یادگیری عمیق

جلسه ۳

مروری بر یادگیری عمیق

A Review on Deep Learning

کاظم فولادی قلعه

دانشکده مهندسی، دانشکدگان فارابی

دانشگاه تهران

http://courses.fouladi.ir/deep

مروری بر یادگیری عمیق

۱

مقدمات

What is Artificial Intelligence?

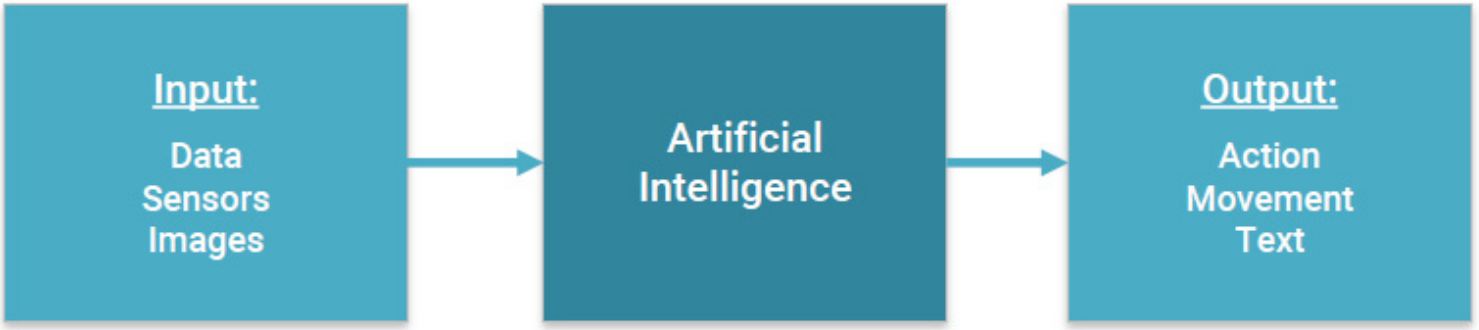

What is Artificial Intelligence?

Artificial Narrow Intelligence (ANI): Machine intelligence that equals or exceeds human intelligence or efficiency at a specific task.

Artificial General Intelligence (AGI): A machine with the ability to apply intelligence to any problem, rather than just one specific problem *(human-level intelligence)*.

Artificial Superintelligence (ASI): An intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom and social skills.

http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html

What is Artificial Intelligence?

Superintelligence

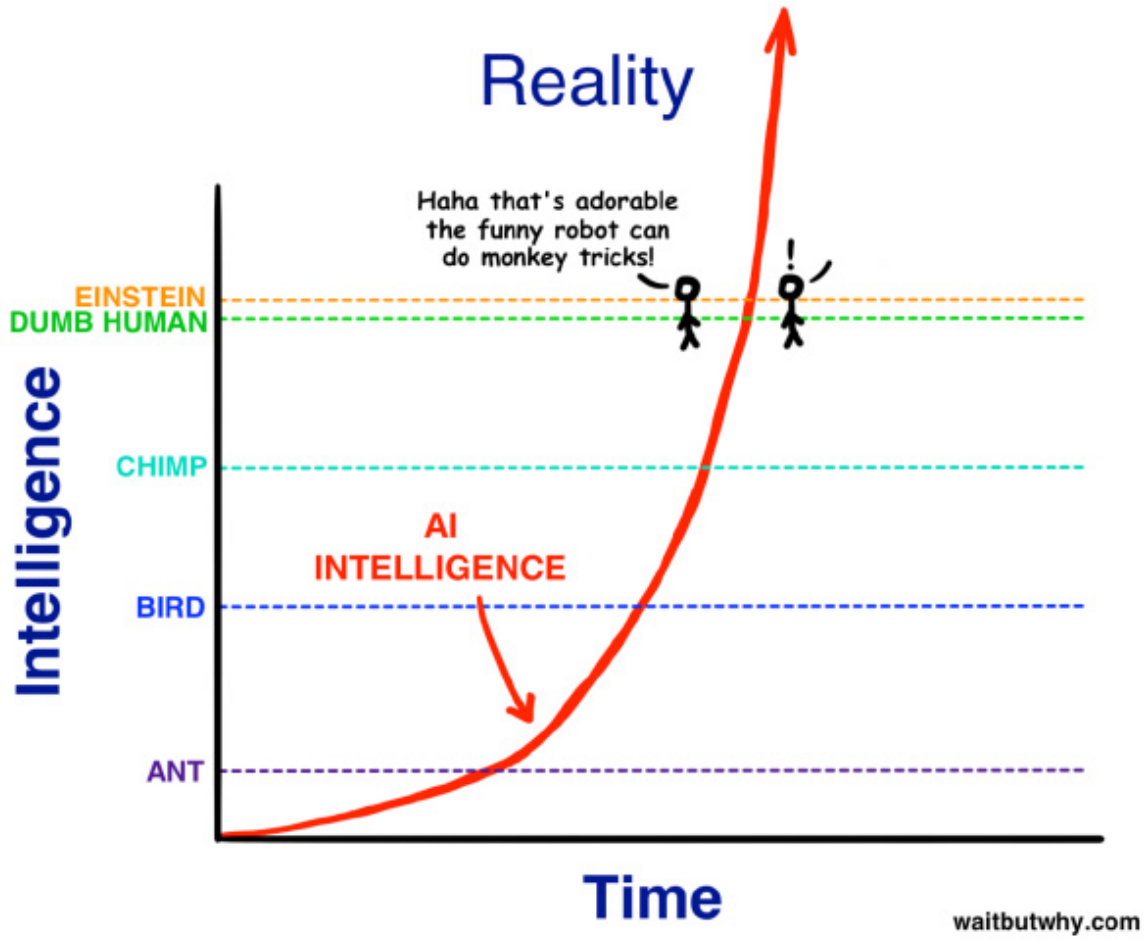

A superintelligence is any intellect that vastly outperforms the best human brains in practically every field, including scientific creativity, general wisdom, and social skills

