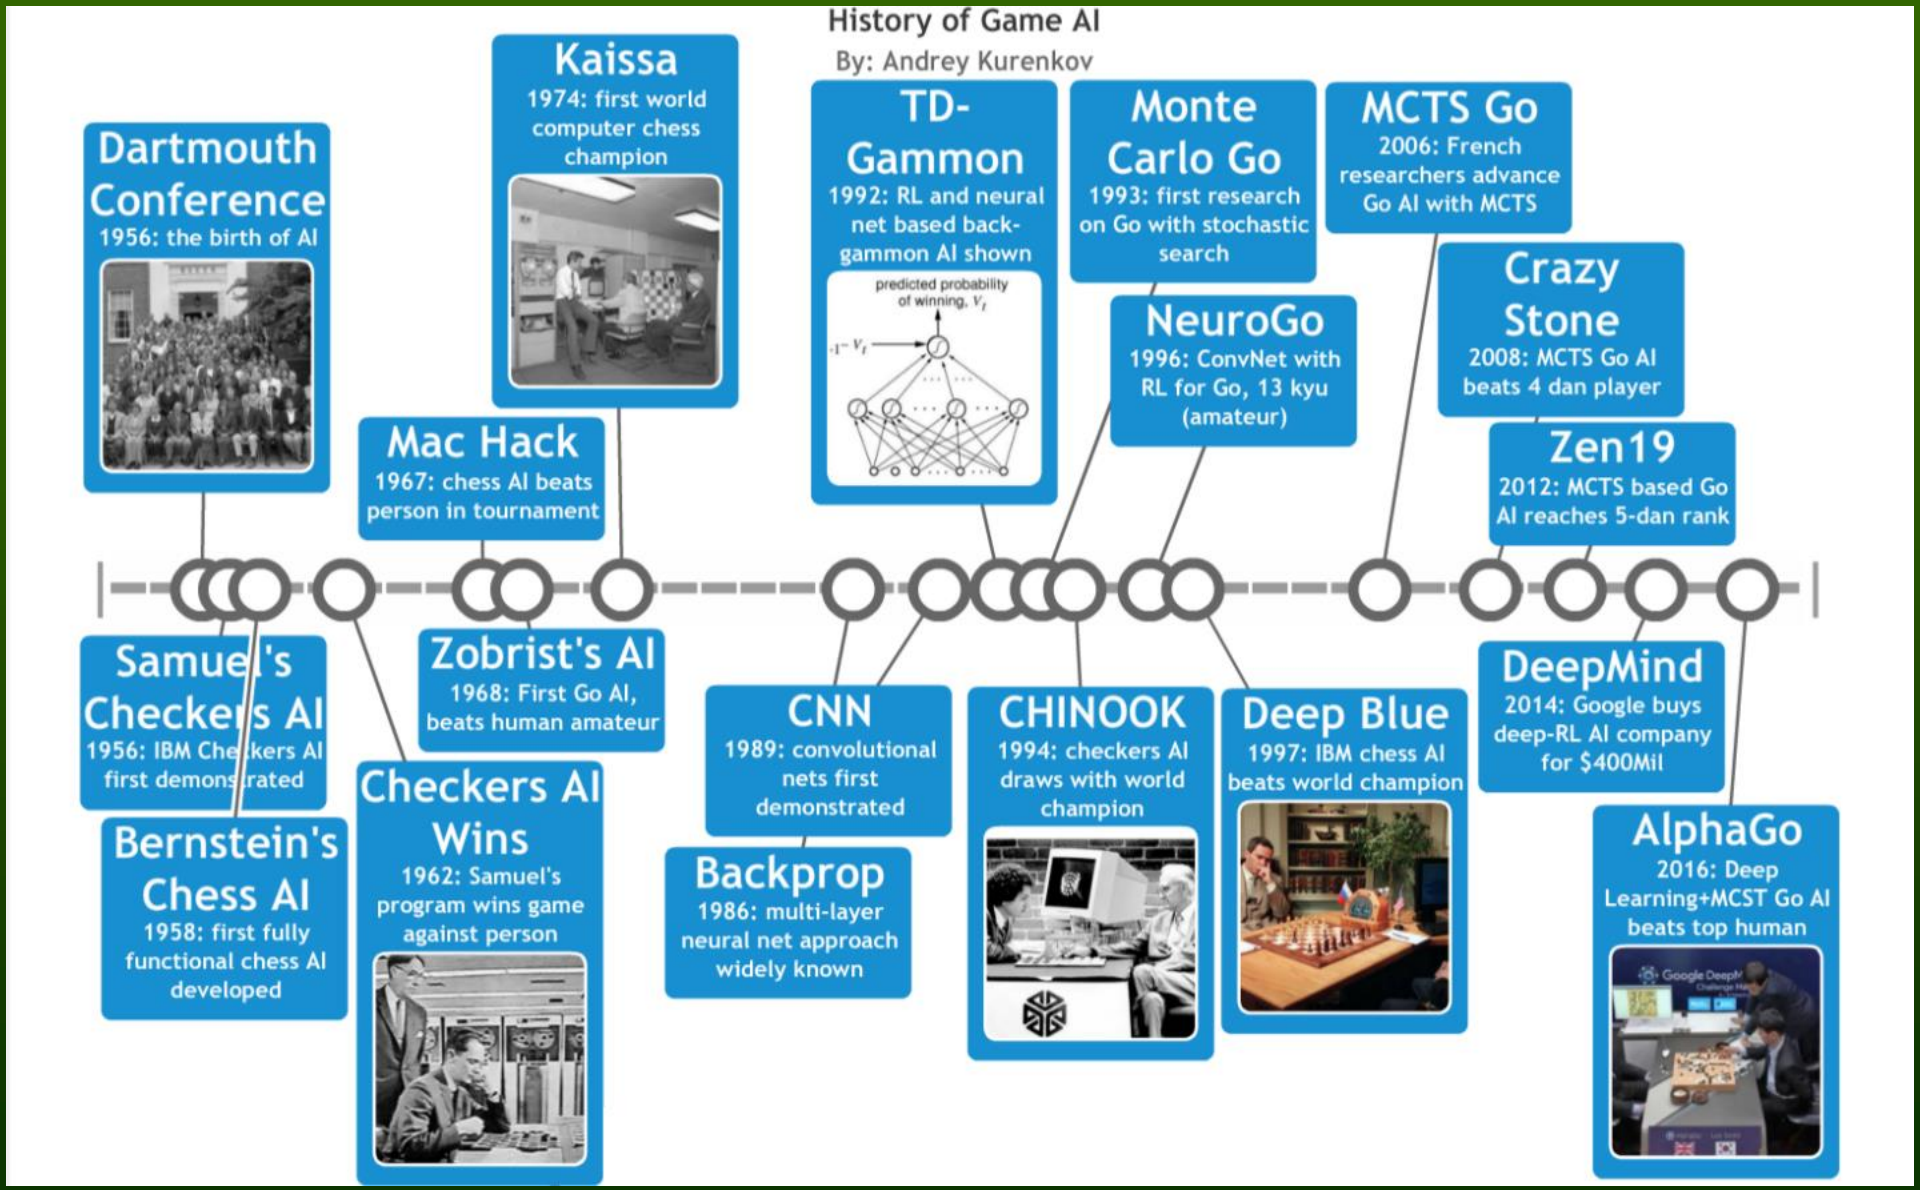
History of Game AI
By: Andrey Kurenkov
Dartmouth Conference
1956: the birth of AI
Kaissa
1974: first world computer chess champion
TD-Gammon
1992: RL and neural net based back-gammon AI shown
predicted probability of winning, V_t
$1-V_t$
Monte Carlo Go
1993: first research on Go with stochastic search
MCTS Go
2006: French researchers advance Go AI with MCTS
Crazy Stone
2008: MCTS Go AI beats 4 dan player
NeuroGo
1996: ConvNet with RL for Go, 13 kyu (amateur)
Zen19
2012: MCTS based Go AI reaches 5-dan rank
Mac Hack
1967: chess AI beats person in tournament
Samuel's Checkers AI
1956: IBM Checkers AI first demonstrated
Zobrist's AI
1968: First Go AI, beats human amateur
CNN
1989: convolutional nets first demonstrated
CHINOOK
1994: checkers AI draws with world champion
Deep Blue
1997: IBM chess AI beats world champion
DeepMind
2014: Google buys deep-RL AI company for $400Mil
Bernstein's Chess AI
1958: first fully functional chess AI developed
Checkers AI Wins
1962: Samuel's program wins game against person
Backprop
1986: multi-layer neural net approach widely known
AlphaGo
2016: Deep Learning+MCST Go AI beats top human

Alpha Go Zero

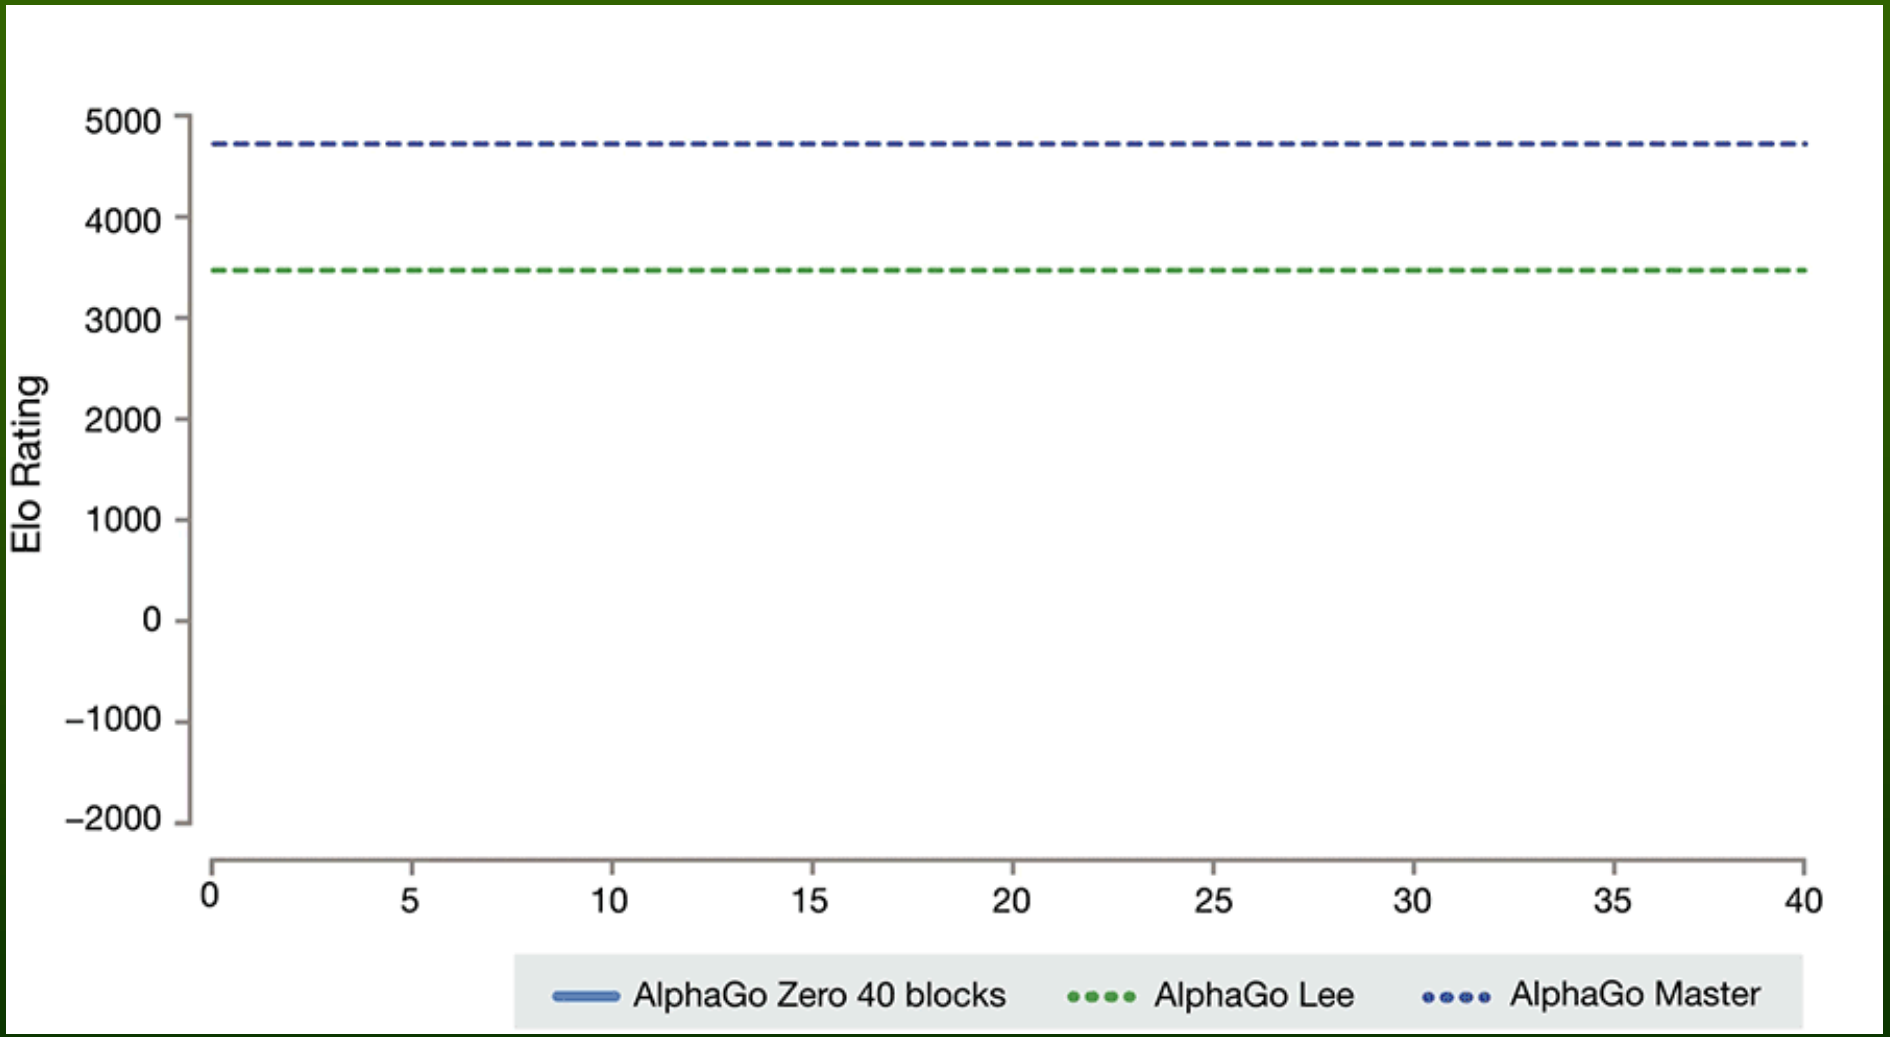

StarCraft II

StarCraft II is one of the most popular games of RTS (real-time strategy), where long-term strategy is important, speed has been limited to give humans a chance, information is incomplete. Complex, open environment.
In Jan. 2019 AlphaStar (DeepMind) has beaten two best human professional players 5:0 – they have been sure they will win.

Deep neural network was trained directly from raw game data via supervised and reinforcement learning techniques.

AlphaStar “use of strategy and fighting techniques had never been seen from human opponents.”

Protein folding

AlphaFold 2 using deep learning predicted more than 2/3 of all protein structures with an accuracy equivalent to experimental!

Nature, 30.11.2020

Structure recognition + learning + inference.

Predicting protein structures based on amino acid sequences is the basis for the search for proteins and the design of drugs with the desired properties.

Prediction of 600 mln protein structures (DM+EMBL-EBI).