Source

Machine Learning - Basics

Introduction

Machine Learning is a type of Artificial Intelligence that provides computers with the ability to learn without being explicitly programmed.

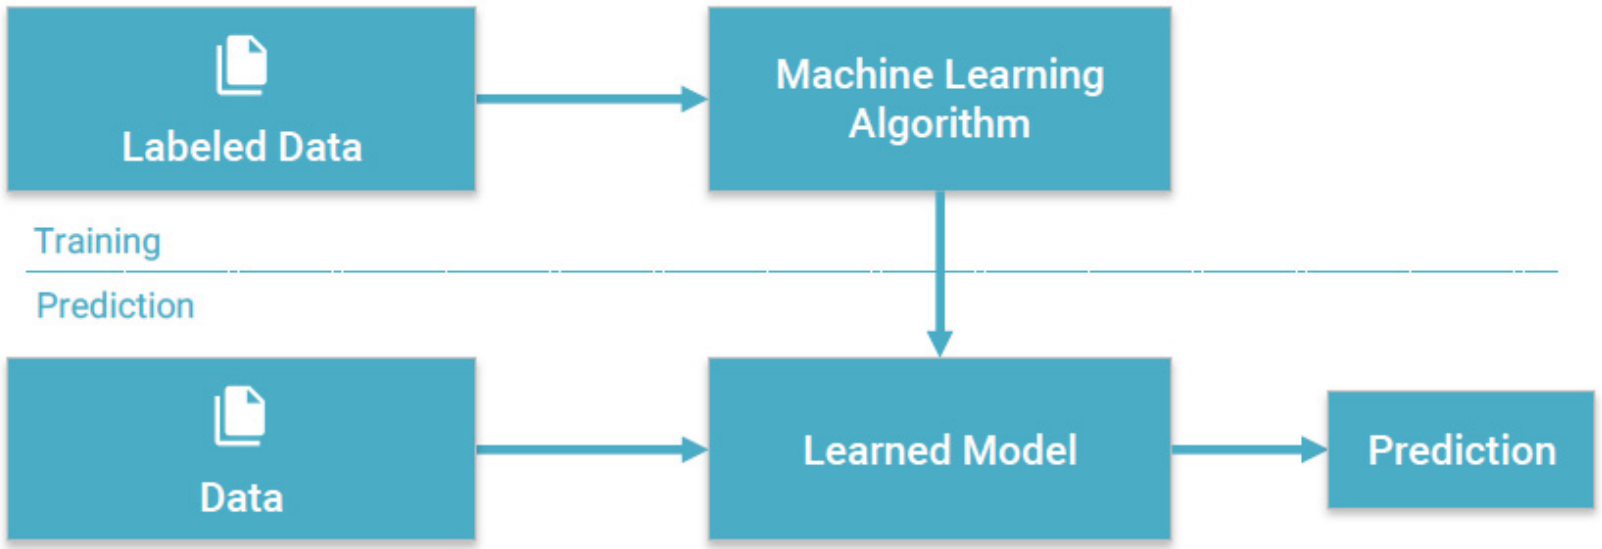

Provides various techniques that can learn from and make predictions on data

Machine Learning - Basics

Learning Approaches

Supervised Learning: Learning with a labeled training set
Example: email spam detector with training set of already labeled emails

Unsupervised Learning: Discovering patterns in unlabeled data
Example: cluster similar documents based on the text content

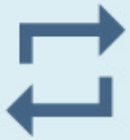

Reinforcement Learning: learning based on feedback or reward
Example: learn to play chess by winning or losing

Machine Learning - Basics

Problem Types

Classification
(supervised – predictive)

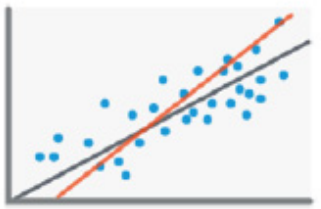

Regression
(supervised – predictive)

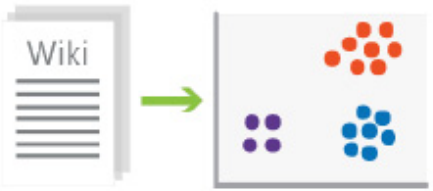

Clustering
(unsupervised – descriptive)

Anomaly Detection
(unsupervised– descriptive)

https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now

Machine Learning - Basics

Problem Types – Mapping from A to B

What Machine Learning Can Do

A simple way to think about supervised learning.

INPUT A	RESPONSE B	APPLICATION
Picture	Are there human faces? (0 or 1)	Photo tagging
Loan application	Will they repay the loan? (0 or 1)	Loan approvals
Ad plus user information	Will user click on ad? (0 or 1)	Targeted online ads
Audio clip	Transcript of audio clip	Speech recognition
English sentence	French sentence	Language translation
Sensors from hard disk, plane engine, etc.	Is it about to fail?	Preventive maintenance
Car camera and other sensors	Position of other cars	Self-driving cars

SOURCE ANDREW NG

Source

www.acttr.com/images/faq/scikit-learn-plot-classifier.png

Machine Learning - Basics

Algorithms Comparison - Classification

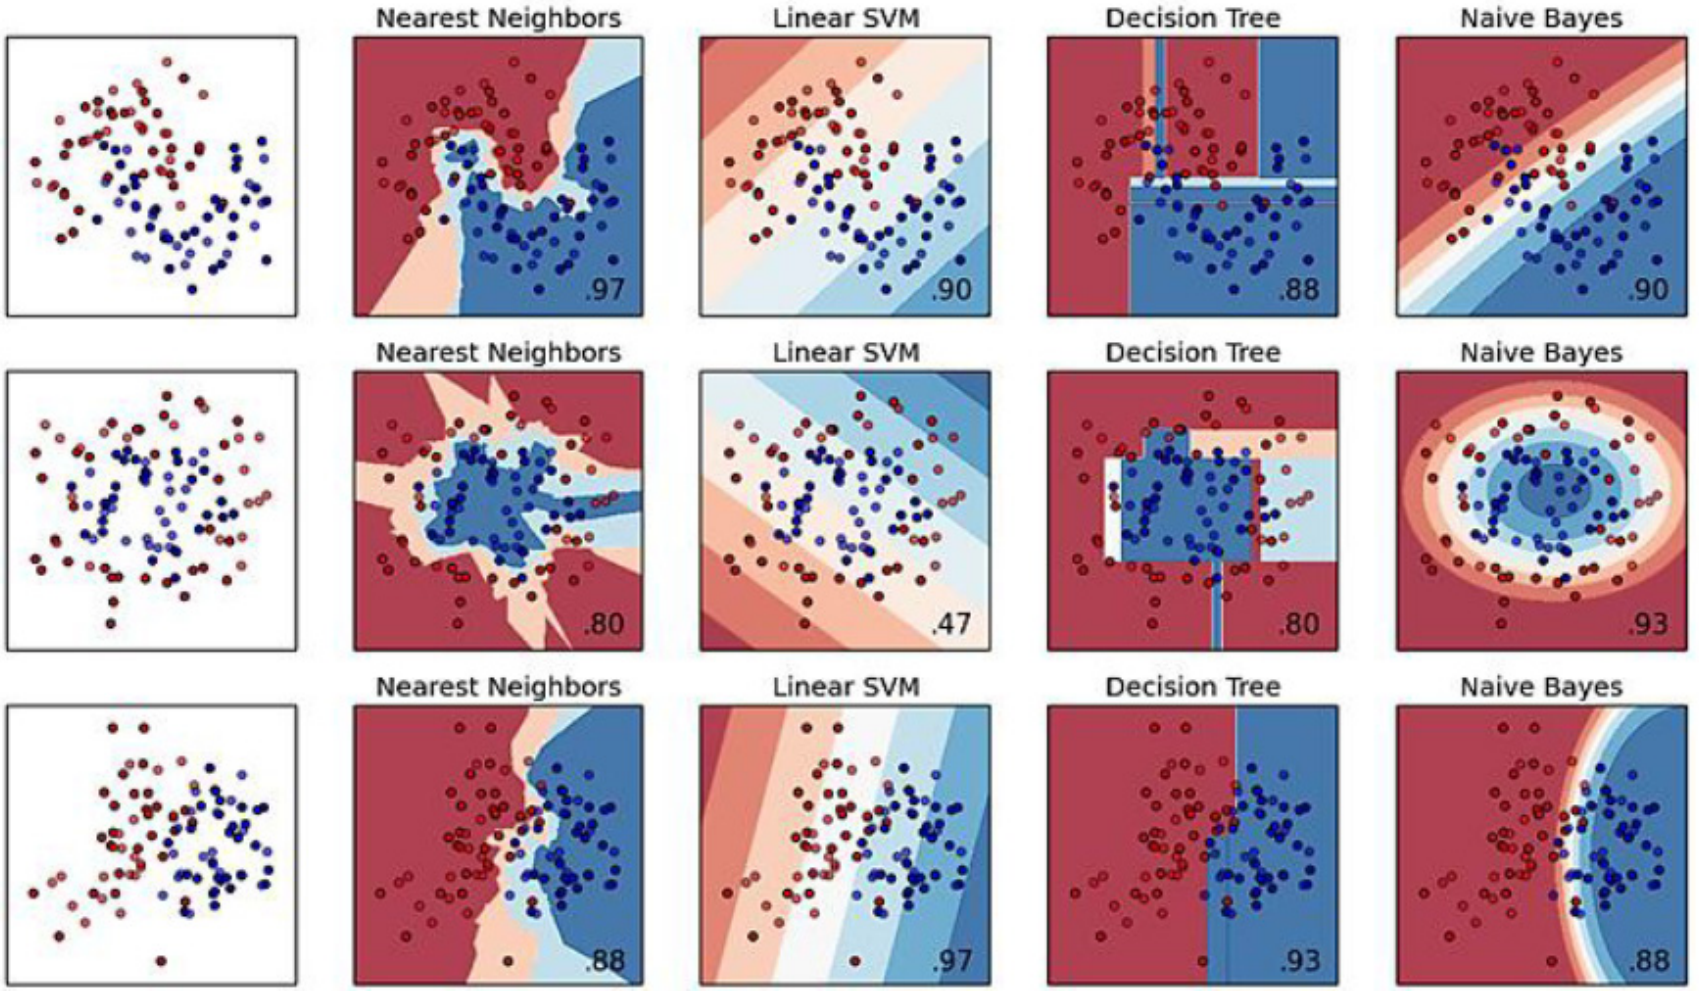

Source

مروری بر یادگیری عمیق

۲

یادگیری عمیق چیست؟

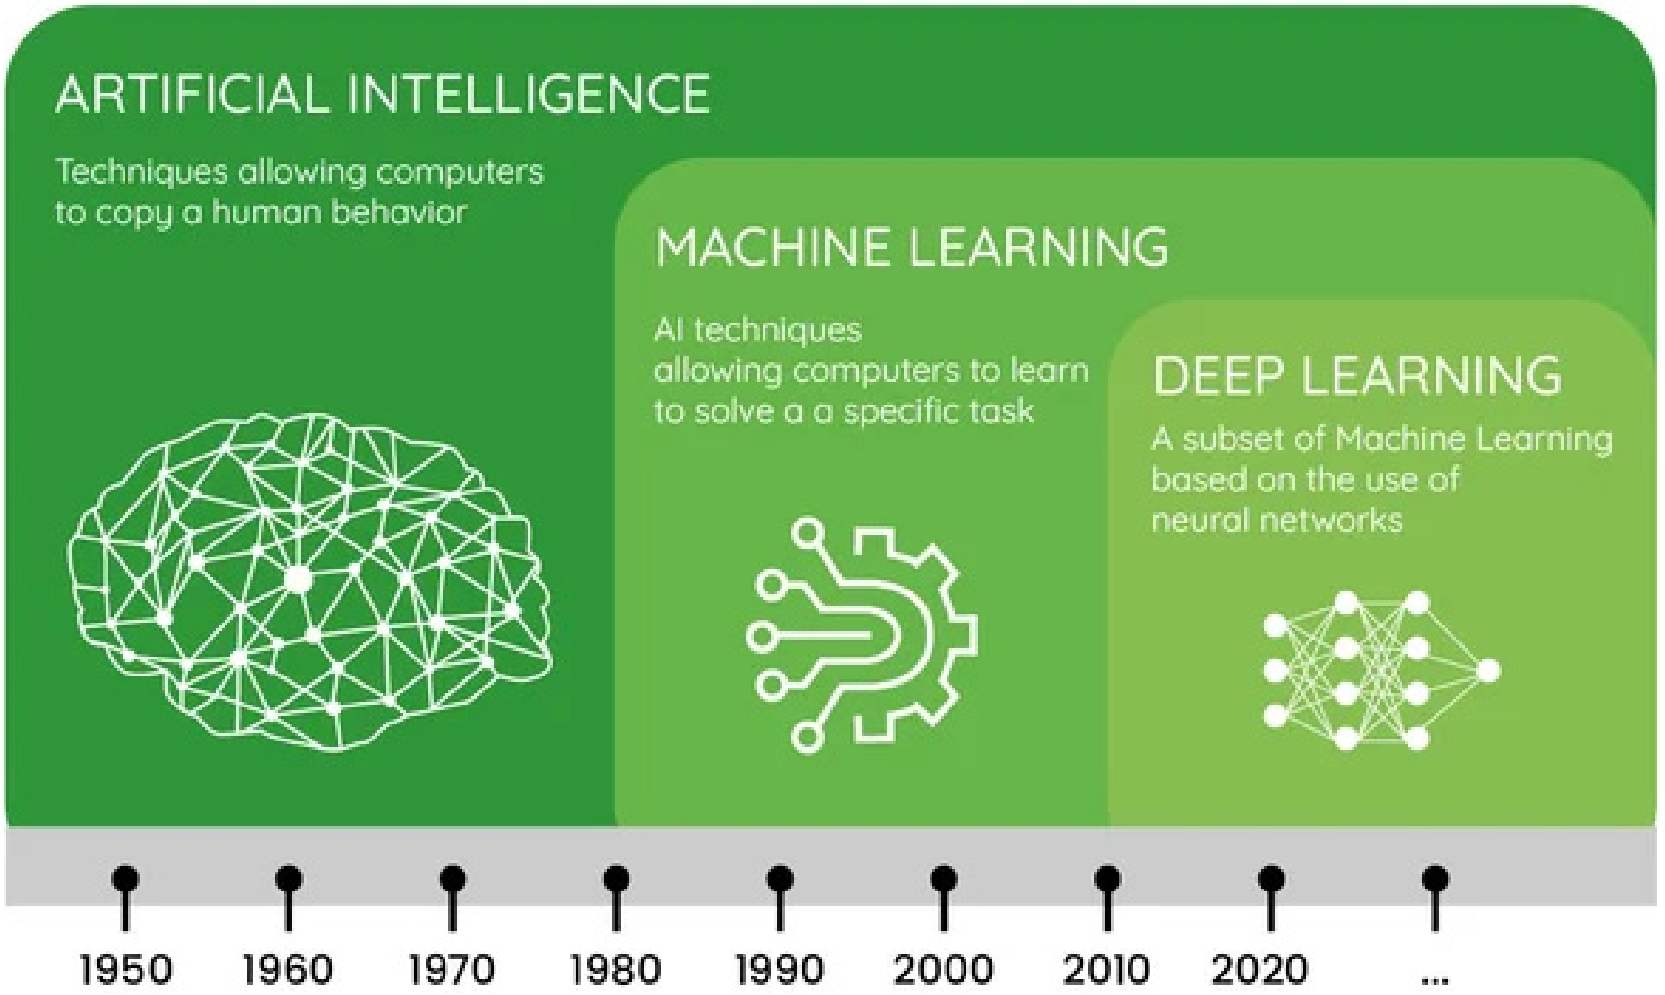
ARTIFICIAL INTELLIGENCE
Techniques allowing computers to copy a human behavior
MACHINE LEARNING
AI techniques allowing computers to learn to solve a a specific task
DEEP LEARNING
A subset of Machine Learning based on the use of neural networks
1950
1960
1970
1980
1990
2000
2010
2020
...