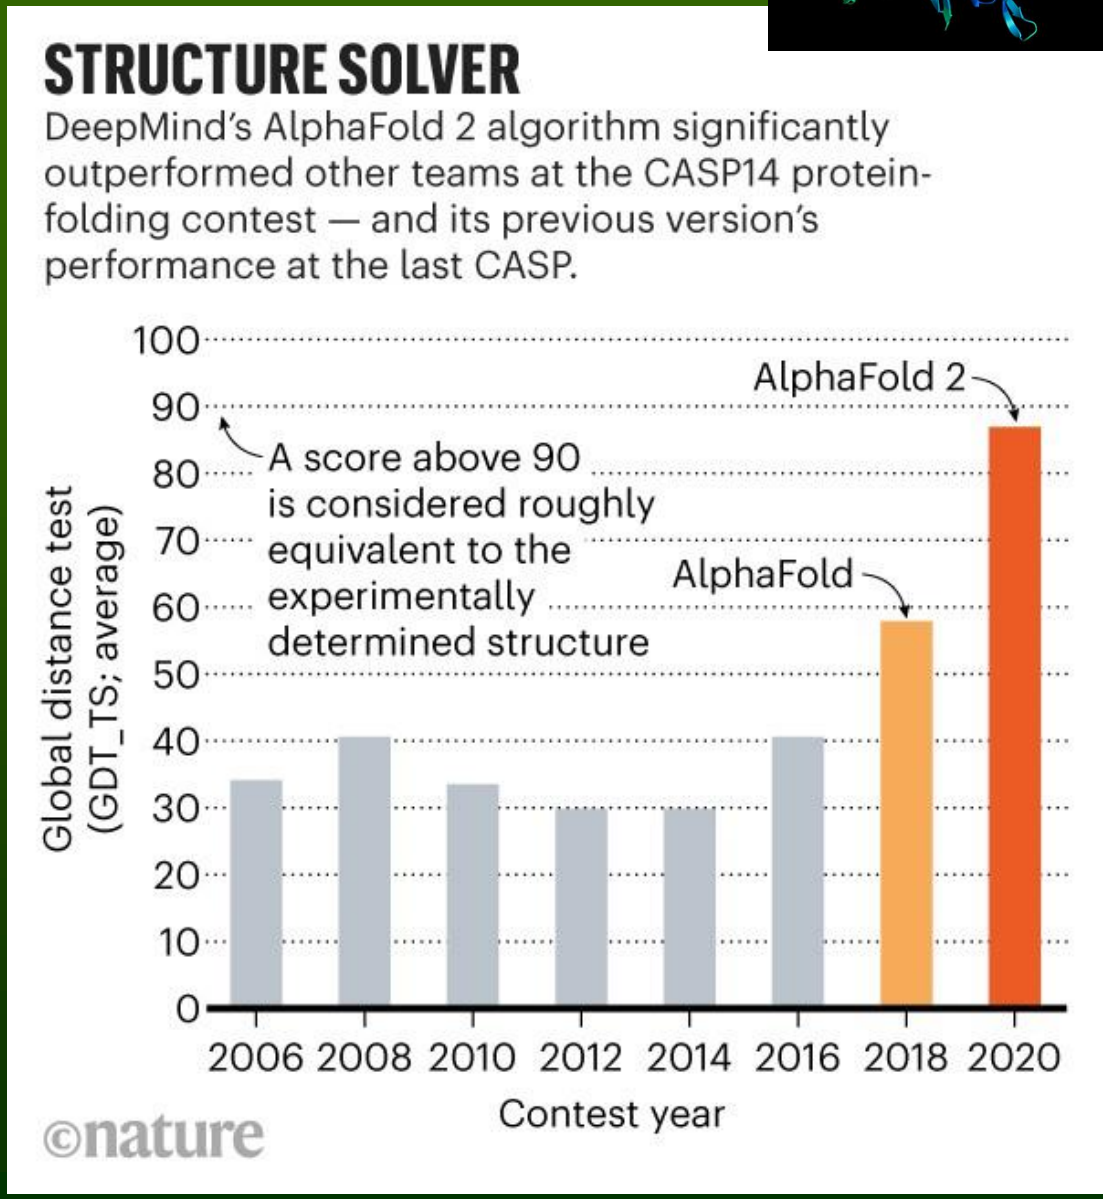

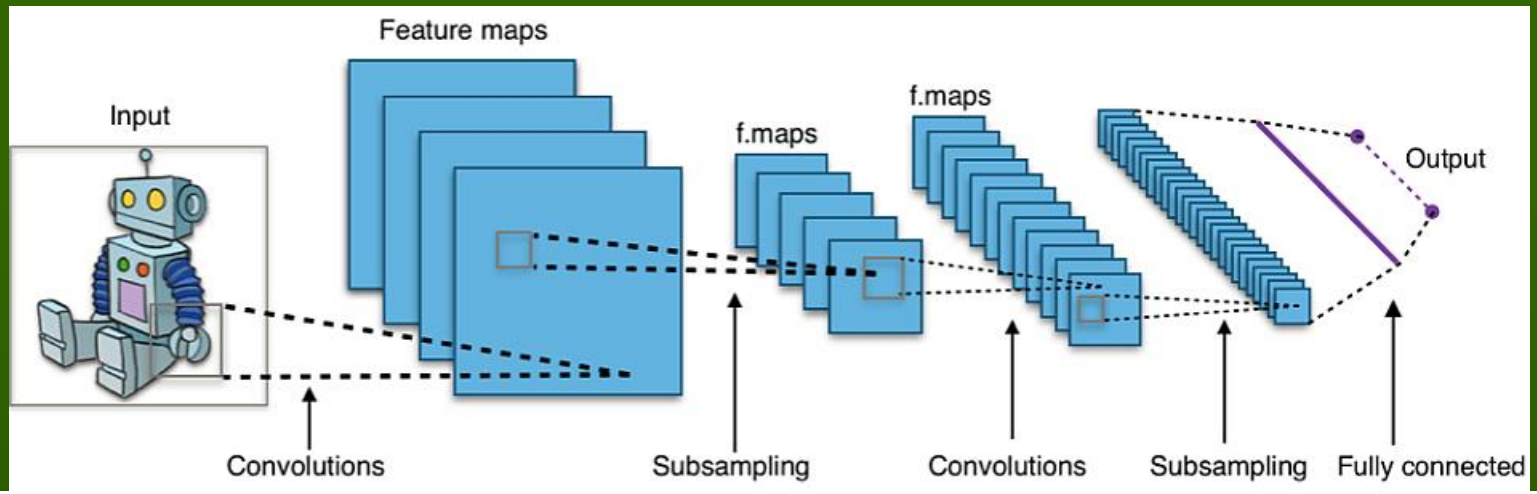

Feature representation

Input data

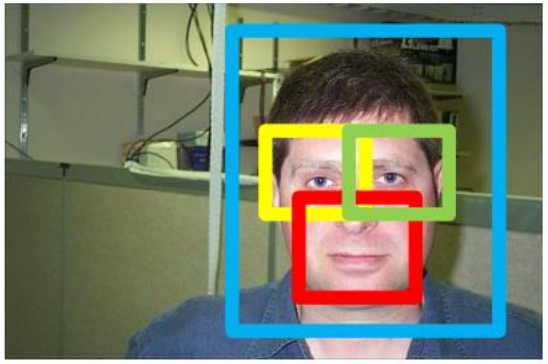

3rd layer
"Objects"

2nd layer
"Object parts"

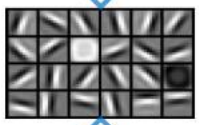

1st layer
"Edges"

Pixels

Lee et al., ICML 2009;
CACM 2011

AI will give us …

Analysis of facial images of >1M people allowed to recognize **conservative vs liberal** orientation in 72%; human judges 55% (M. Kosiński, Sci. Rep. 2021).

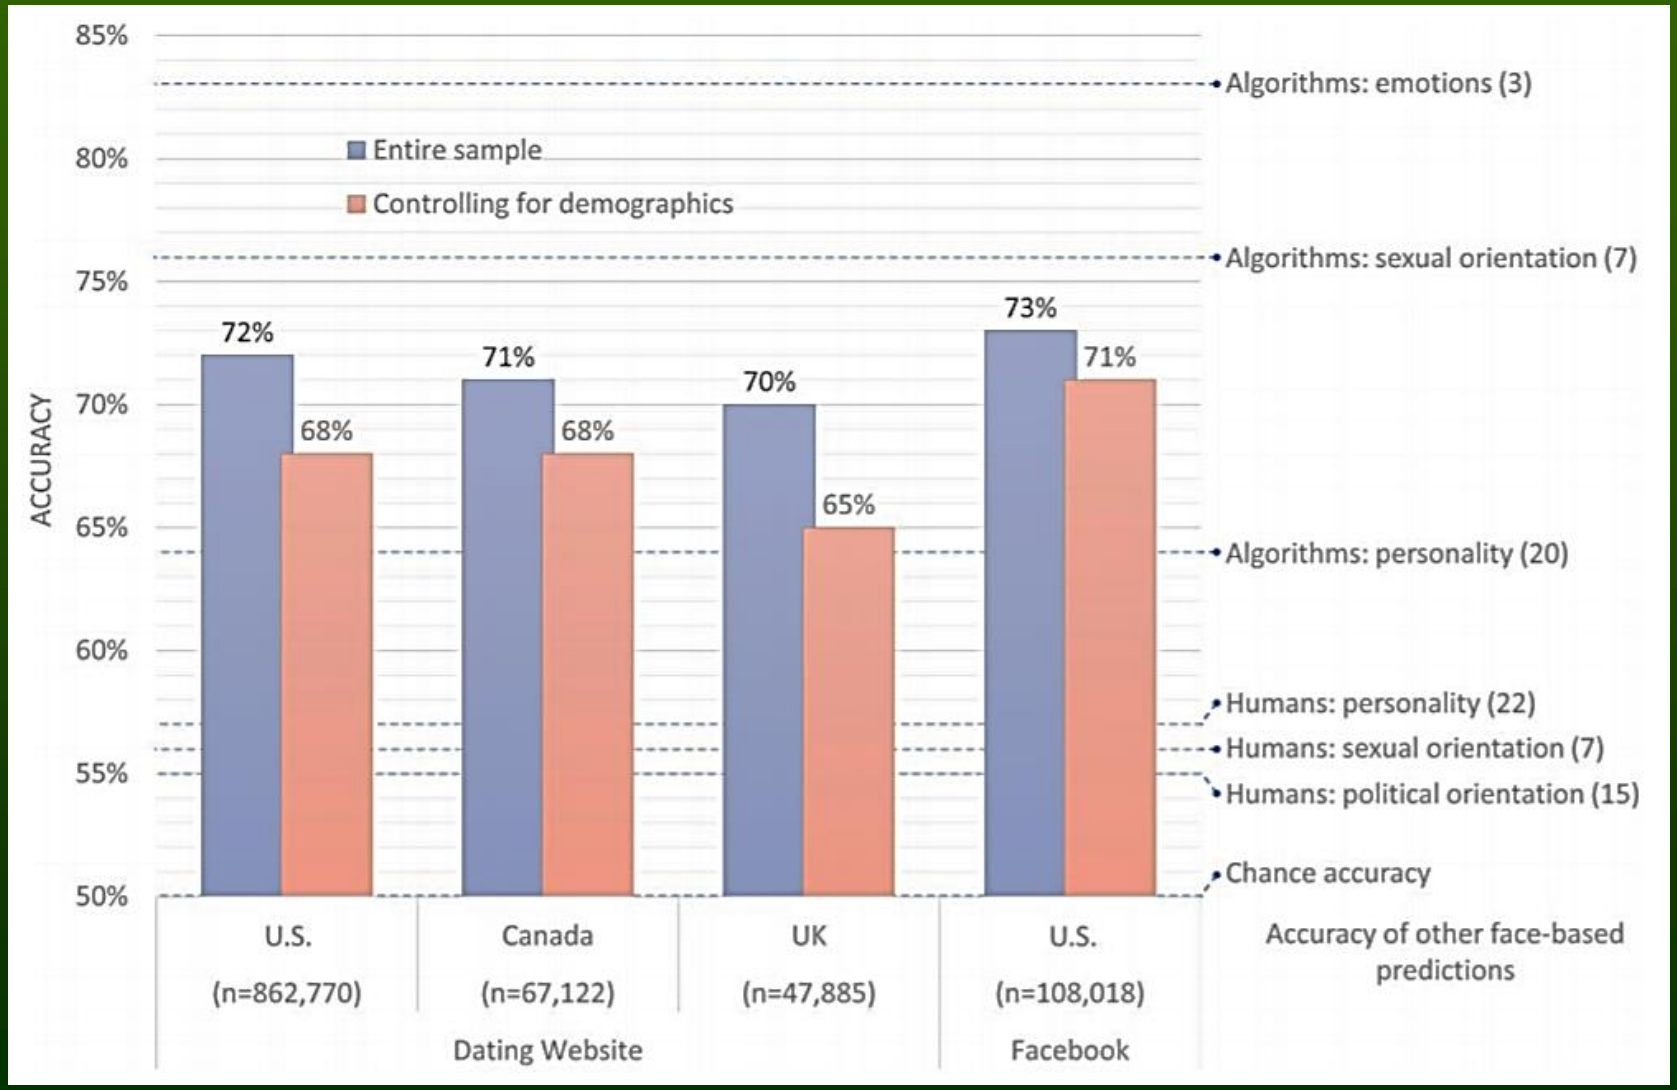

Science in the new era

1st paradigm	2nd paradigm	3rd paradigm	4th paradigm	
Empirical science	**Theoretical science**	**Computational science**	**Big data-driven science**	**Accelerated discovery**
Observations Experimentation	Scientific laws Physics Biology Chemistry	Simulations Molecular dynamics Mechanistic models	Big data Machine learning Patterns Anomalies Visualzation	Scientific knowledge at scale AI-generated hypotheses Autonomous testing
Pre-Renaissance	**~1600s**	**~1950**	**~2000**	**~2020**

Increasing speed, automation, and scale

IBM Science and Technology Outlook 2021.

Increasingly complex data models: CyC, IBM Watson, GPT-3, Google Mixture of Experts (MoE), WuDao, models with more than trillion parameters ...

Control: robots

Behavioral intelligence: training a robot from “infancy”.

Cog Project, MIT Brooks lab, 1994-2003. iCube (EU). Now we have **Atlas**.

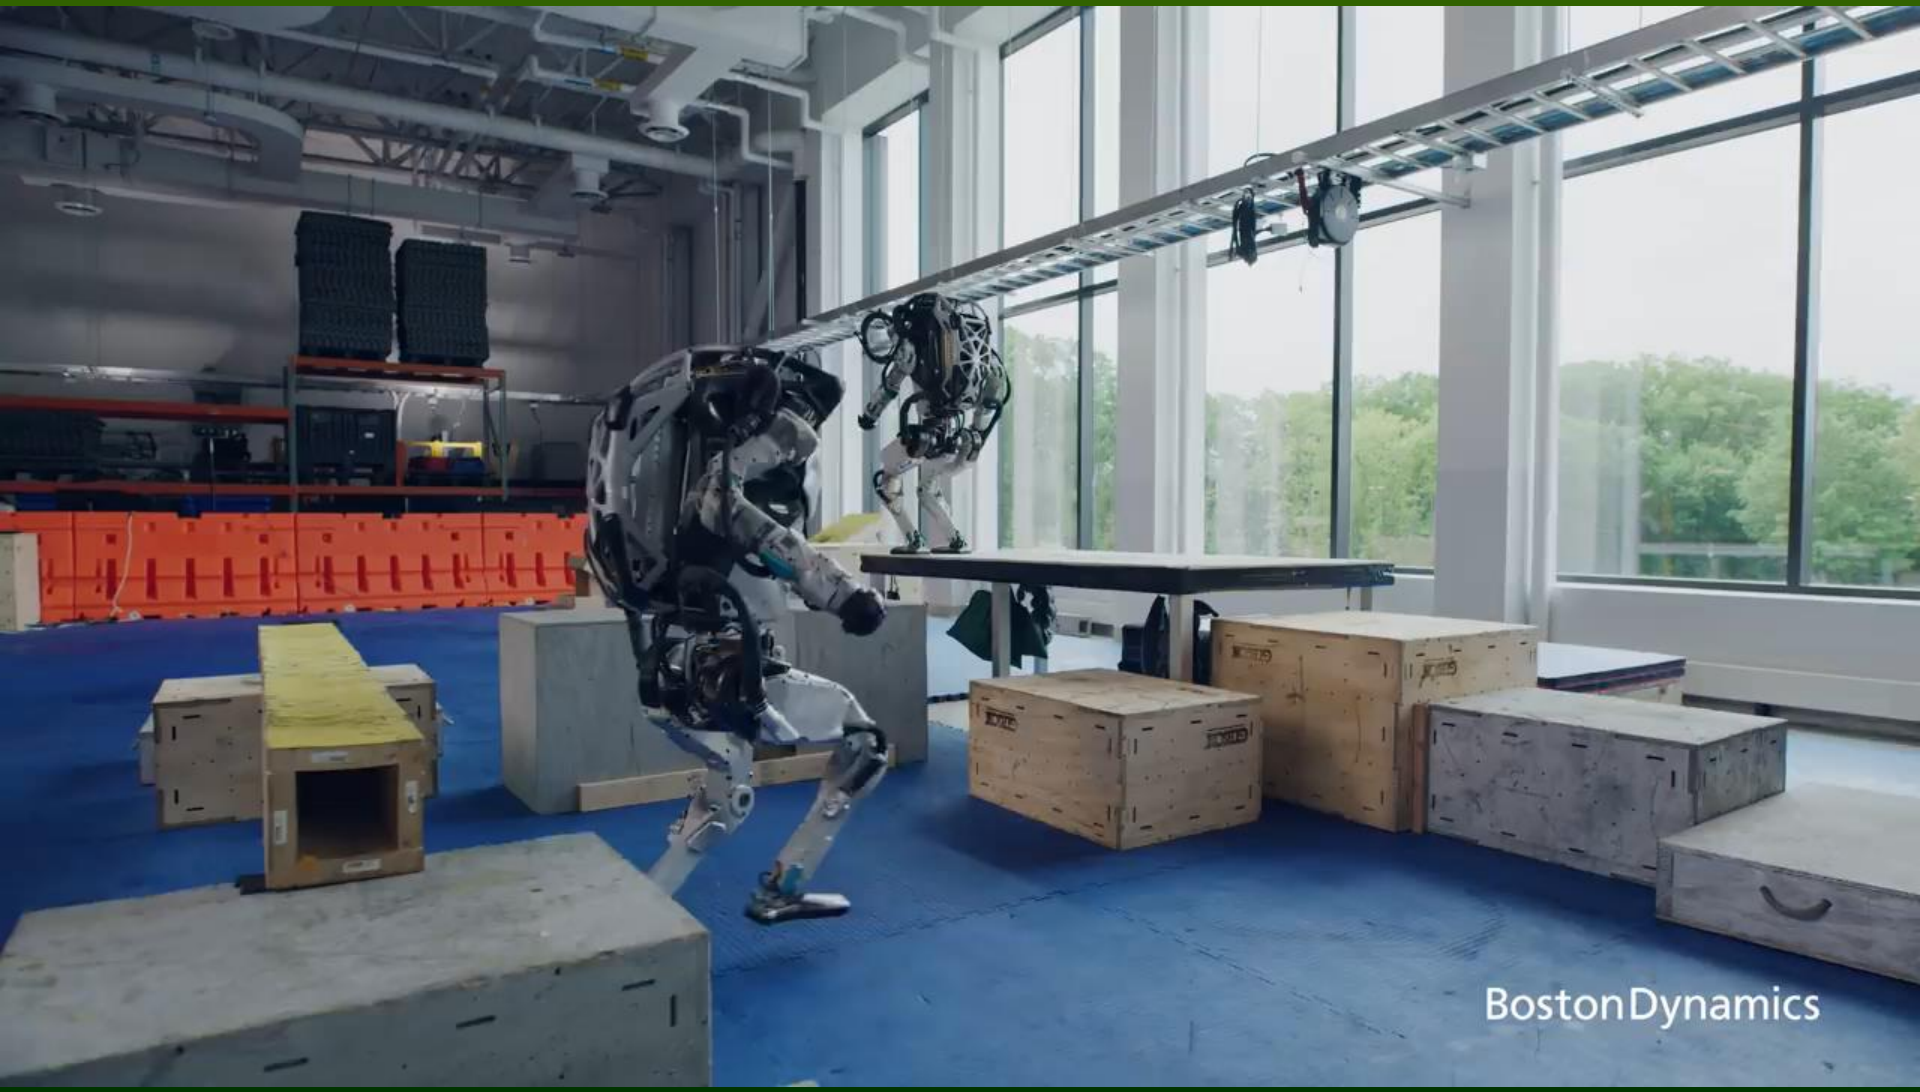

Goal: AGI & BICA

From an engineer's perspective, to understand the brain is to build a working model that exhibits the same functions. Needed: spatial models of phenomena, actions and their causes, real world imagery.

AGI = Artificial General Intelligence, learn many different things.

BICA (Brain-Inspired Cognitive Architecture) brain-like intelligence.

Duch, Oentaryo, Pasquier, Cognitive architectures: where do we go from here?

"We'll never have true AI without first understanding the brain"

Jeff Hawkins (2020).