1950

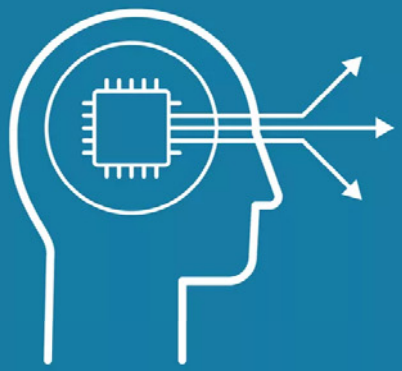

ARTIFICIAL INTELLIGENCE

ENGINEERING OF MACHINES
THAT MIMIC COGNITIVE FUNCTIONS

1980

MACHINE LEARNING

ABILITY TO PERFORM TASKS
WITHOUT EXPLICIT INSTRUCTIONS
AND RELYING ON PATTERNS

2010

DEEP LEARNING

MACHINE LEARNING BASED
ON ARTIFICIAL NEURAL NETWORKS

What is Deep Learning?

Part of the machine learning field of learning representations of data. Exceptional effective at learning patterns.

Utilizes learning algorithms that derive meaning out of data by using a hierarchy of multiple layers that mimic the neural networks of our brain.

If you provide the system tons of information, it begins to understand it and respond in useful ways.

یادگیری عمیق چیست؟

WHAT IS DEEP LEARNING?

بخشی از حوزه‌ی یادگیری ماشینی که به یادگیری بازنمایی داده‌ها می‌پردازد؛
و به‌طور استثنایی در یادگیری الگوها مؤثر است ...

یادگیری عمیق از الگوریتم‌هایی بهره می‌برد که
با استفاده از یک سلسله مراتب از چندین لایه، از داده‌ها معنا استخراج می‌کند
(با تقلید از شبکه‌های عصبی مغز انسان).

اگر برای این سیستم مقدار زیادی اطلاعات فراهم کنیم،
شروع به فهم آنها می‌کند و به‌صورت معمول پاسخ می‌دهد.

Inspired by the Brain

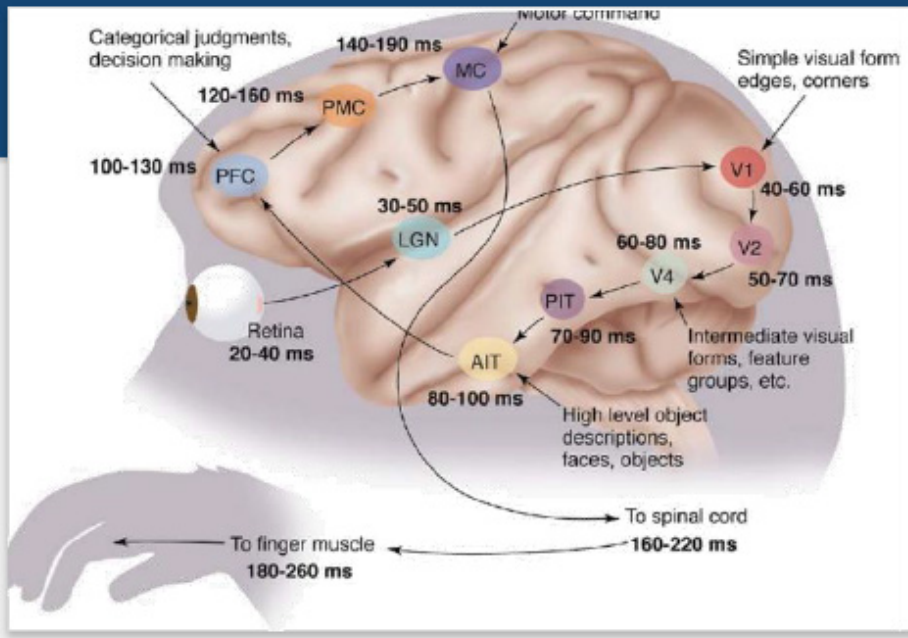

The first hierarchy of neurons that receives information in the visual cortex are sensitive to specific edges while brain regions further down the visual pipeline are sensitive to more complex structures such as faces.

Our brain has lots of neurons connected together and the strength of the connections between neurons represents long term knowledge.

Deep Learning - Basics

Architecture

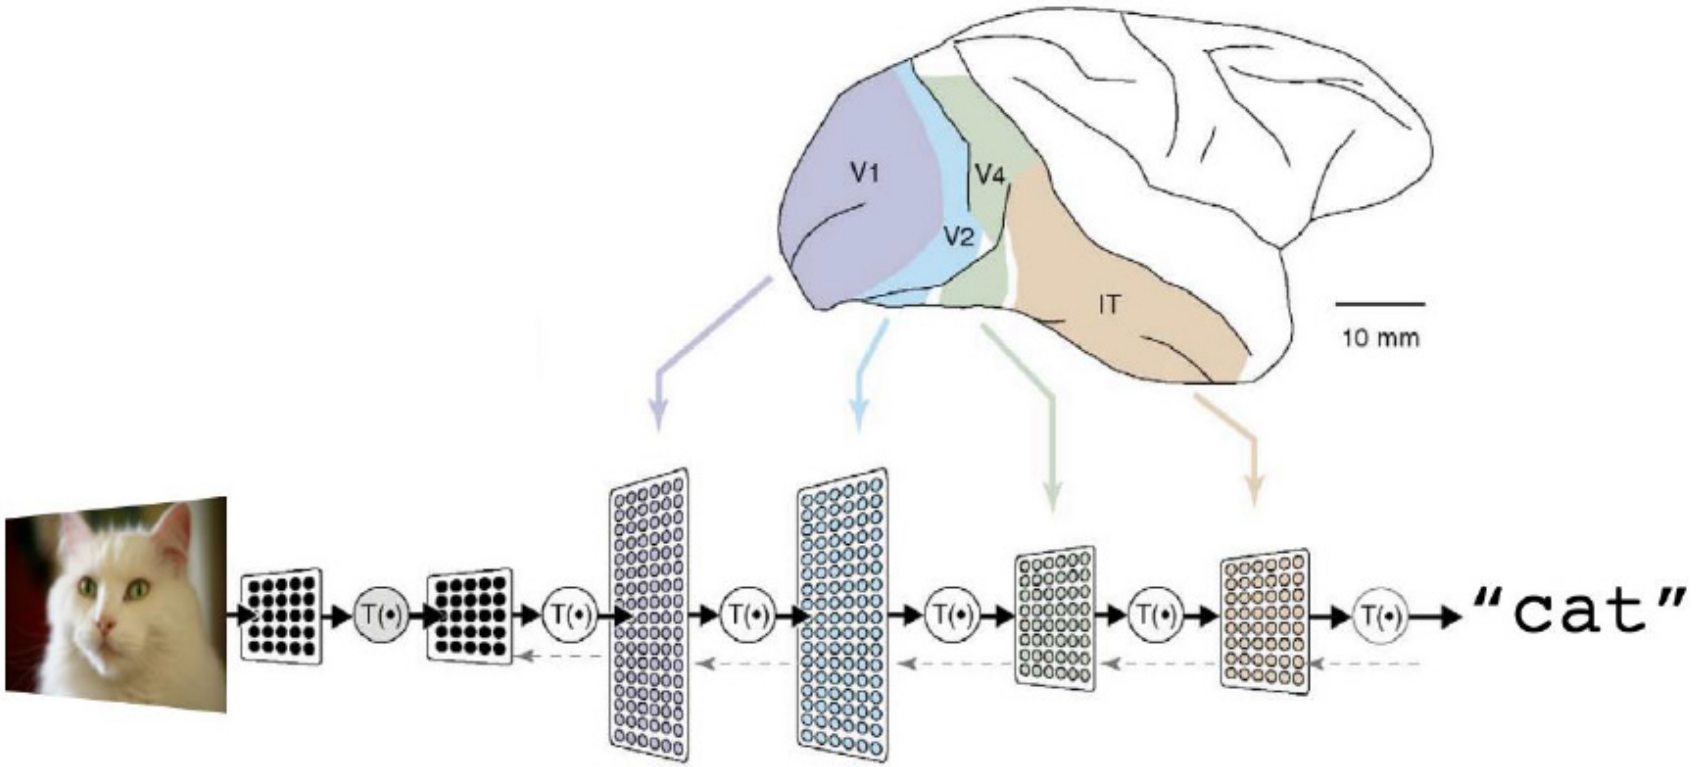

A deep neural network consists of a hierarchy of layers, whereby each layer transforms the input data into more abstract representations (e.g. edge -> nose -> face). The output layer combines those features to make predictions.

Deep Learning - Basics

What did it learn?