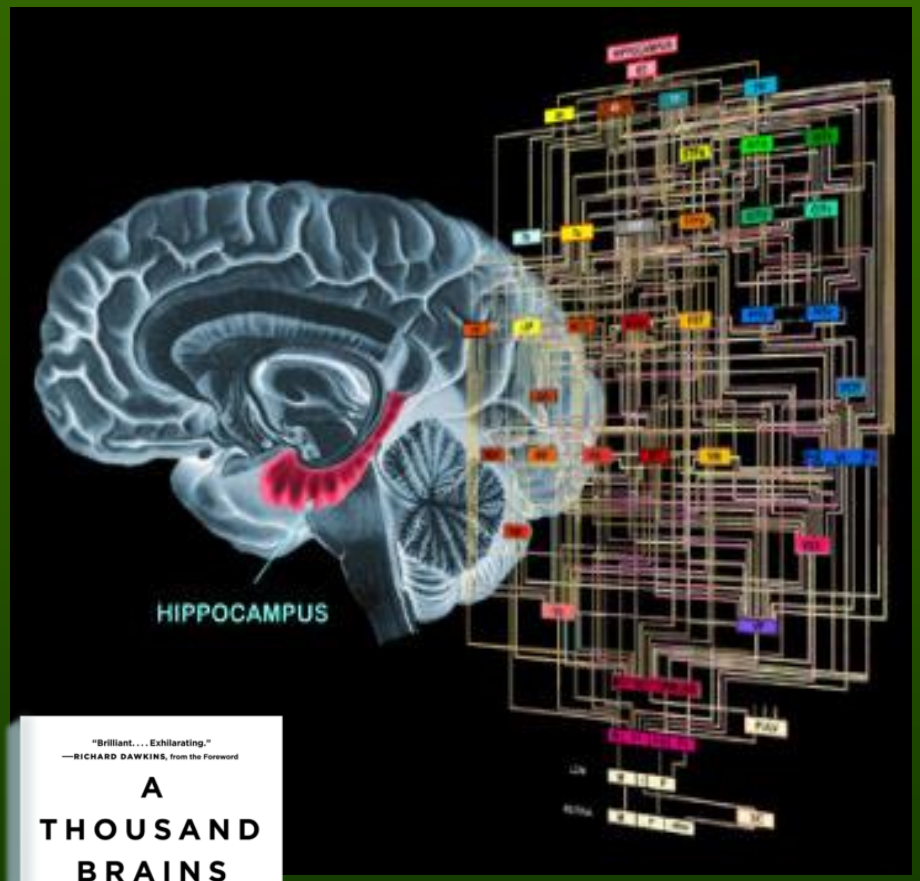

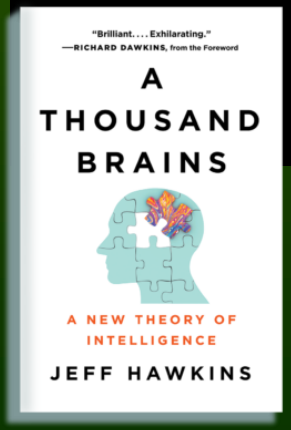

Artificial General Intelligence (AGI), Memphis 2008

First step to AGI: 2022, DeepMind Gato, a relatively small model, 1.2 B parameters. Multi-modal, multi-task, multi-embodiment, learned simultaneously over 600 tasks, from games to robot tasks.

AGI - when?

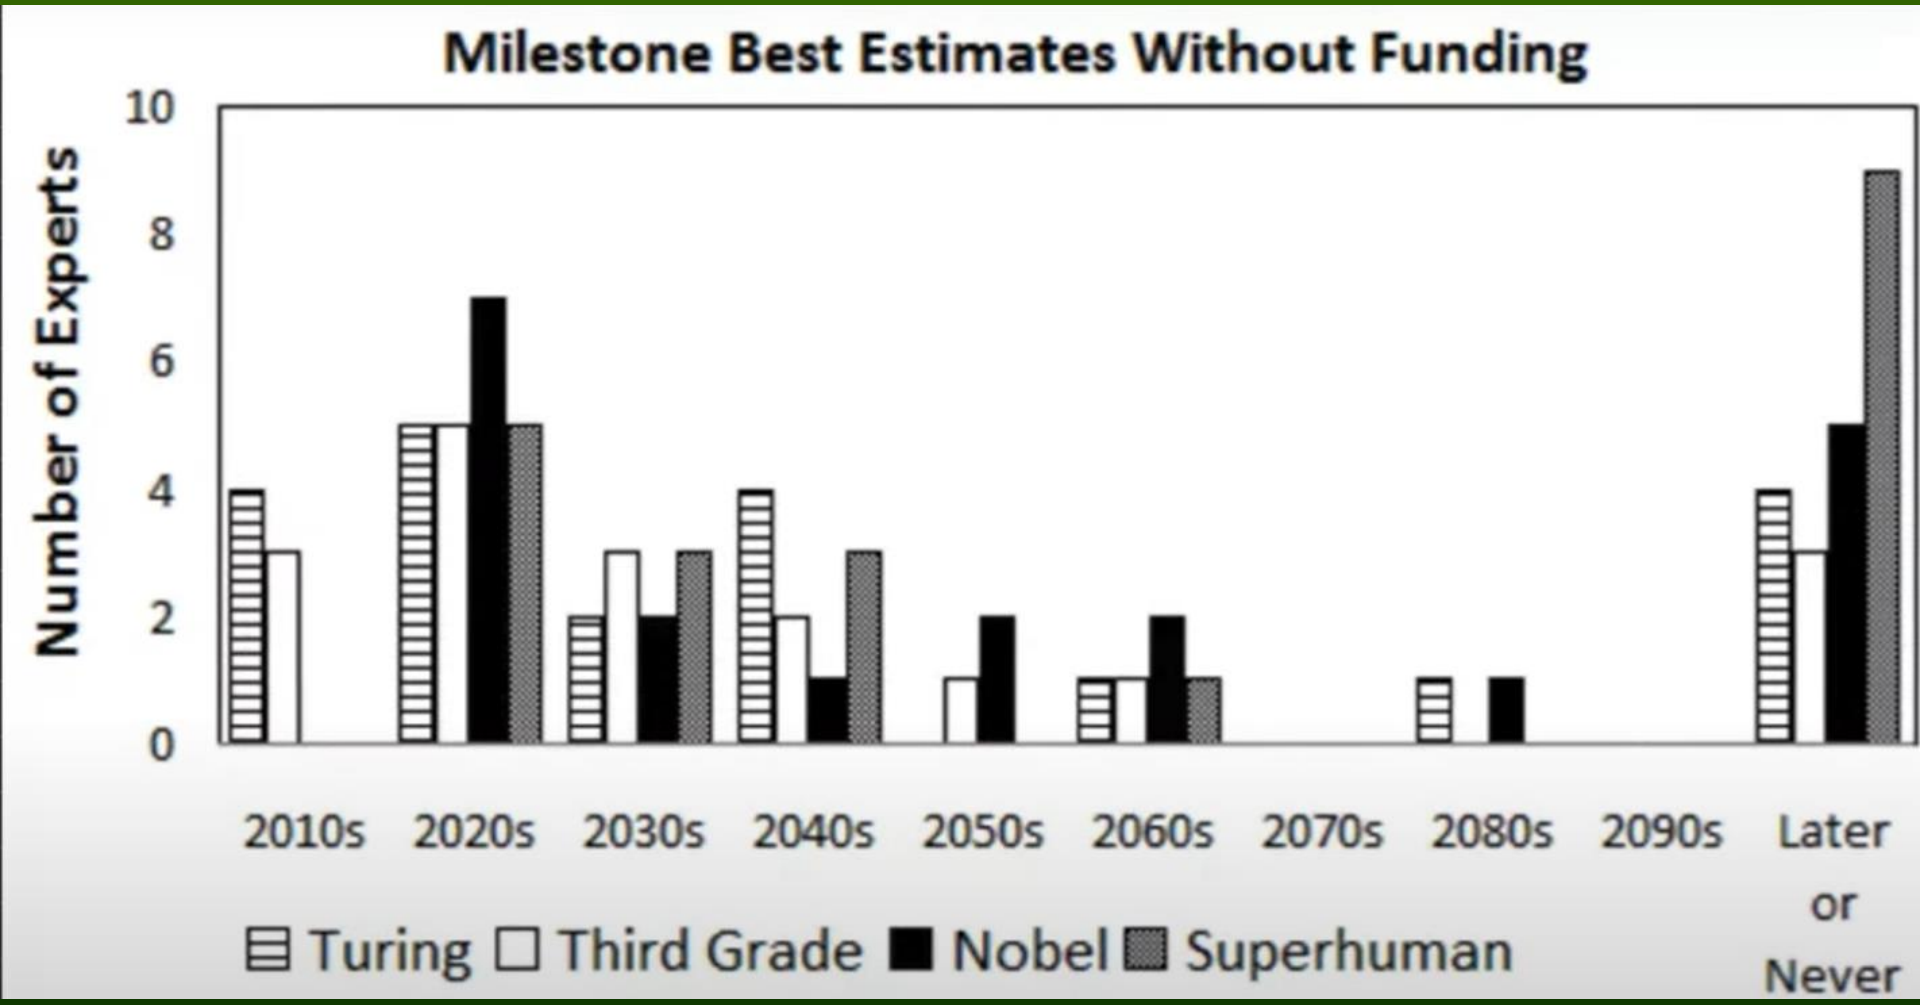

In 2009 many people thought that superhuman milestone will never be reached.

AGI-09 Survey – AI Impacts

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).

Sophia and Mika: CEO of Dictador

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

Neural classifiers

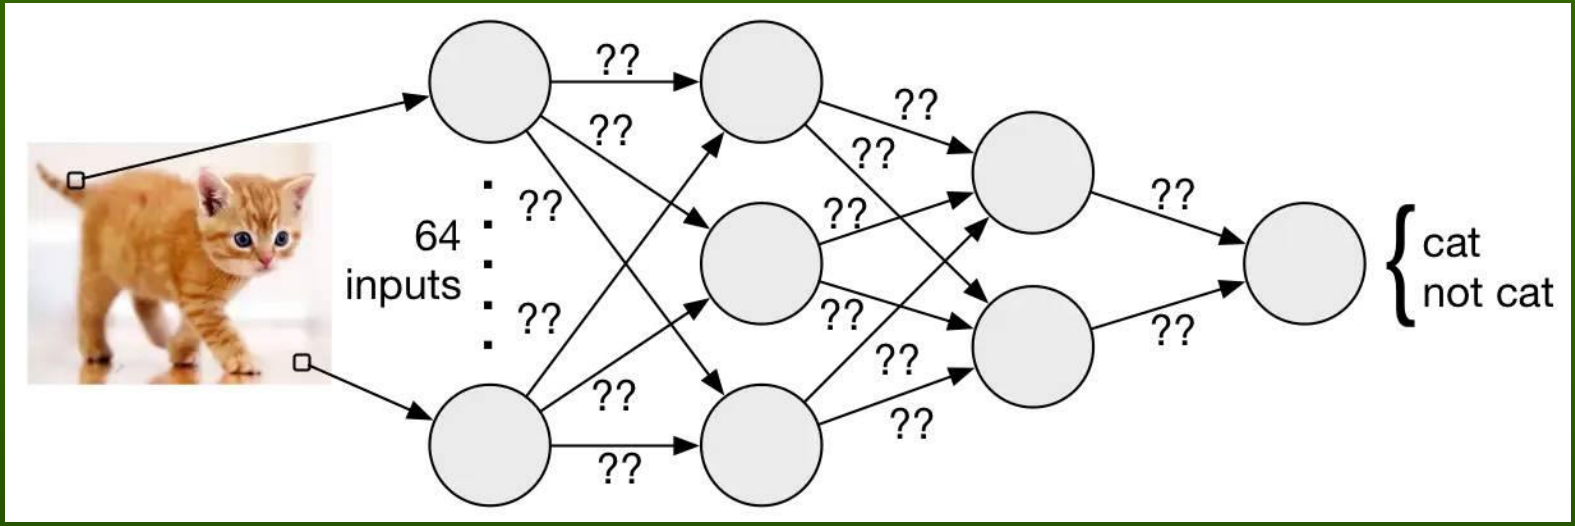

Words, image patches => networks with adjustable parameters
=> training to recognize patterns => object classification, diagnosis.

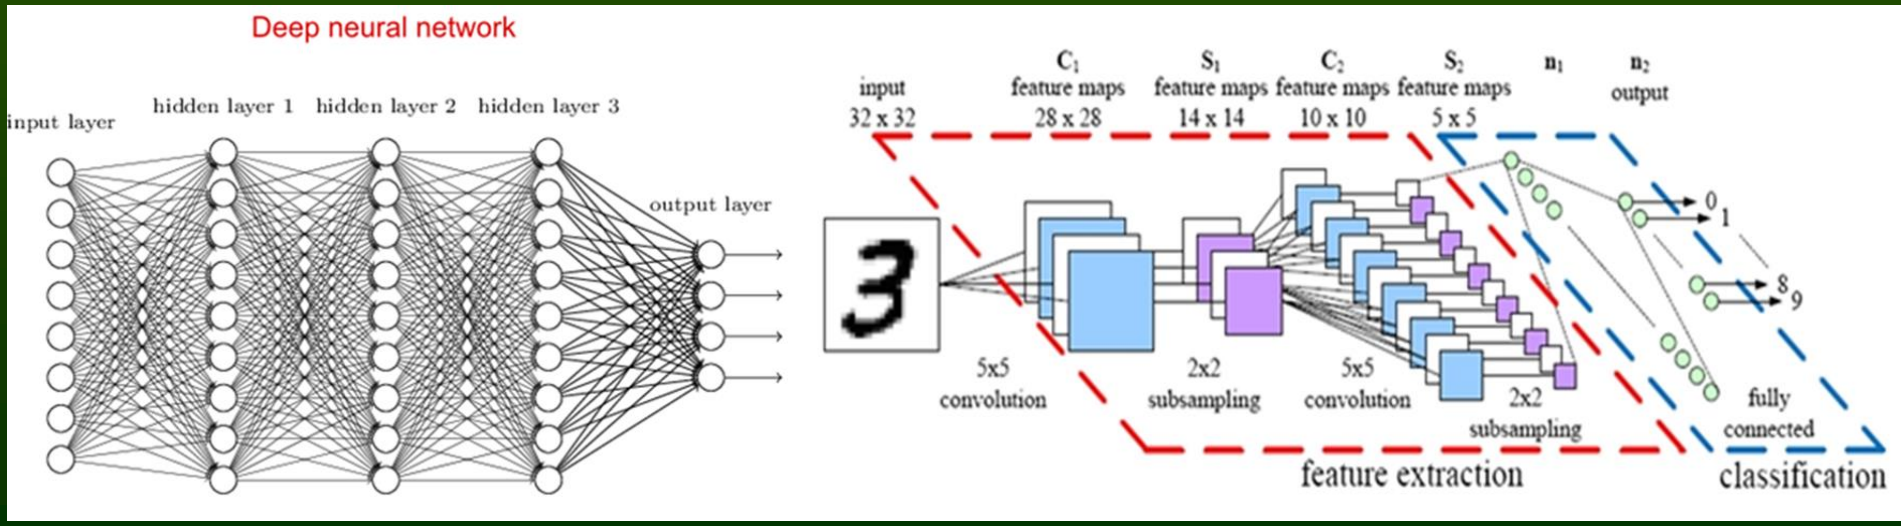

Third wave of AI

Layering neural networks

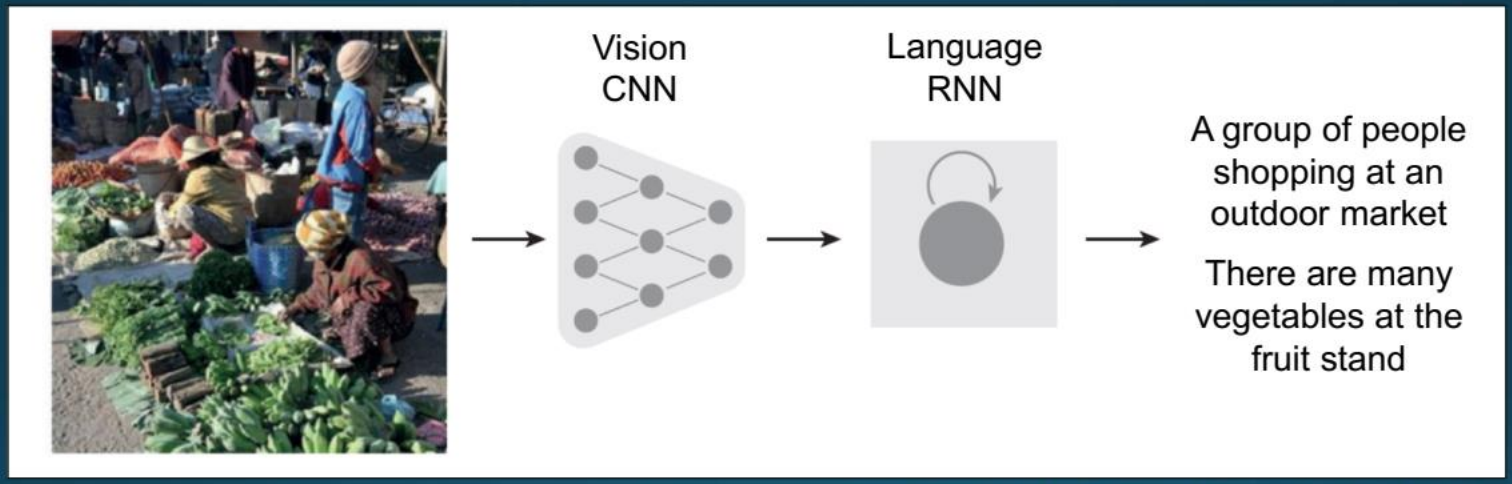

Yann LeCun, Yoshua Bengio, & Geoffrey Hinton (2015). Deep Learning, Nature, Vol. 521, (pp. 436-444)

A deep convolution neural net (CNN) produces a set of outputs (abstract "words")

A language-generating recurrent neural net (RNN) "translates" the abstract "words" into captions

Invariant image analysis and interpretation.

Third wave of AI

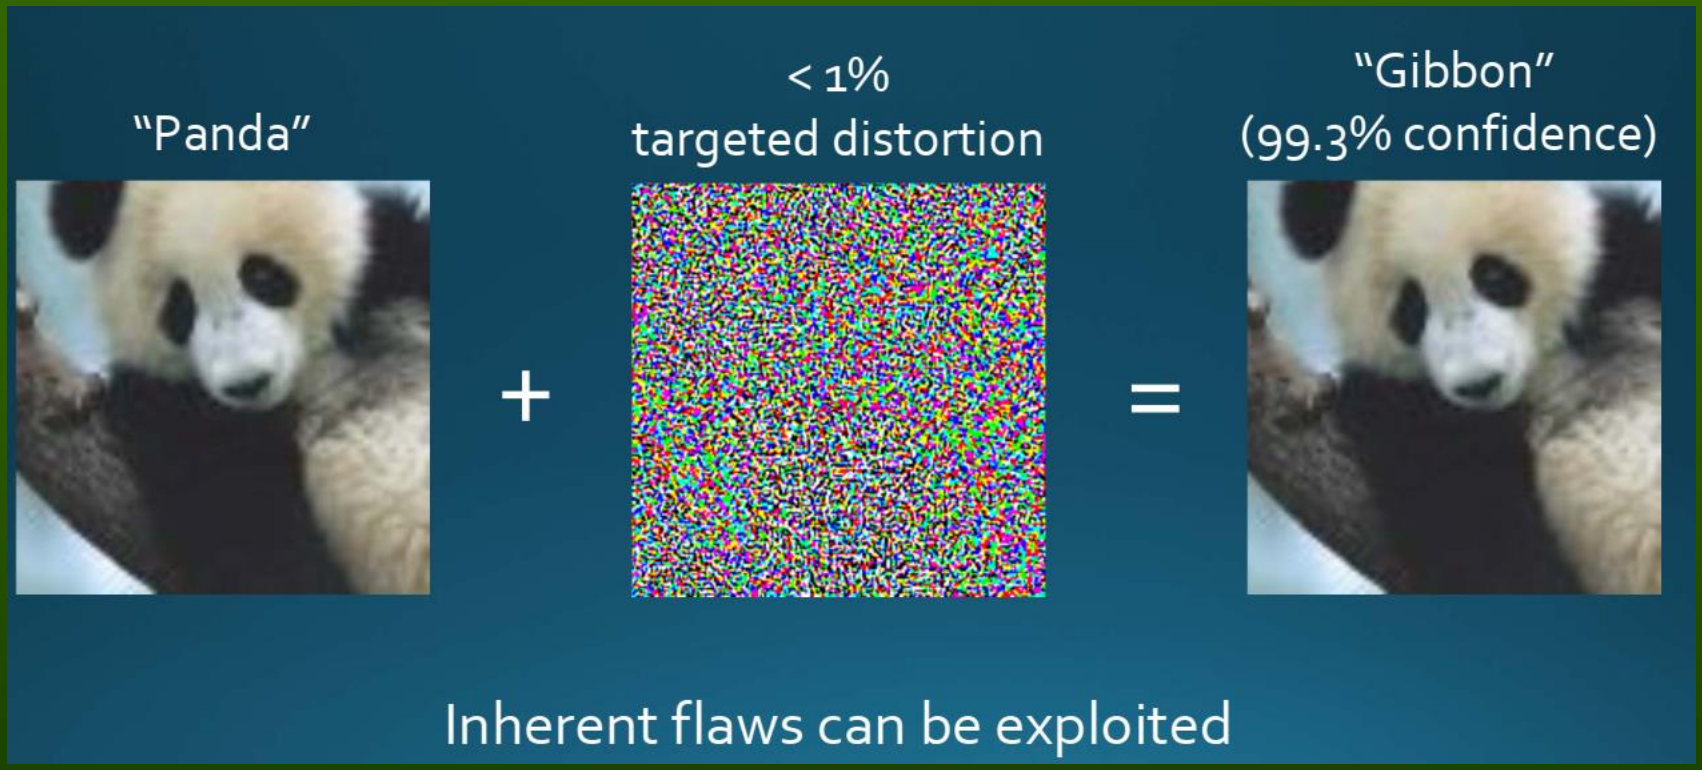

Easy to distort, not clear what has been learned.
With wrong data biased solution, ex: famous bot that learned from fascist groups, or mistaking some black people for gorillas …

Need for contextual models of real world phenomena – developed in 2022/23.

Third AI wave and brains

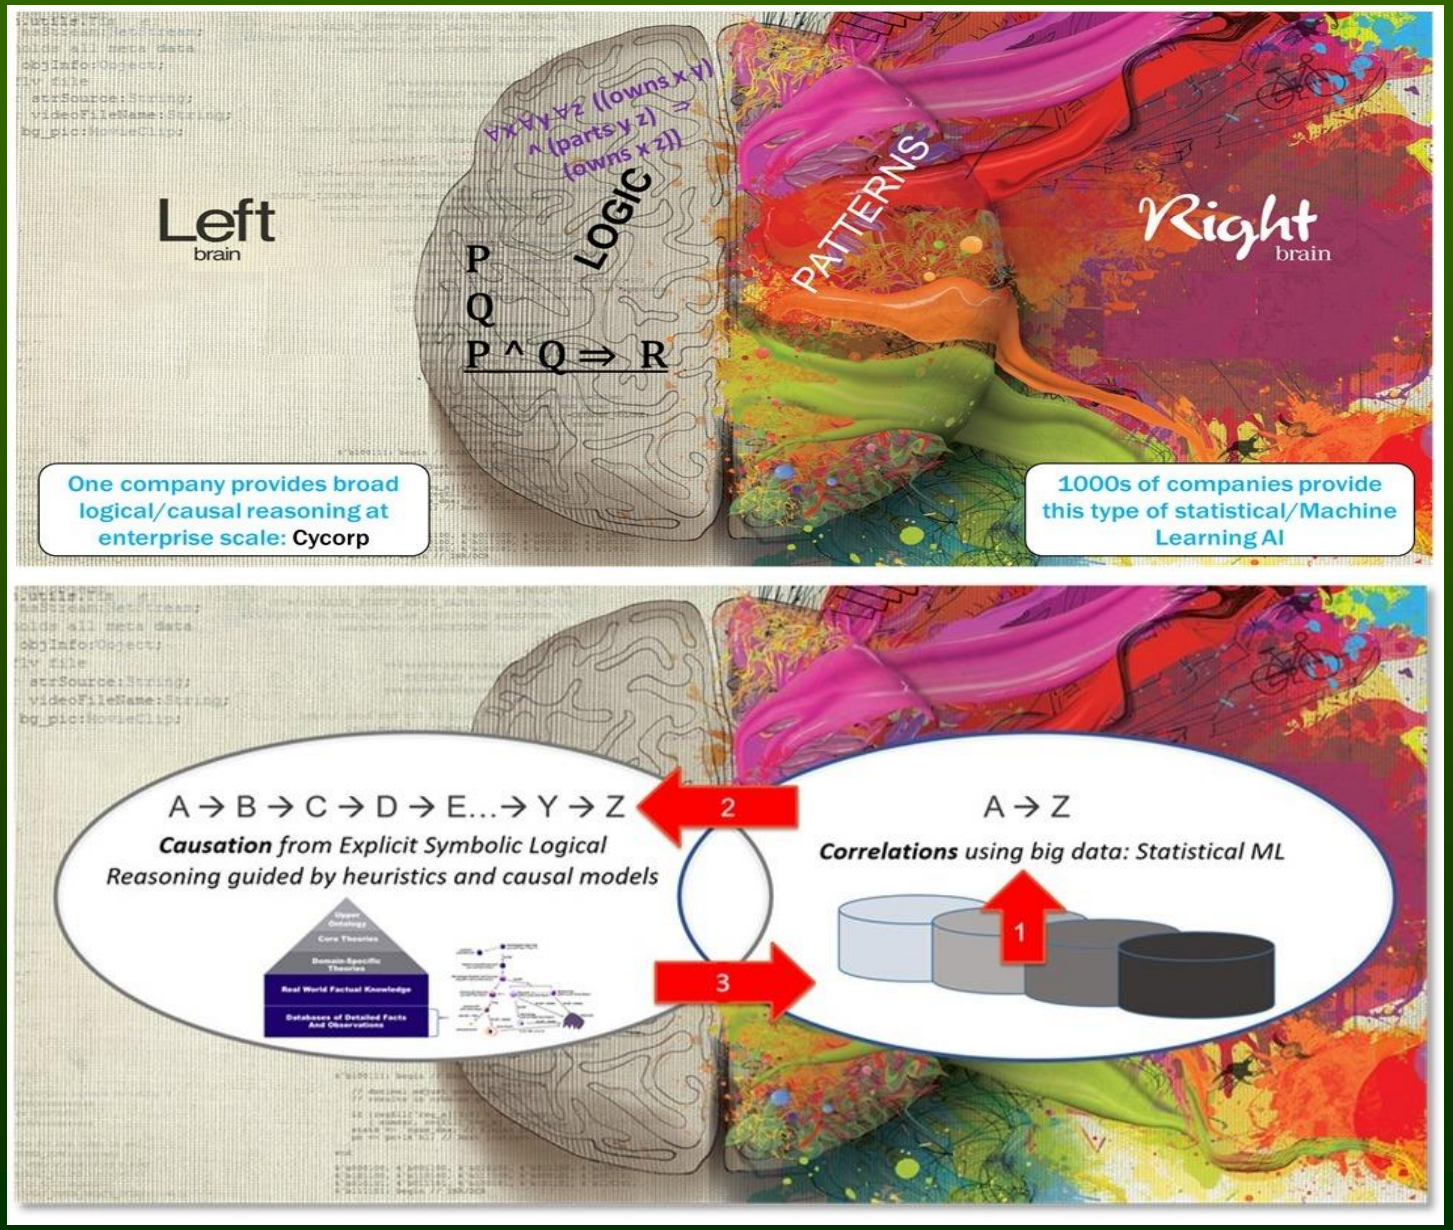

Third wave of AI

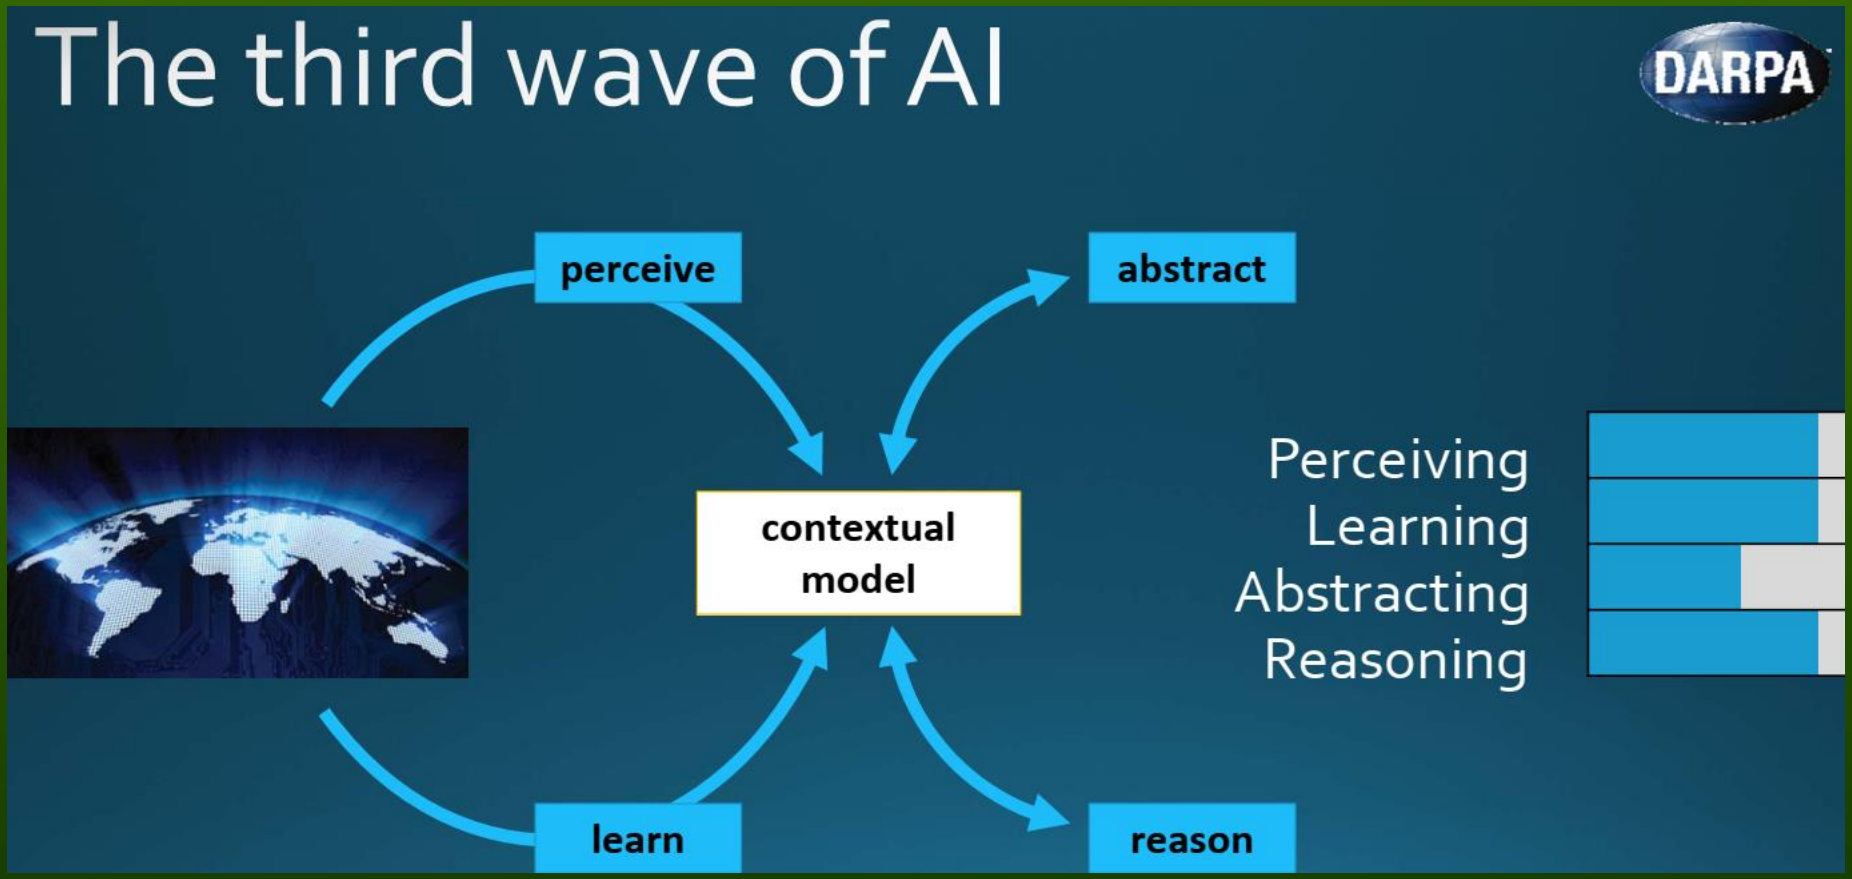

GAN, Generative Adversarial Networks, one network creates false examples distorting learning data, another network learns to distinguish them from natural ones. Building models of objects and situations is the next step.

Yesterday: transformers

LLM timeline

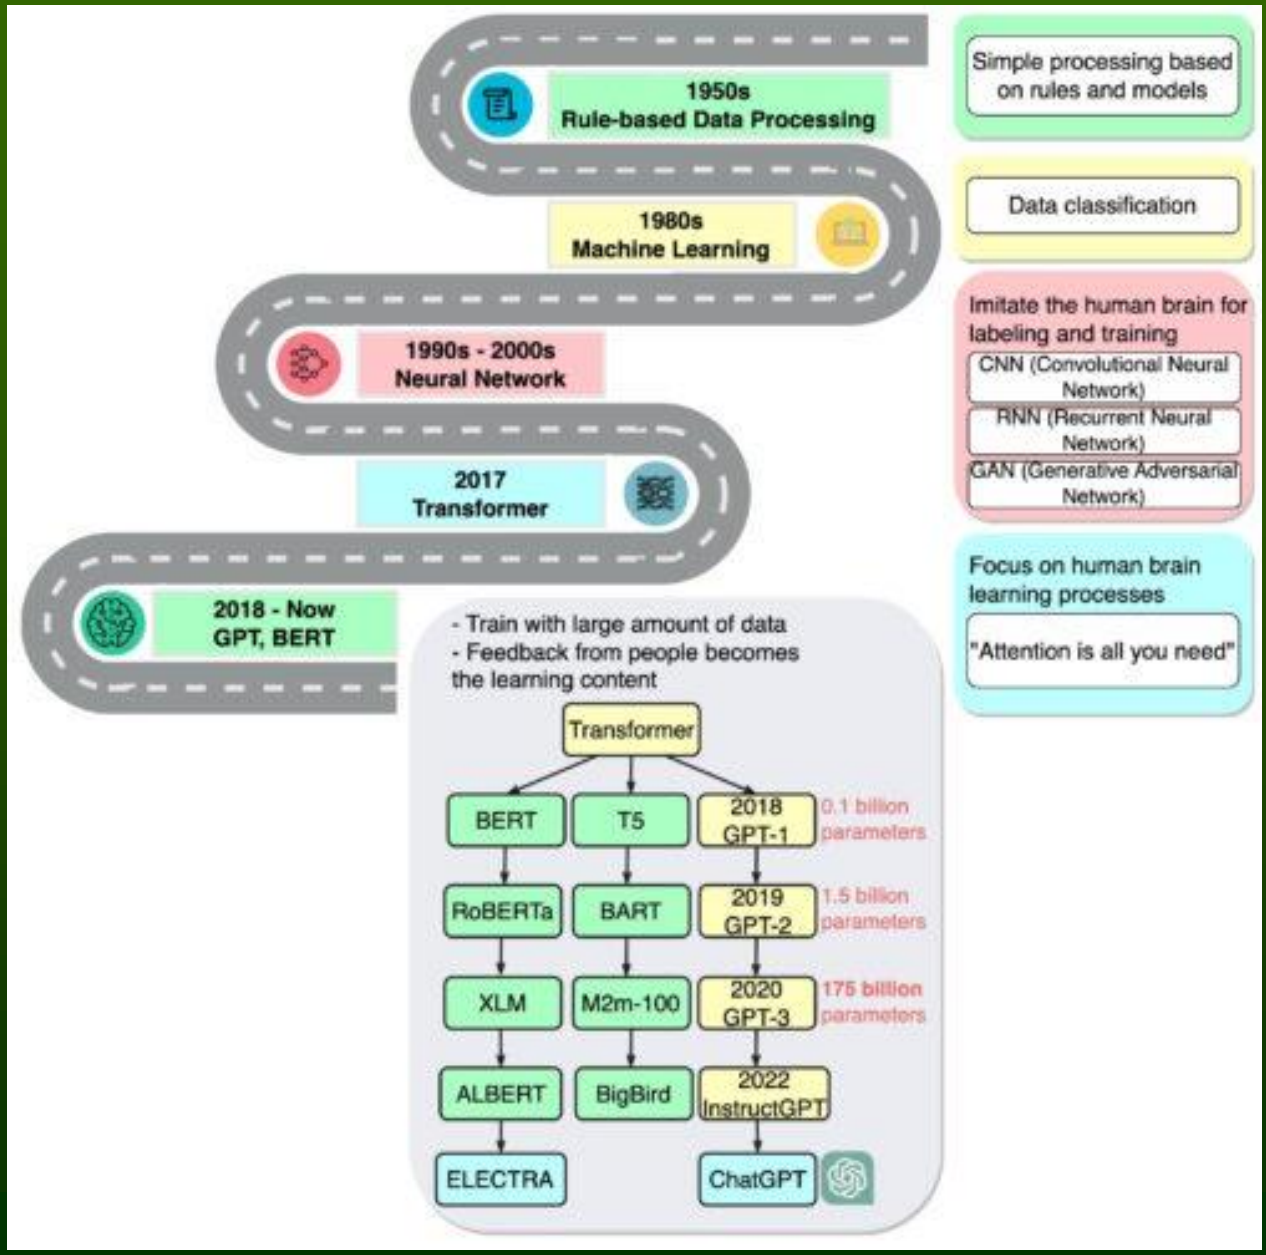

Language algorithms

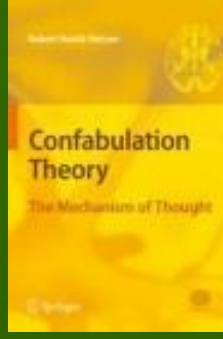

Language models: relation of words in complex network structures.
In 2018, to gain a general-purpose “language understanding”,
Google created BERT, model pre-trained on a very large text corpus.

- **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.