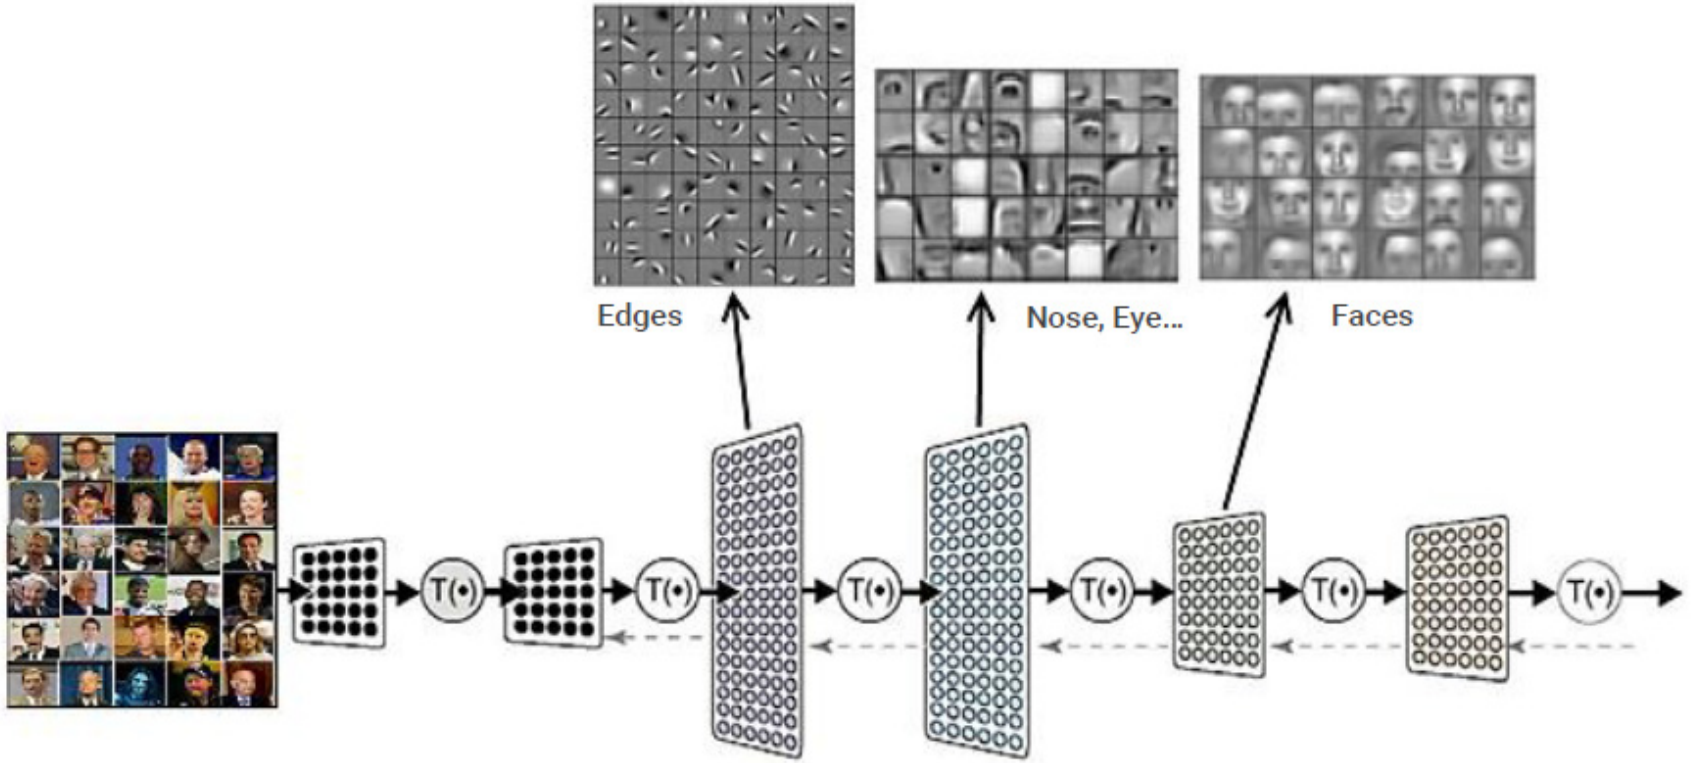

Deep Learning - Basics

No more feature engineering