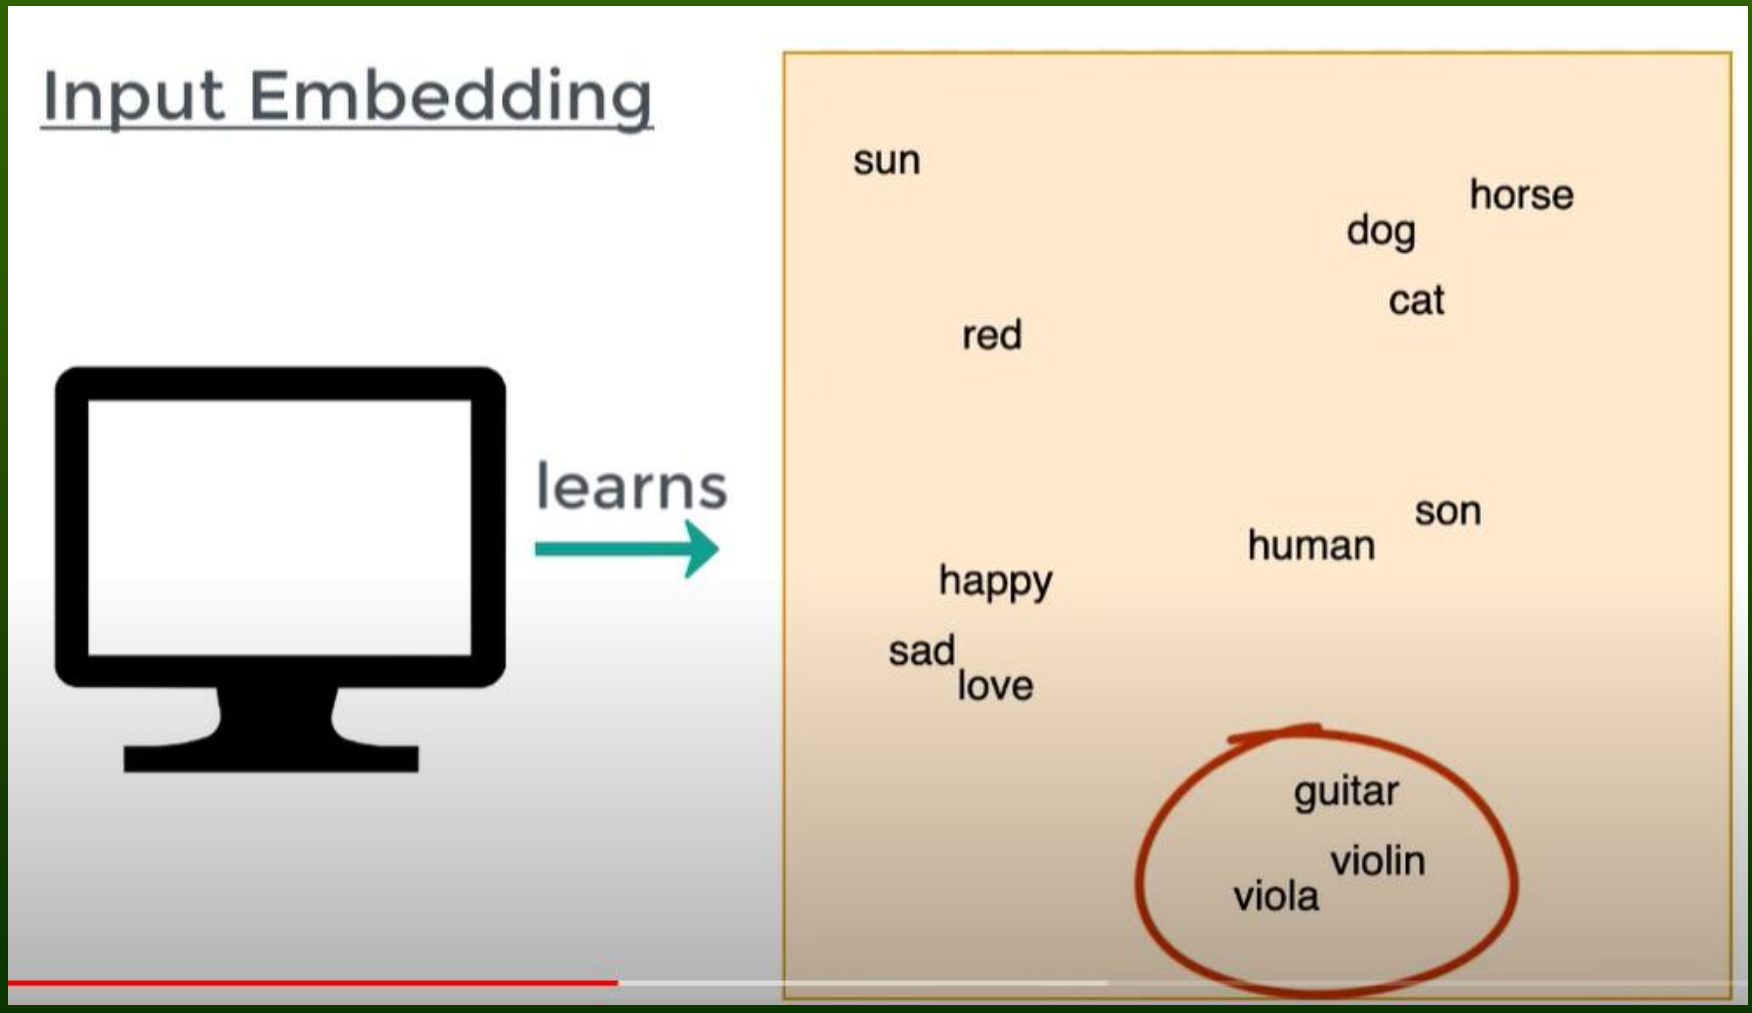

Transformers

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (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 works in a sentence. For example:

Input: English sentence;
Output: Polish sentence.

Google BERT has used this.

Simple intro on Youtube.

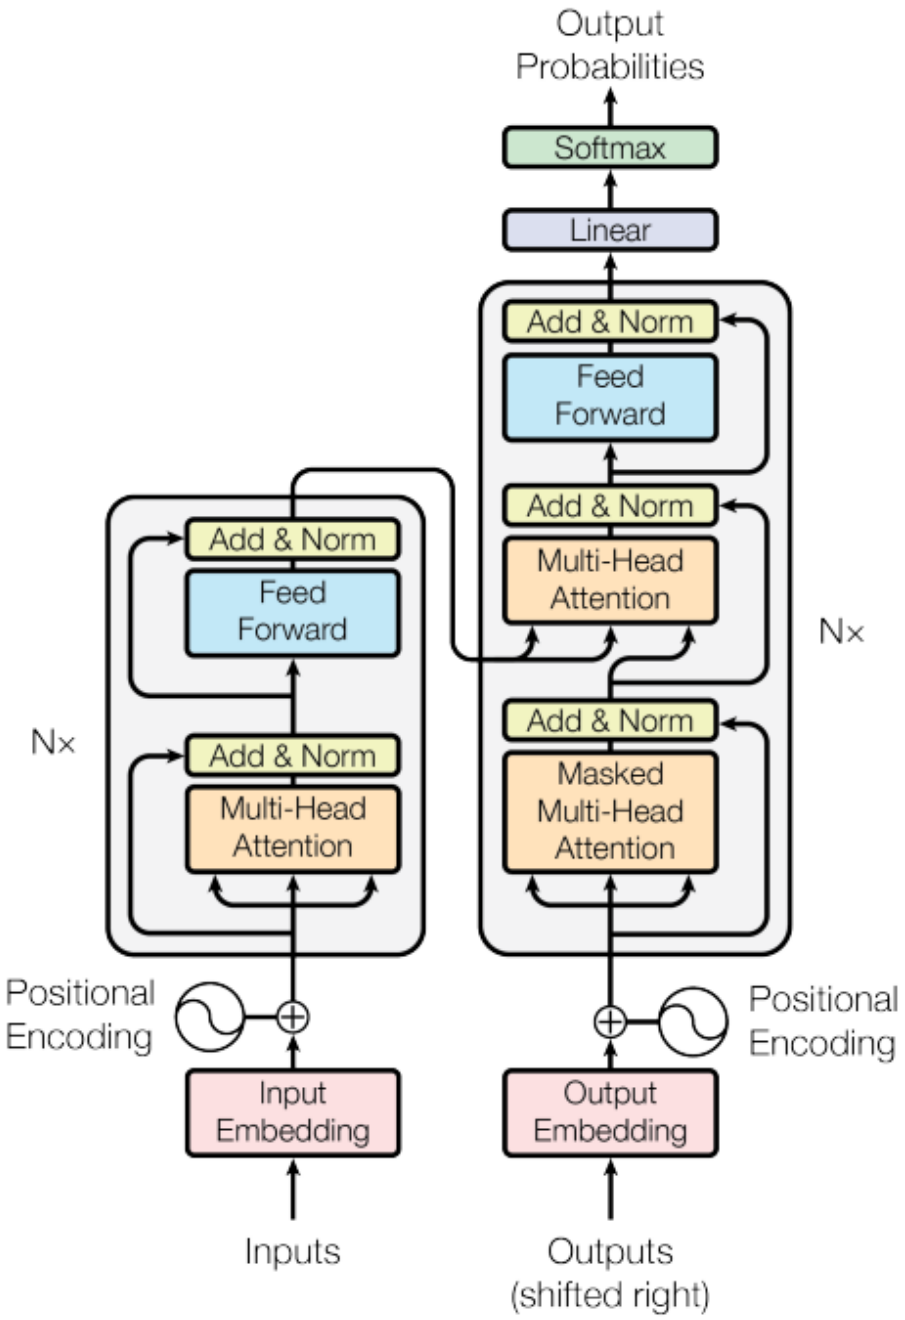

Figure 1: The Transformer - model architecture.

Q/A state of the art

Results for 100,000 questions from the database Stanford Question Answering Dataset (SquAD) are better than the results achieved by humans.

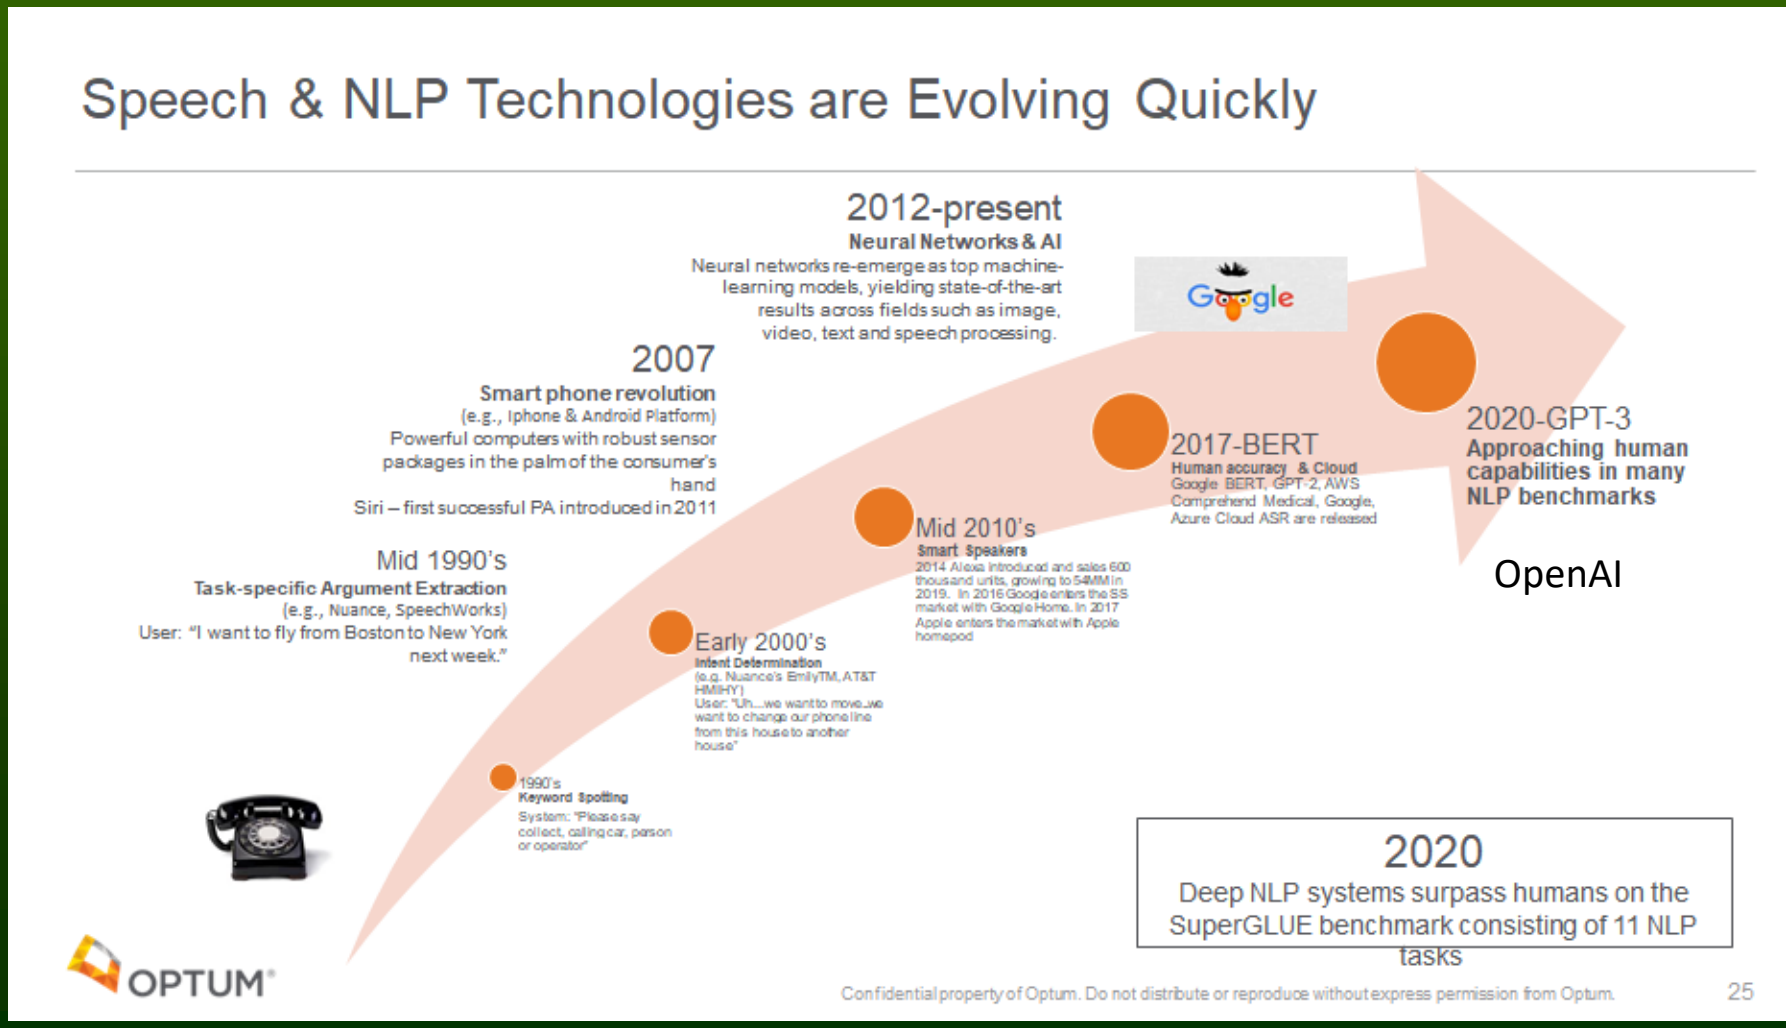

NLP supermodels

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.

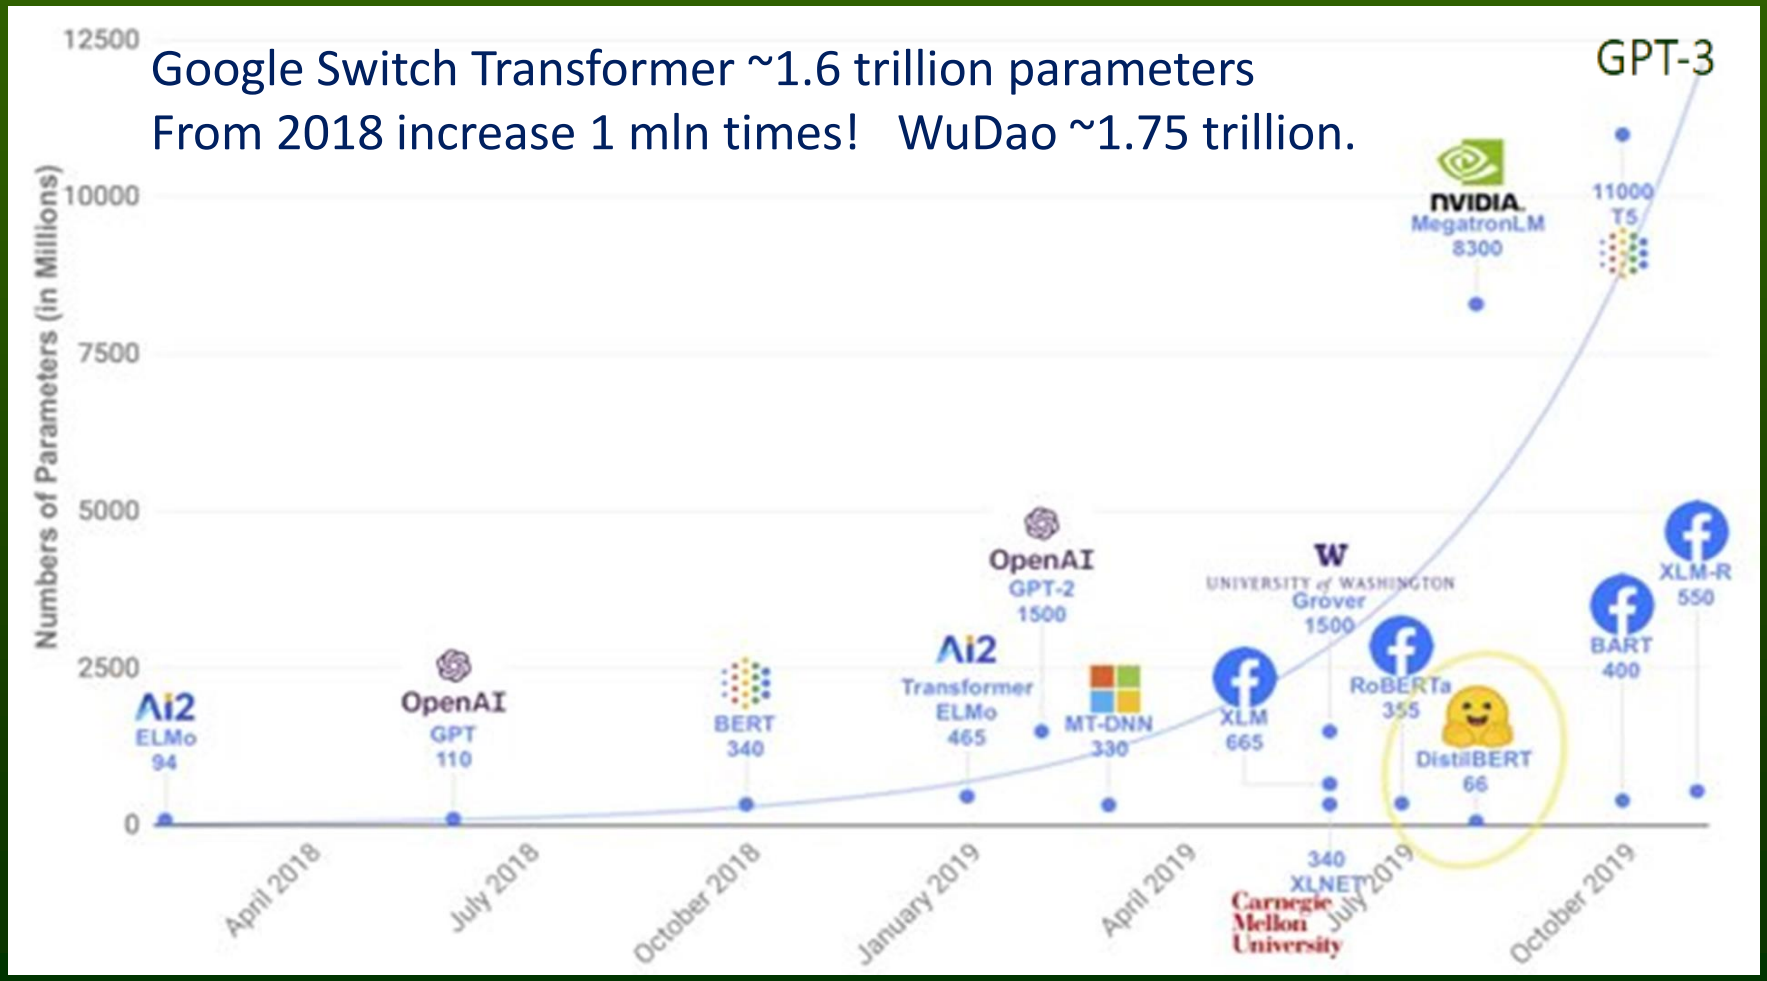

Large Language Models

TEXT GENERATIVE AND CONVERSATIONAL AI LANDSCAPE*

Companies with ChatGPT-like Functions

Text Analysis

Text Summarization

scholarcy, TLDR this, agolo, IntelliPPT, genei, spaCy, yabble, bundleIQ, upword*, Summarizer, QuillBot, IRIS.AI

Sentiment Analysis

Repustate data in sight, MonkeyLearn, meaning cloud, Rosette TEXT ANALYTICS, RAPIDMINER, Talkwalker, Lexalytics, Dandelion API, opentext, co:here, awario, {}Komprehend

Text Translation

DeepL, SMARTLING, modernMT, textunited, Taia.io, SONIX, Phrase, Reverso, AppTek, Smartcat, TRANSPERFECT, RWS

Conversational AI

Virtual Assistants

Siri, alexa, SoundHound, Bard, Google Assistant, Claude, Socratic, fireflies.ai, Replika, augment, 24me, ELSA, youper, clara

Chatbot Building Platforms

avaamo, nlx, Openstream.ai, aivo, imperson, OneReach.ai, boost·ai, Omilia, inbenta, AMELIA, Landbot, EBI.AI

Chatbot Frameworks & NLP Engines

amazon LEX, LUIS, botpress, JOVO, COGNIGY, RASA, Microsoft Bot Framework, kore.ai, IBM Watson, yellow.ai, Dialogflow

Writers

Jasper, Writesonic, articleforge, AI-WRITER, INK, Rytr, CLOSERS COPY, Scalenut, copy.ai, frase, WordAi, wordtune, copysmith, grammarly

Coders

Codiga, tabnine, GitHub Copilot, Intellicode, Amazon CodeWhisperer, replit, Ponicode, K.Explorer, BLACKBOX, AIXcoder, Ask Codi, mutable.ai, codeium, Kodezi

Search

neeva, Perplexity, YOU, nuclia, algolia, coveo, Lucidworks, Hebbia, YeXT, raffle, Andi, vespa, KOMO, DIFFBOT

Language Models

GPT-4, LaMDA, DeepMind CHINCHILLA, T5, GPT Neo, XLNet, BLOOM, MEGATRON-TURING NLG, ELECTRA, XLM-RoBERTa, OPT-IML, Google BERT

*Source: DaybreakInsights.com, March 2023. This graphic shows only a sampling of companies from our full Text Generative and Conversational AI research.

Large Language Models

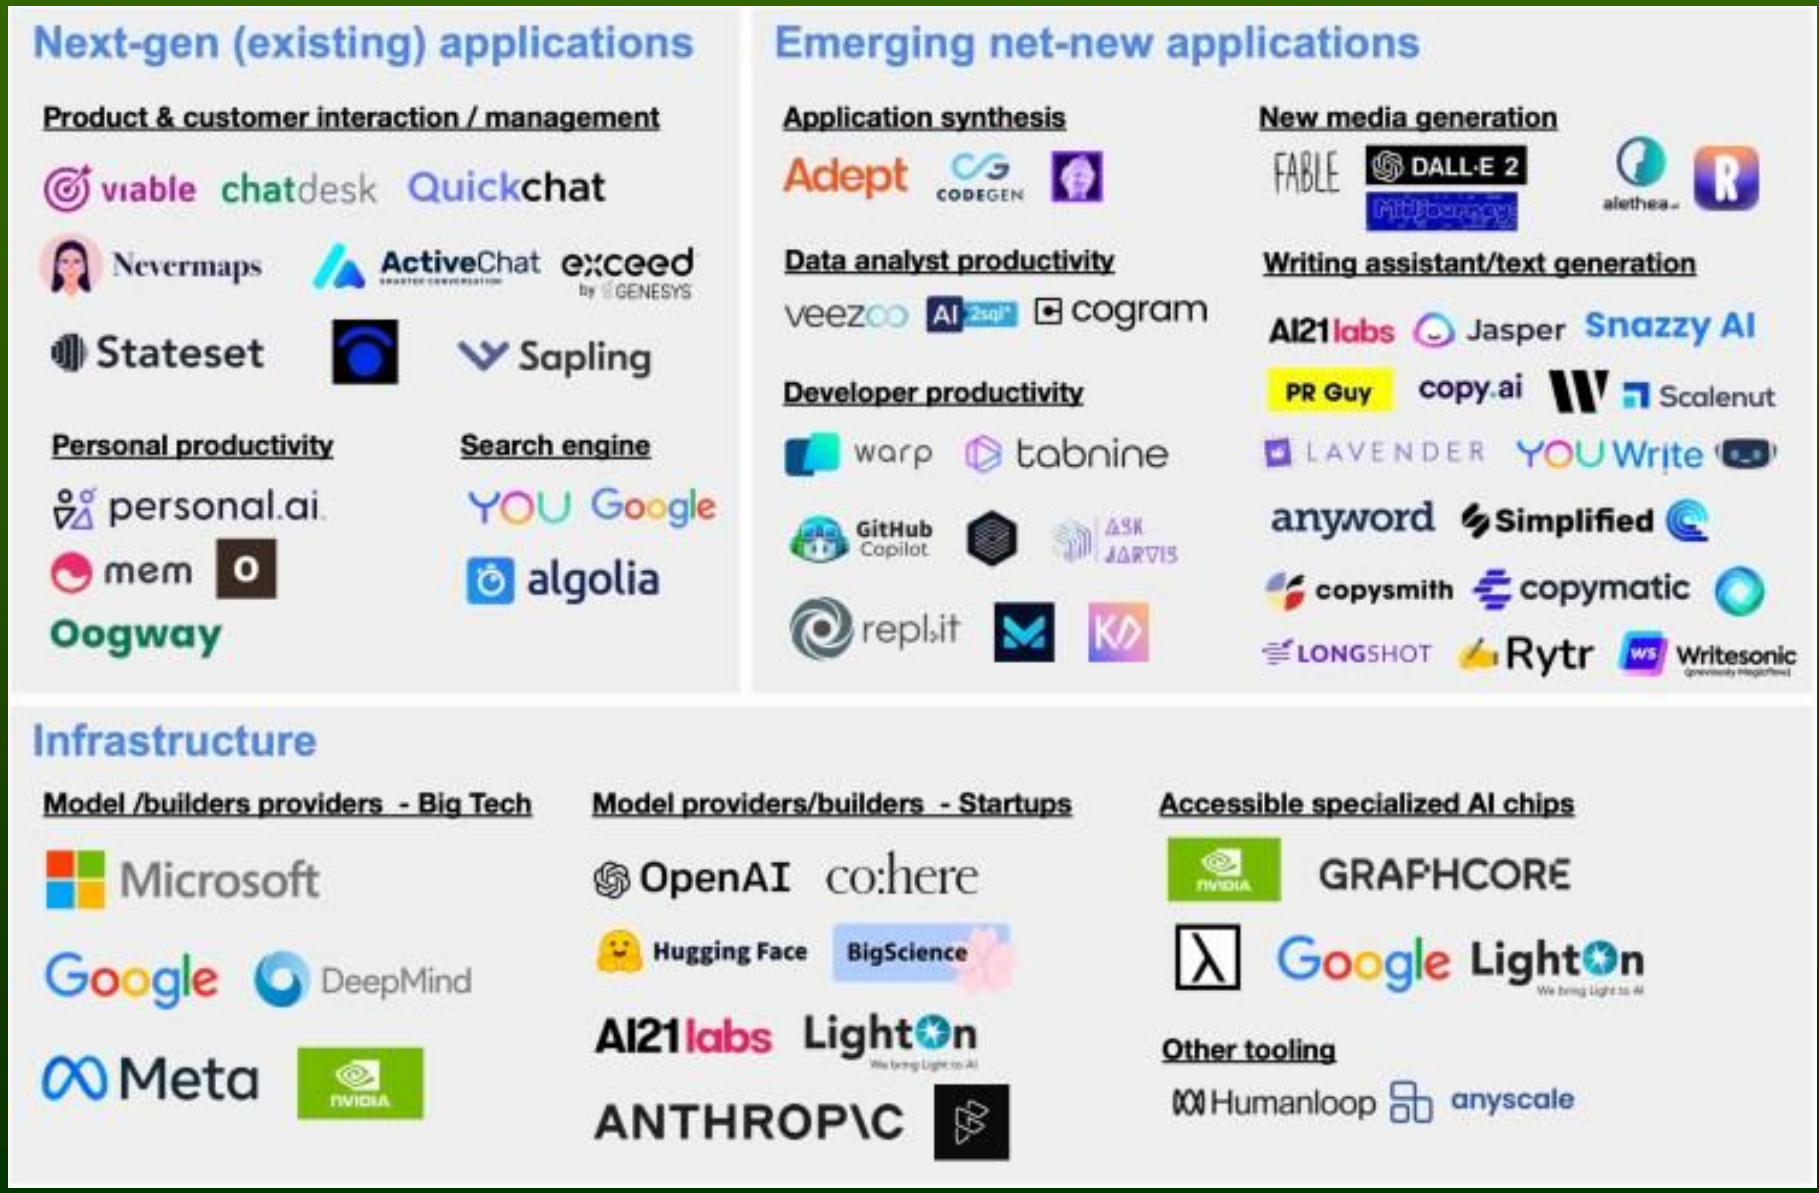

GPT-3 as philosopher

Eric Schwitzgebel, David Schwitzgebel, Anna Strasser, Creating a Large Language Model of a Philosopher, arXiv: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.

ChatPDF

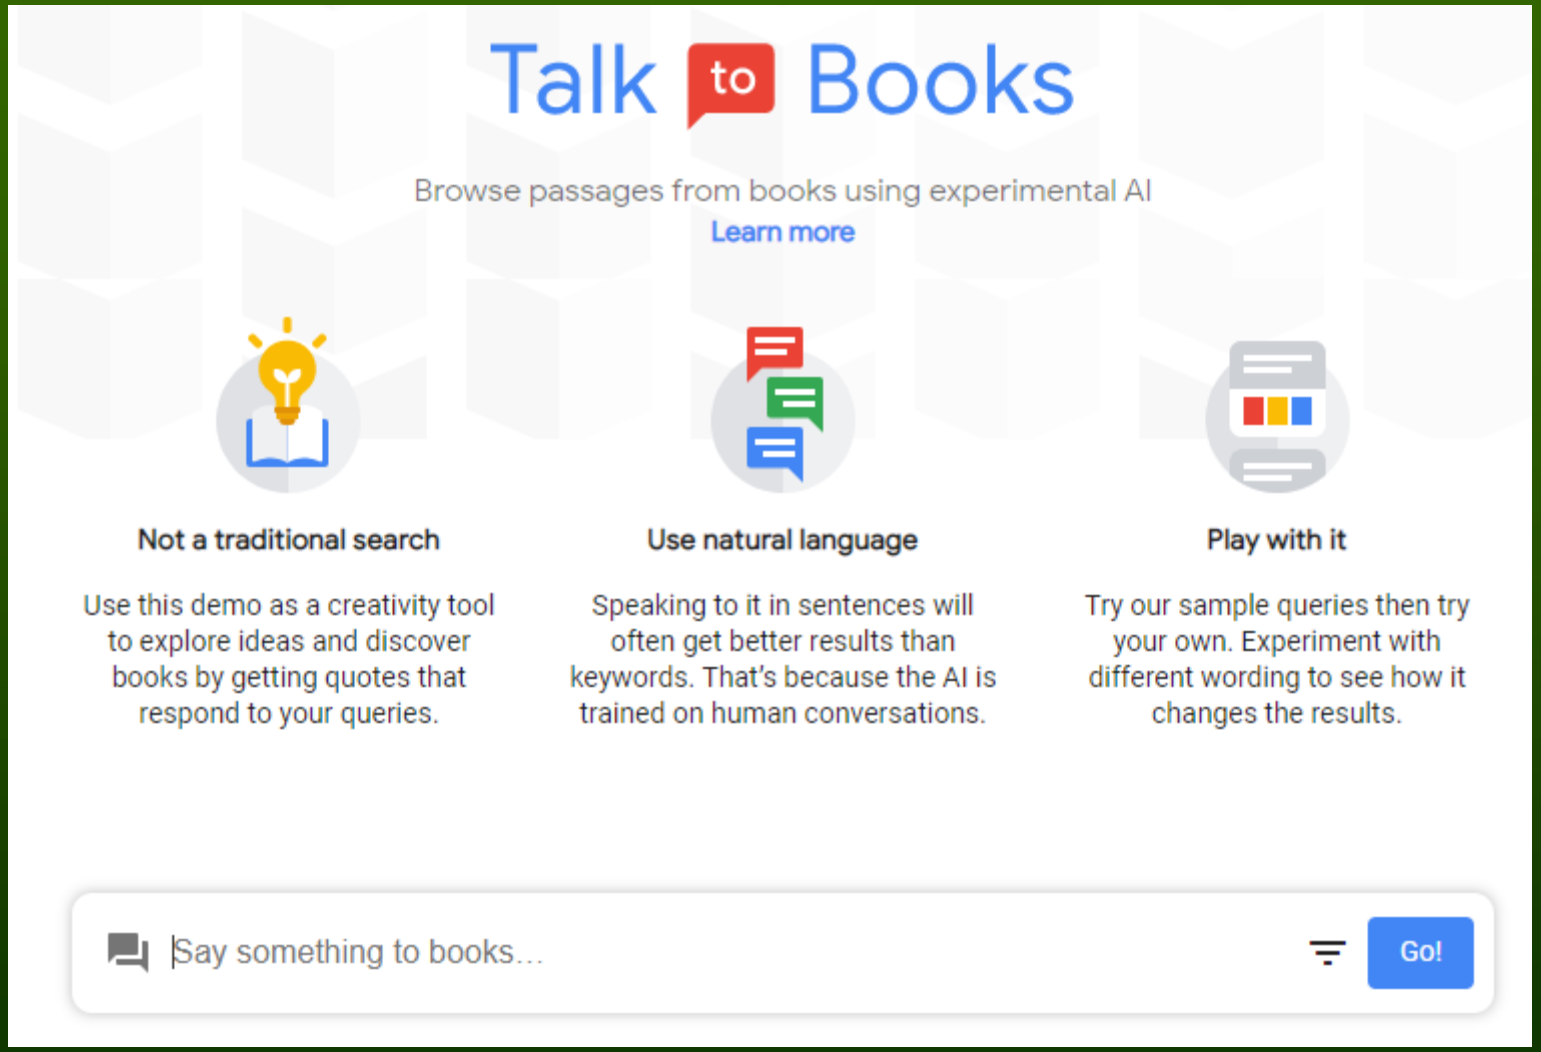

Galactica trained on science, and Consensus for evidence-based answers.

Generative Networks

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

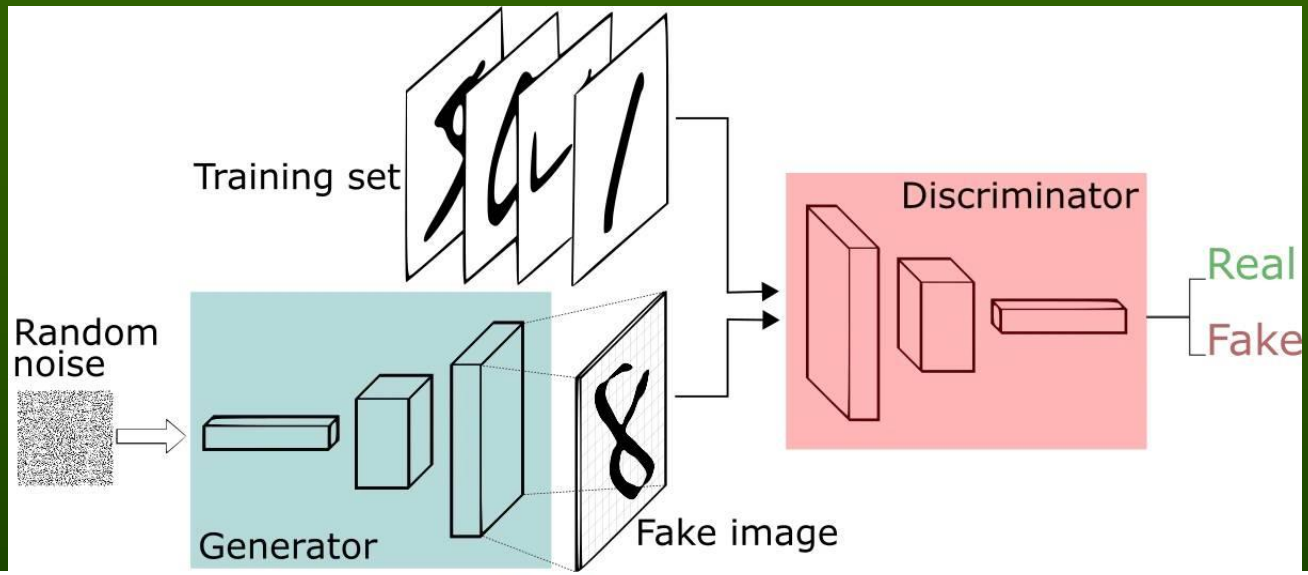

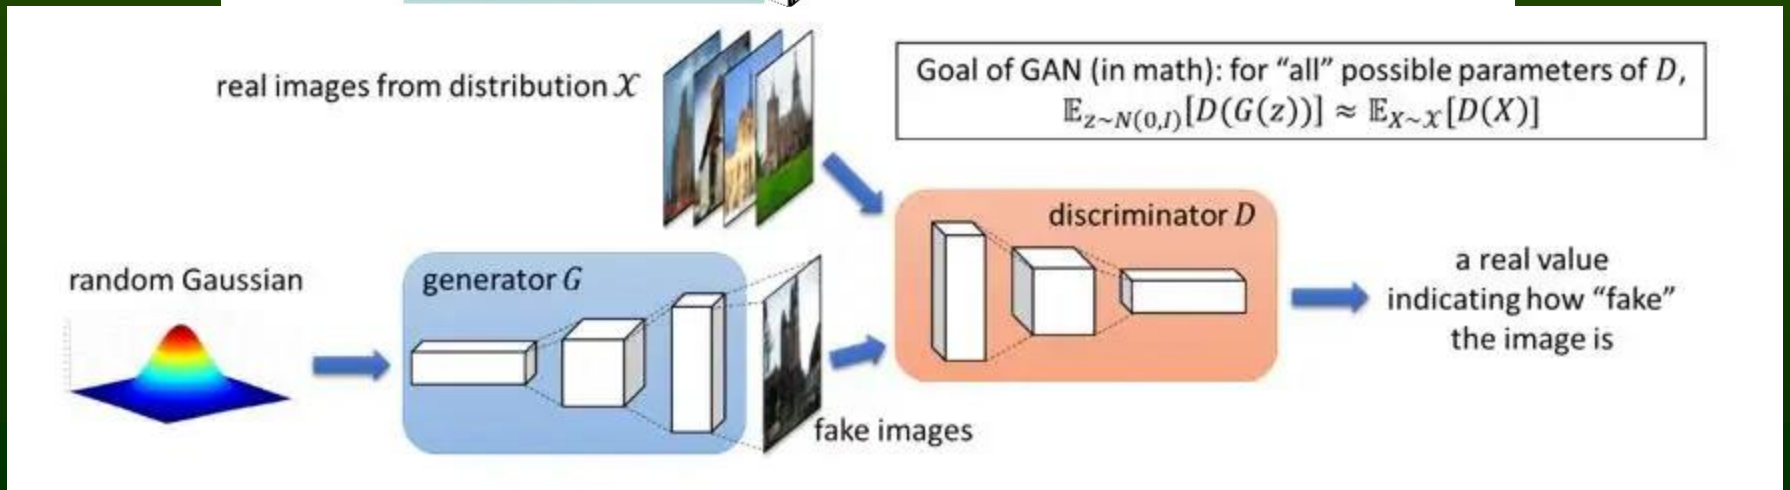

Generative Networks

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.

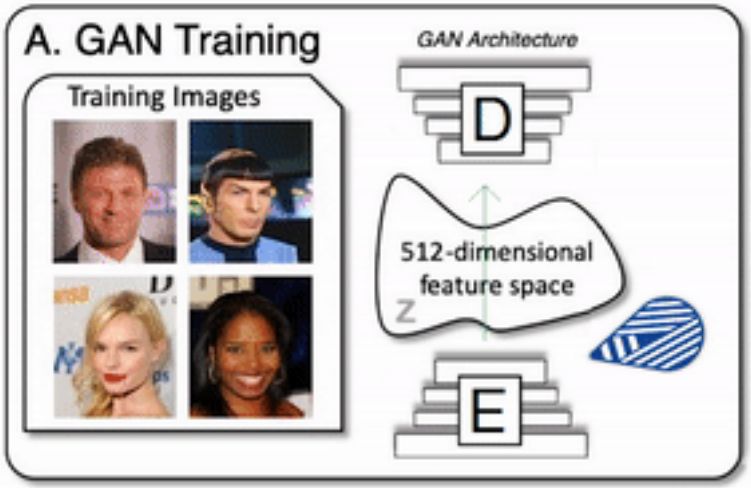

Imagery: Deep Dream

Artificial imagery: Google Deep Dream/Deep Style & Generator, Gallery
LA Gatys, AS Ecker, M Bethge, A Neural Algorithm of Artistic Style (2015)

GAN, Generative Adversarial Networks

Idea (2014): one network generates false examples by distorting training data, the other evaluates whether it is real data. To see is to believe? Not anymore!

Vision-language models

Vision-Language Pre-Trained Models (VL-PTMs): convergence of language, vision, and multimodal pretraining. General-purpose foundation models can be easily adapted to multiple diverse tasks with minimal training.

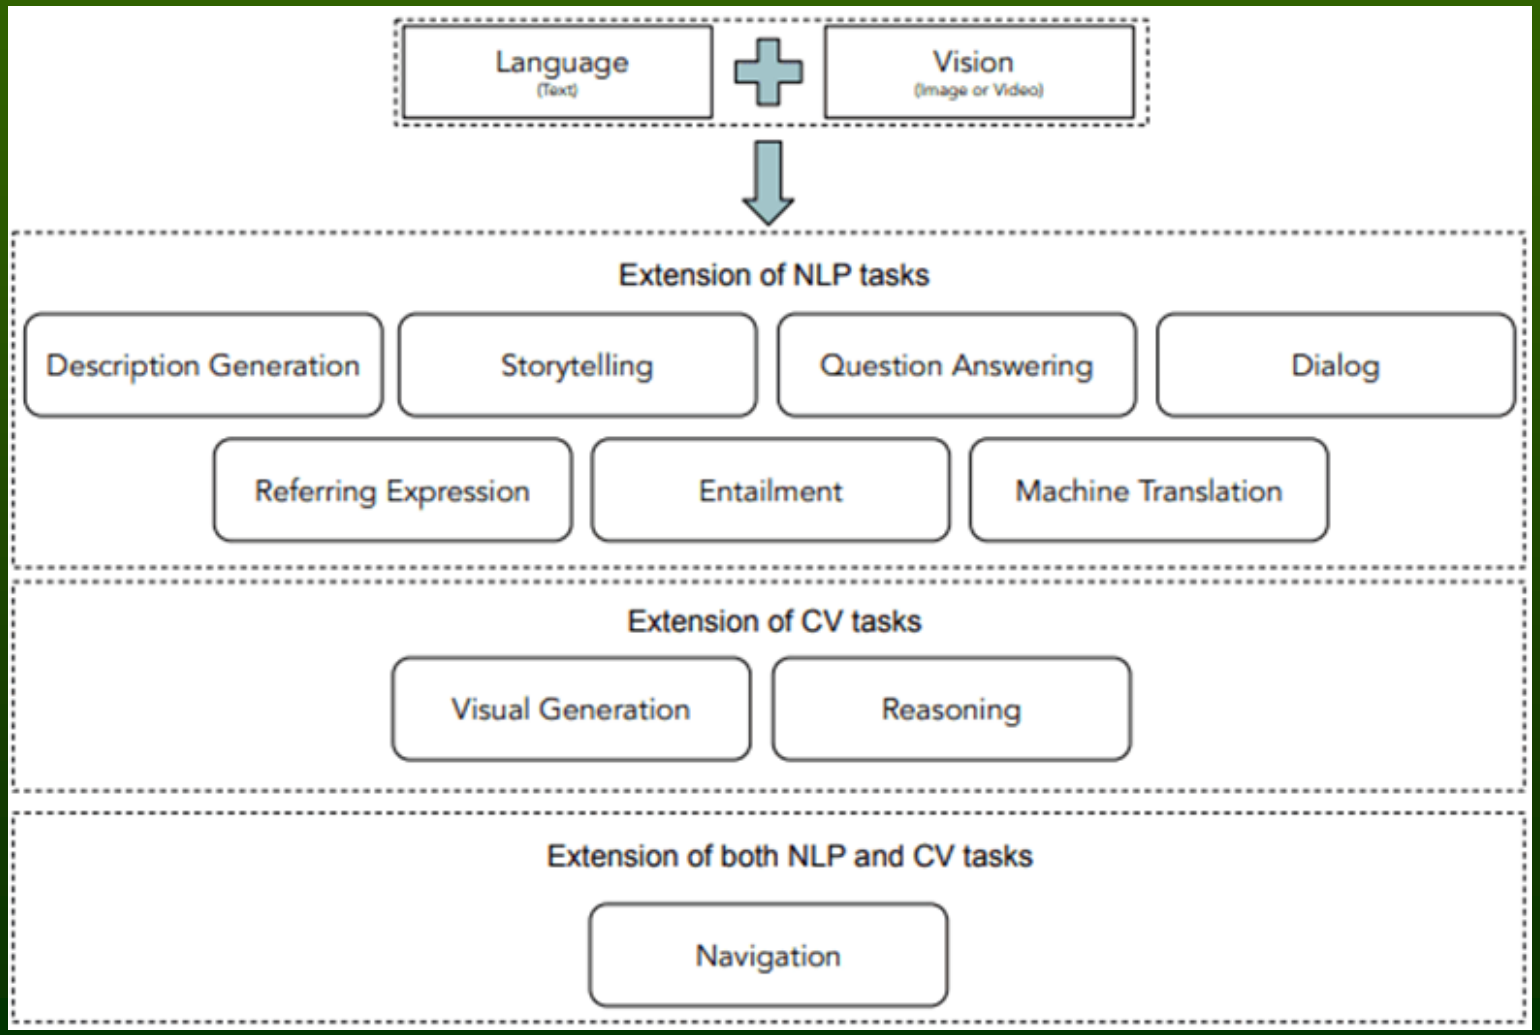

Vision-language models

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.

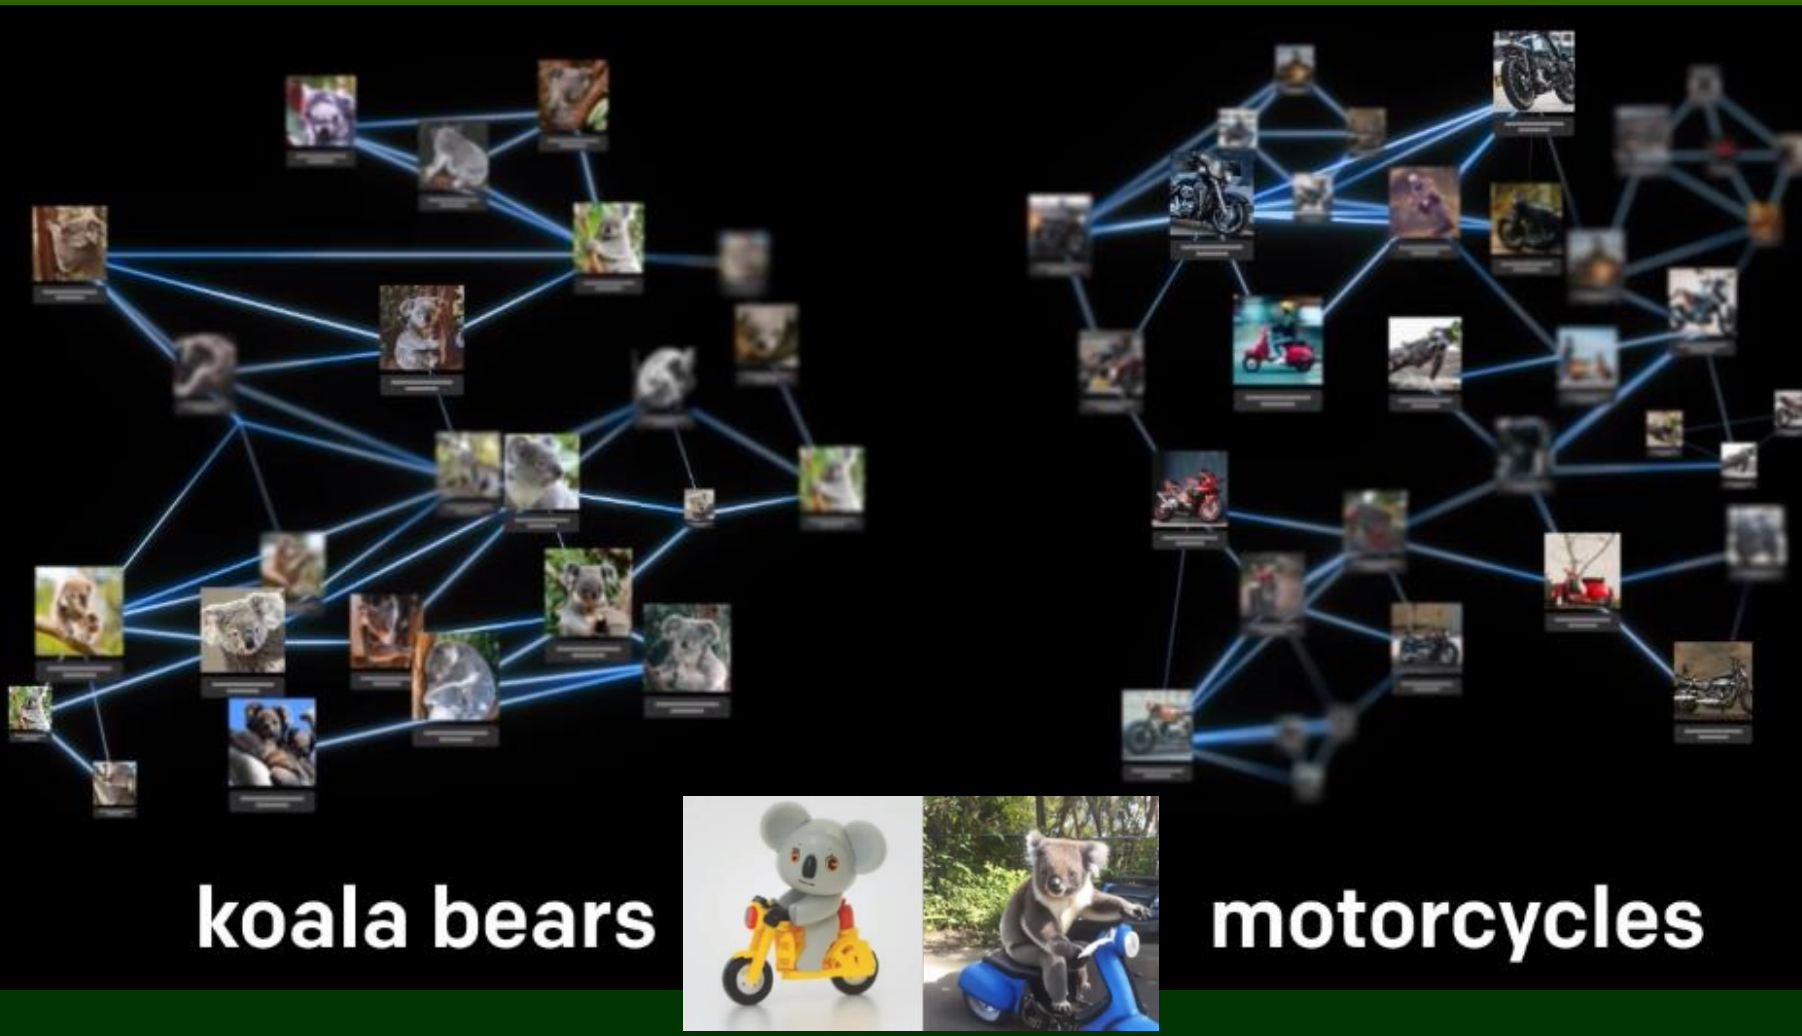

Vision-language generative models

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

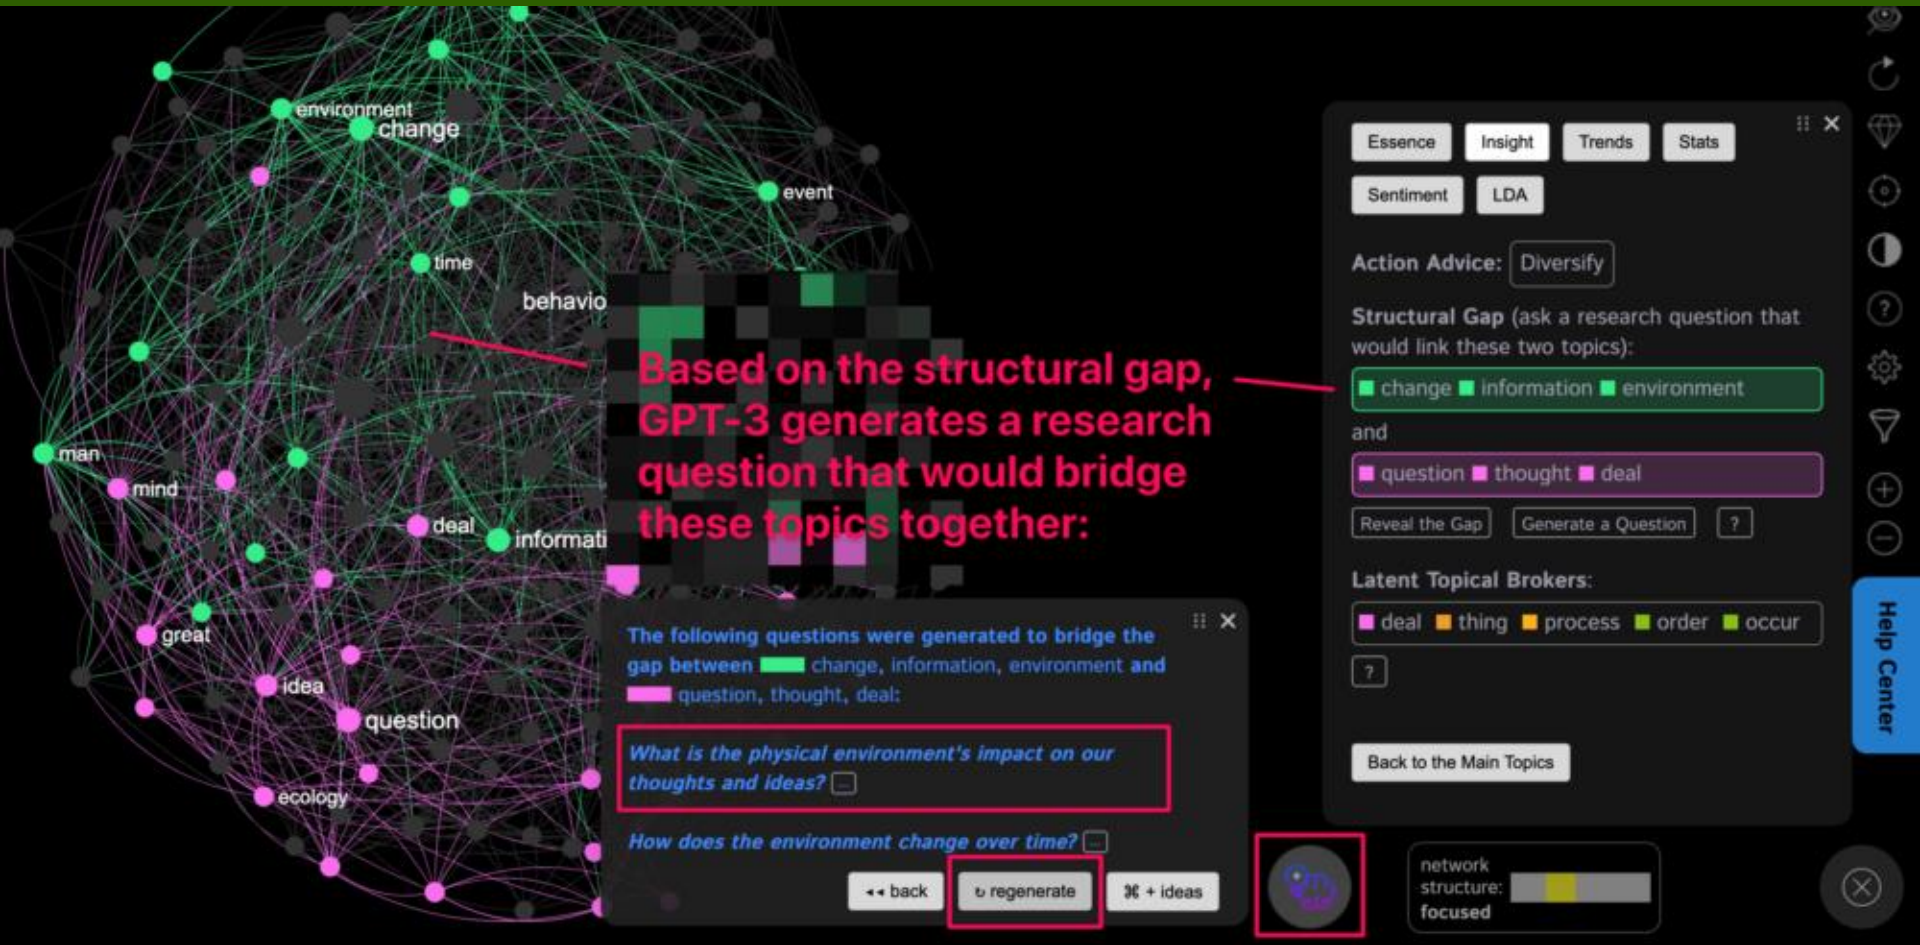

InnerEye

- Transfer experts' knowledge to support AI model personalization and optimization

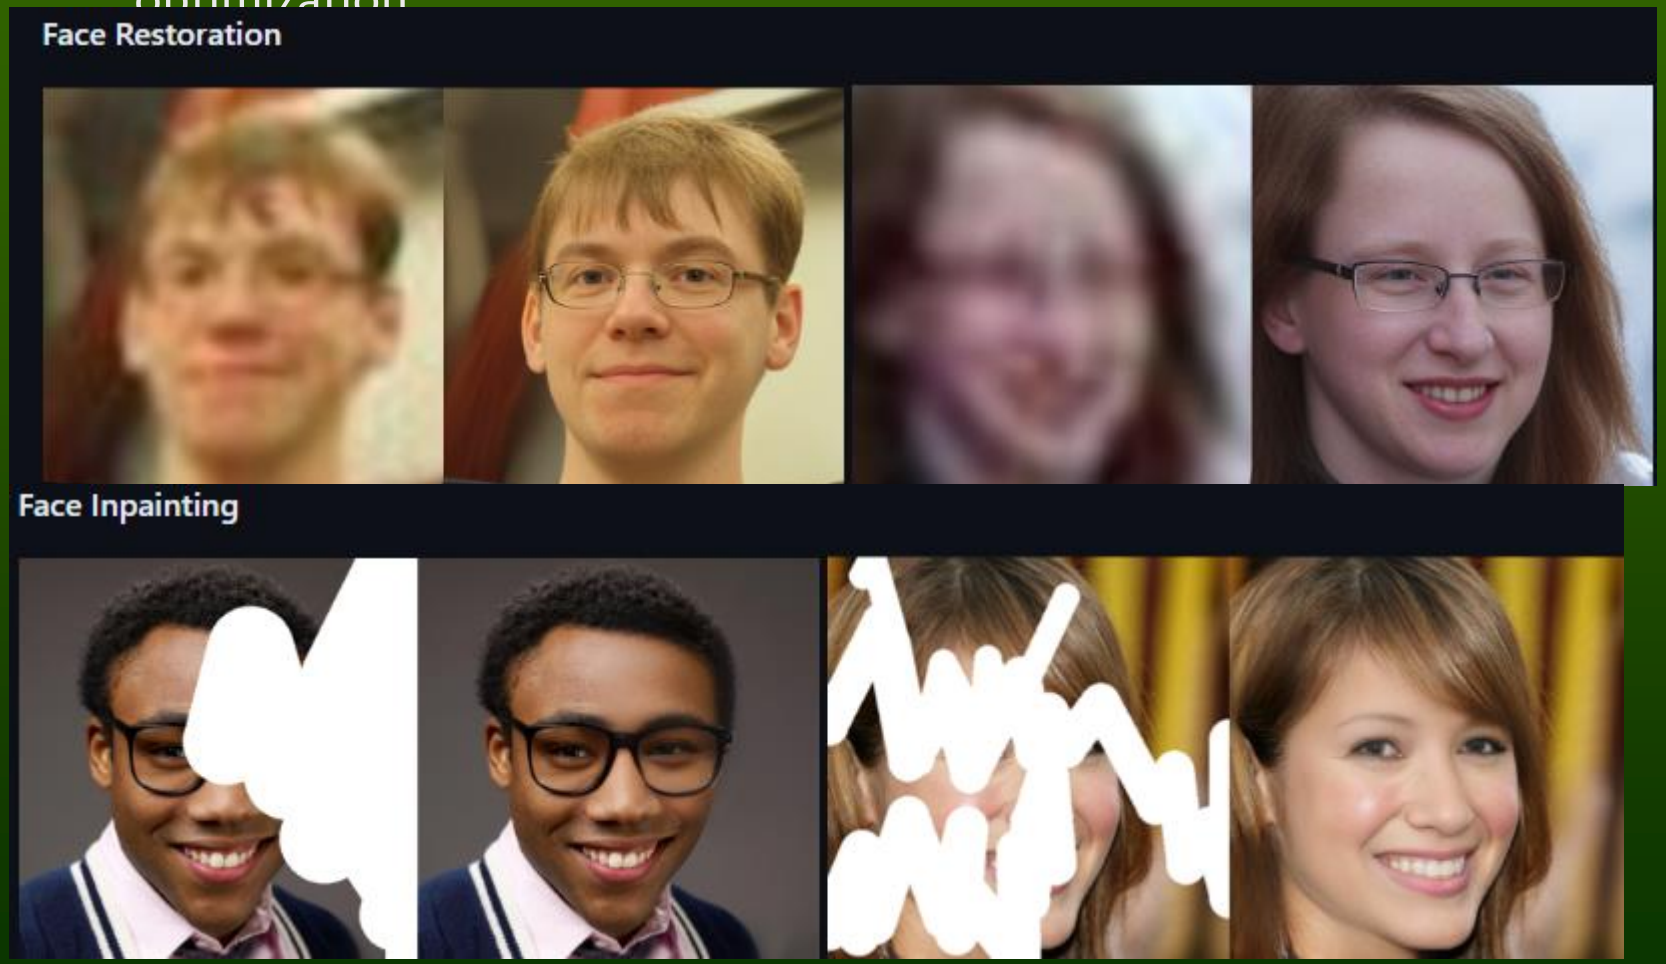

AI imagery

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

These images were created from prompt:

The painting American Ghotic, 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.

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.

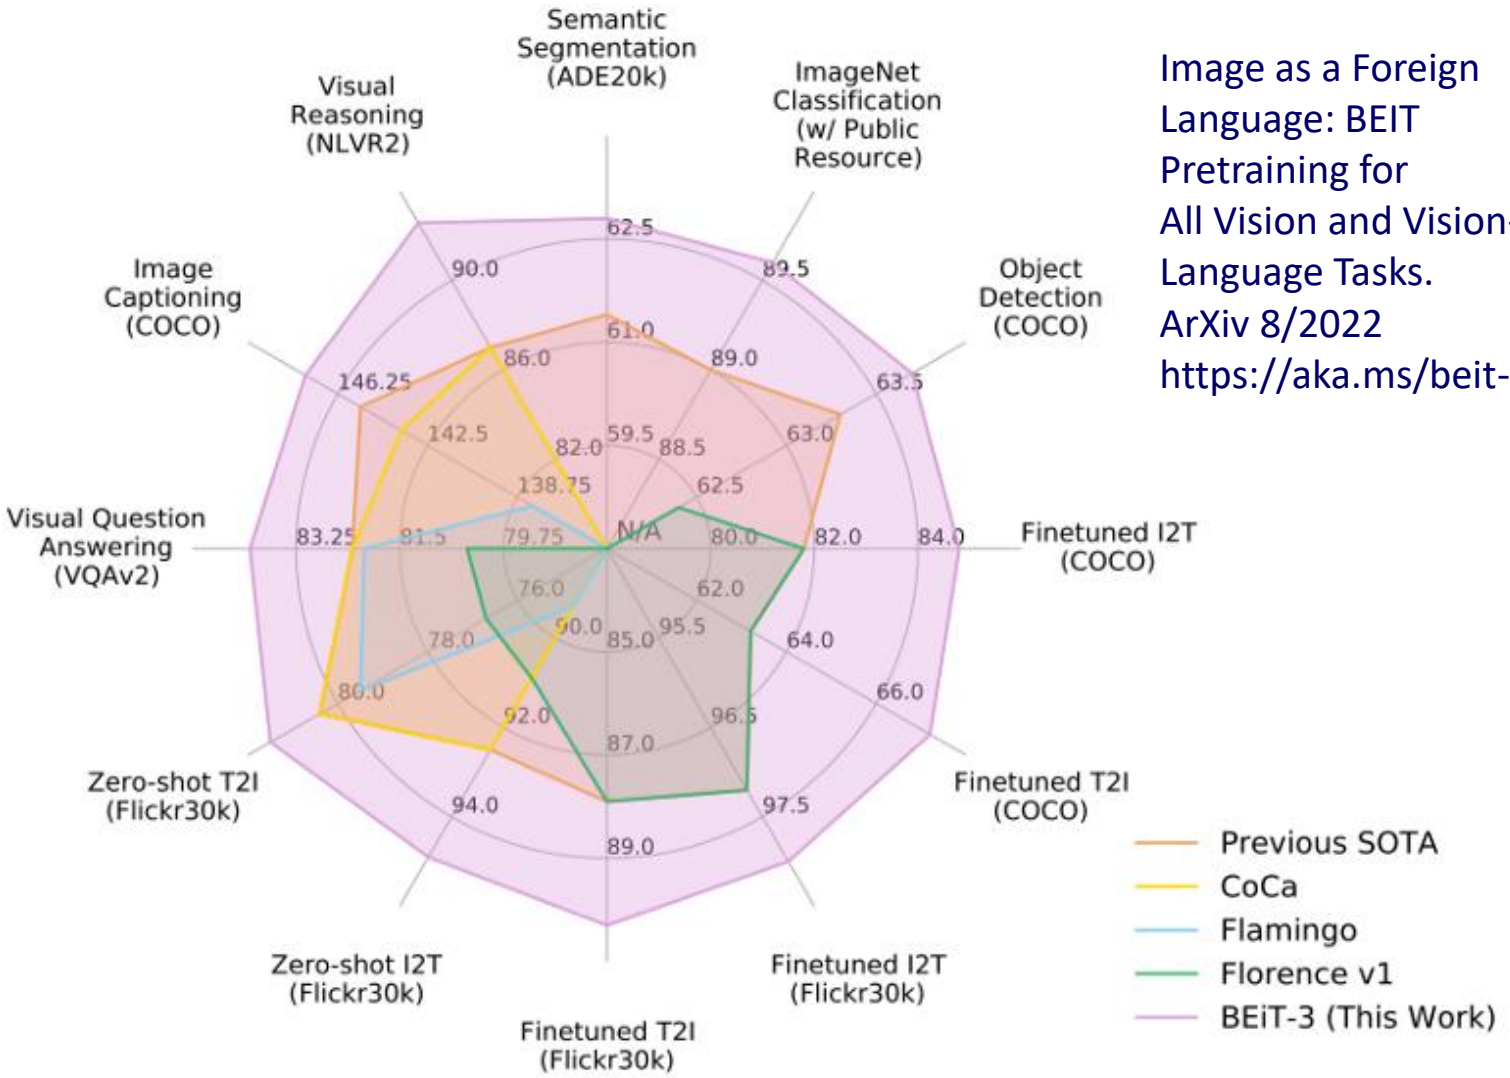

Image as a Foreign Language: BEIT Pretraining for All Vision and Vision-Language Tasks. ArXiv 8/2022 https://aka.ms/beit-3

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 i Apple channel

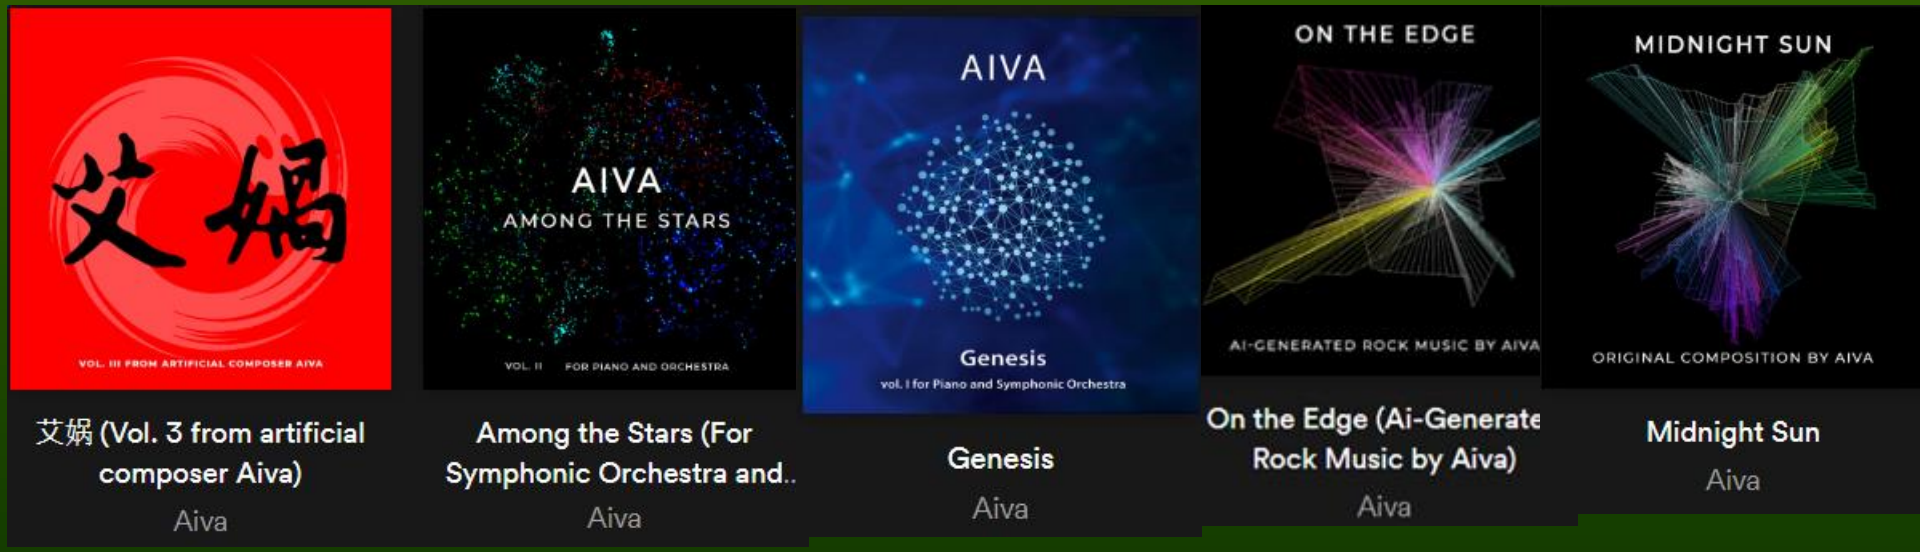

AI completed Beethoven X Symphony in 2021.

Duch W, *Intuition, Insight, Imagination and Creativity*.
IEEE Computational Intelligence Magazine 2(3), August 2007, pp. 40-52

Music AI Ecosystem

THE CREATIVE PROCESS

Water & Music

SONGWRITING/IDEATION

amper · SPLASH · NSYNTH SUPER · Melody Sauce 2 · Loudly · boomy · USICO · LifeScore · SOUNDRAW · soundful · MuseNet · COSO · BandLab SongStarter · Aiva Technologies · ORB Producer Suite

SOUND/SAMPLE SEARCH

splice · COSMOS · AudioStellar · MUSIIO · CYANITE

AUDIO TRANSCRIPTION

BASIC PITCH

AUDIO SYNTHESIS

Dance Diffusion · Never Before Heard Sounds · dyna score · MUSIKA! · SampleRNN · Jukebox · Mubert · Mawf

SOURCE SEPARATION

Spleeter by deezer · Usample · Demucs · audioshake · Audionamix SEPARATE2CREATE

VOICE/SPEECH SYNTHESIS

descript · supertone · voiceful · UBERDUCK · X Studio · SOLARIA

MIXING/MASTERING

FAST LIMITER · songmastr · PreTube · TAIP · Nectar · master channel · Ozone · LANDR · RoEx · BandLab Mastering

PRODUCTION (OTHER)

neutone · AmpliTube · BRONZE · LALAL.AI · xln audio · Chameleon · Splitter.ai

LYRICS/TEXT

These Lyrics Do Not Exist · BRAIN./RAP · sudo write

waterandmusic.com

Part II

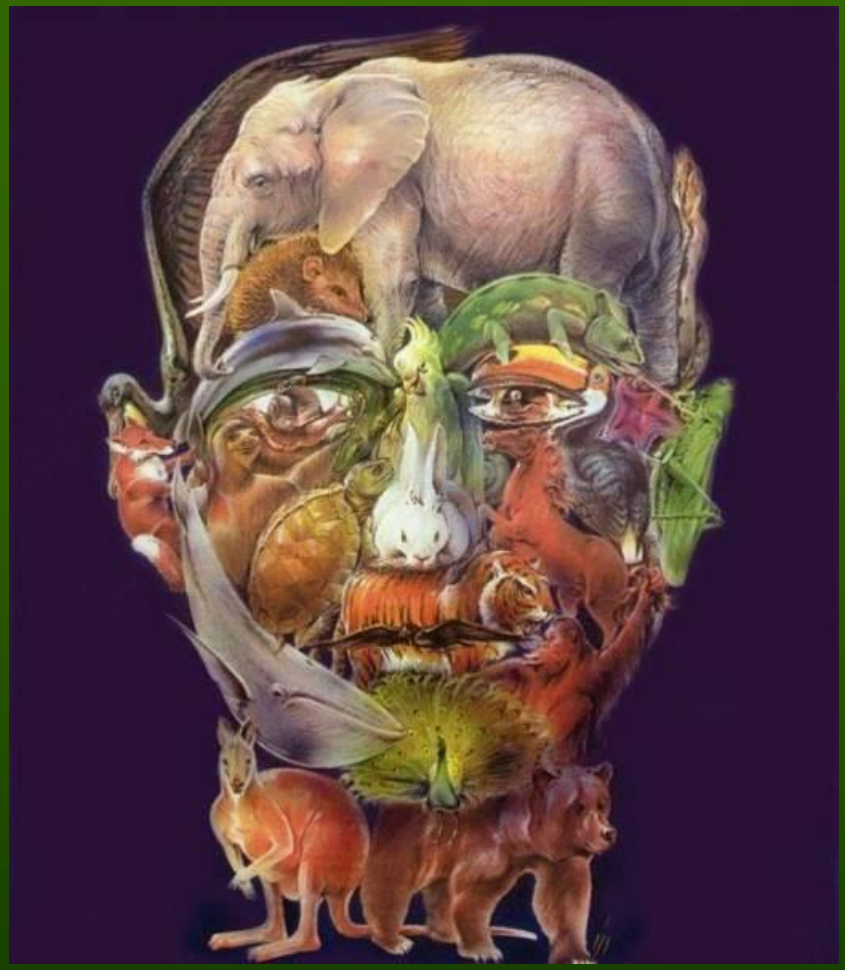

Google: Wlodek Duch
=> talks, papers, lectures ...