مهندسی ویژگی‌ها

حذف مرحله‌ی مهندسی ویژگی‌ها در یادگیری عمیق

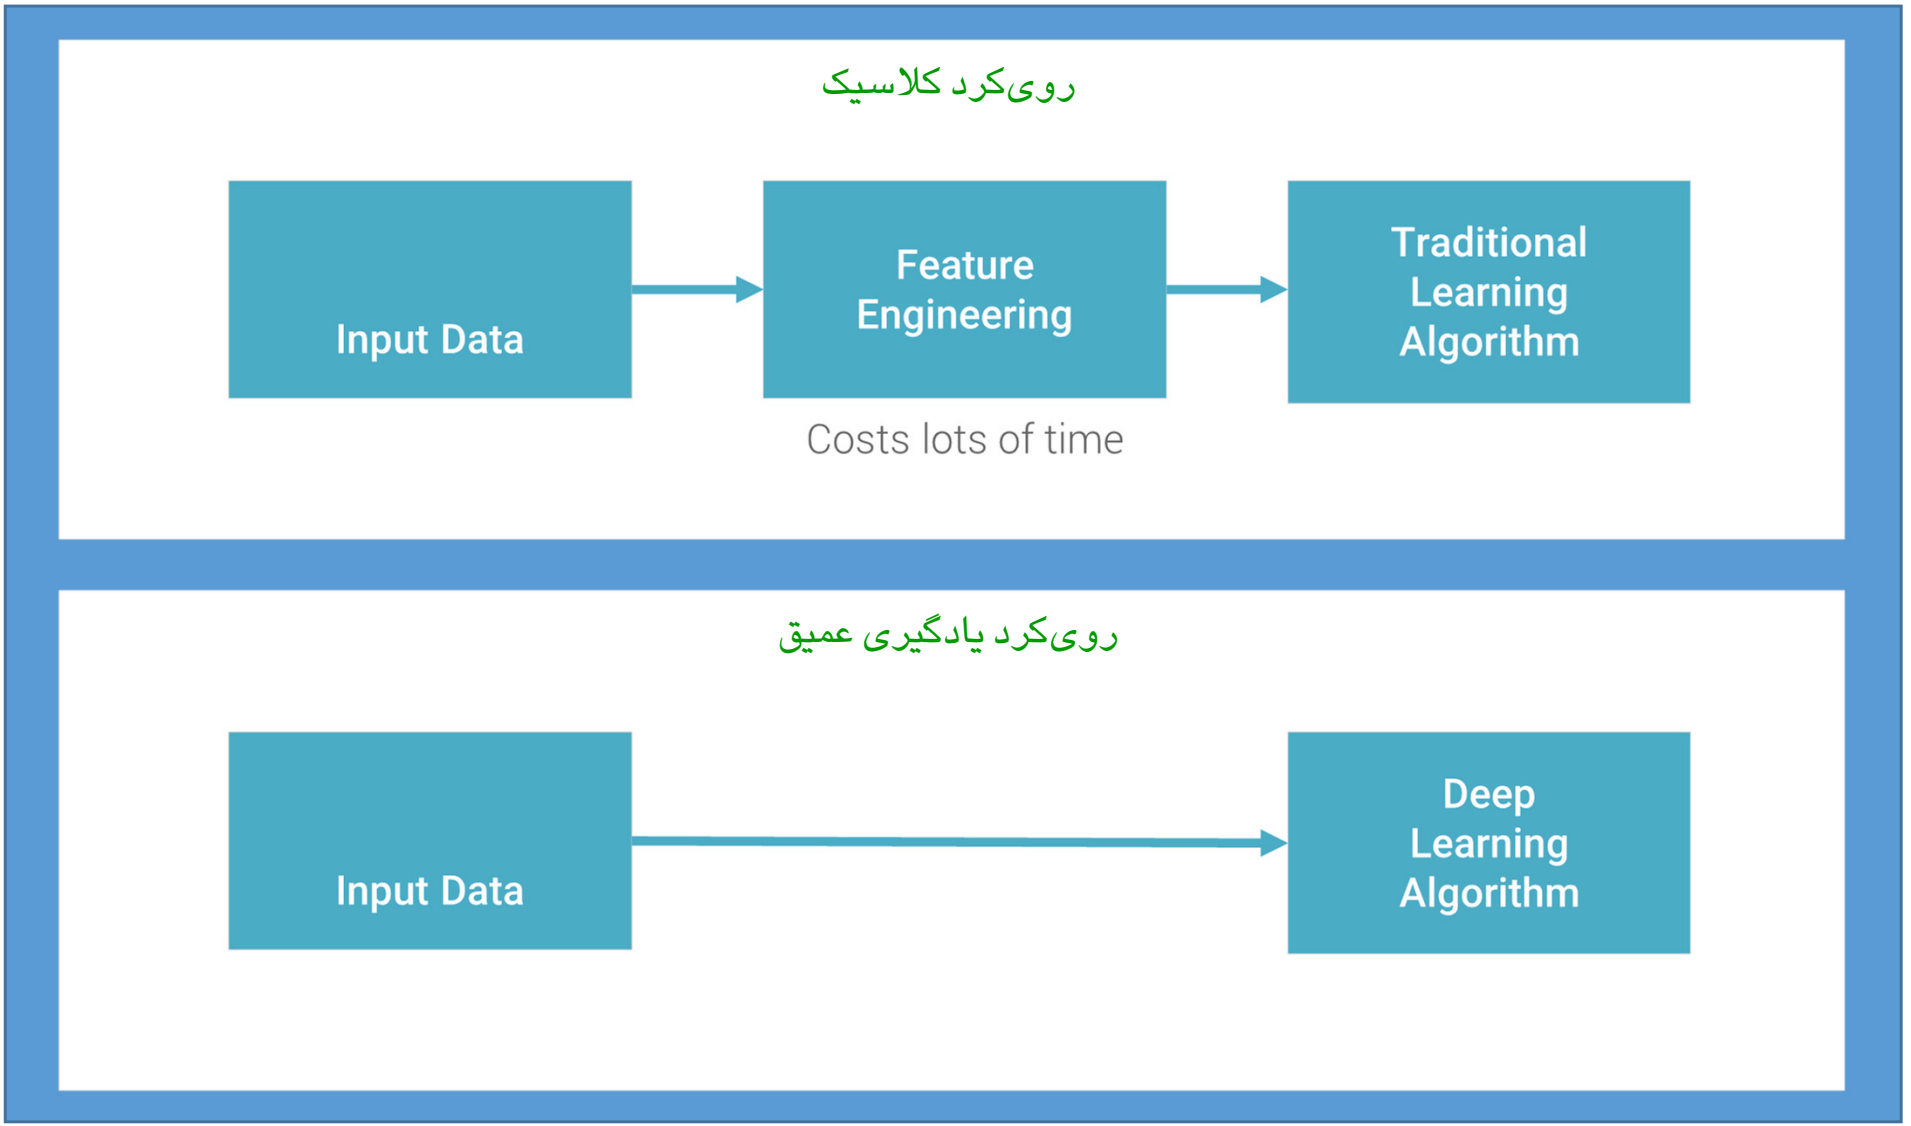

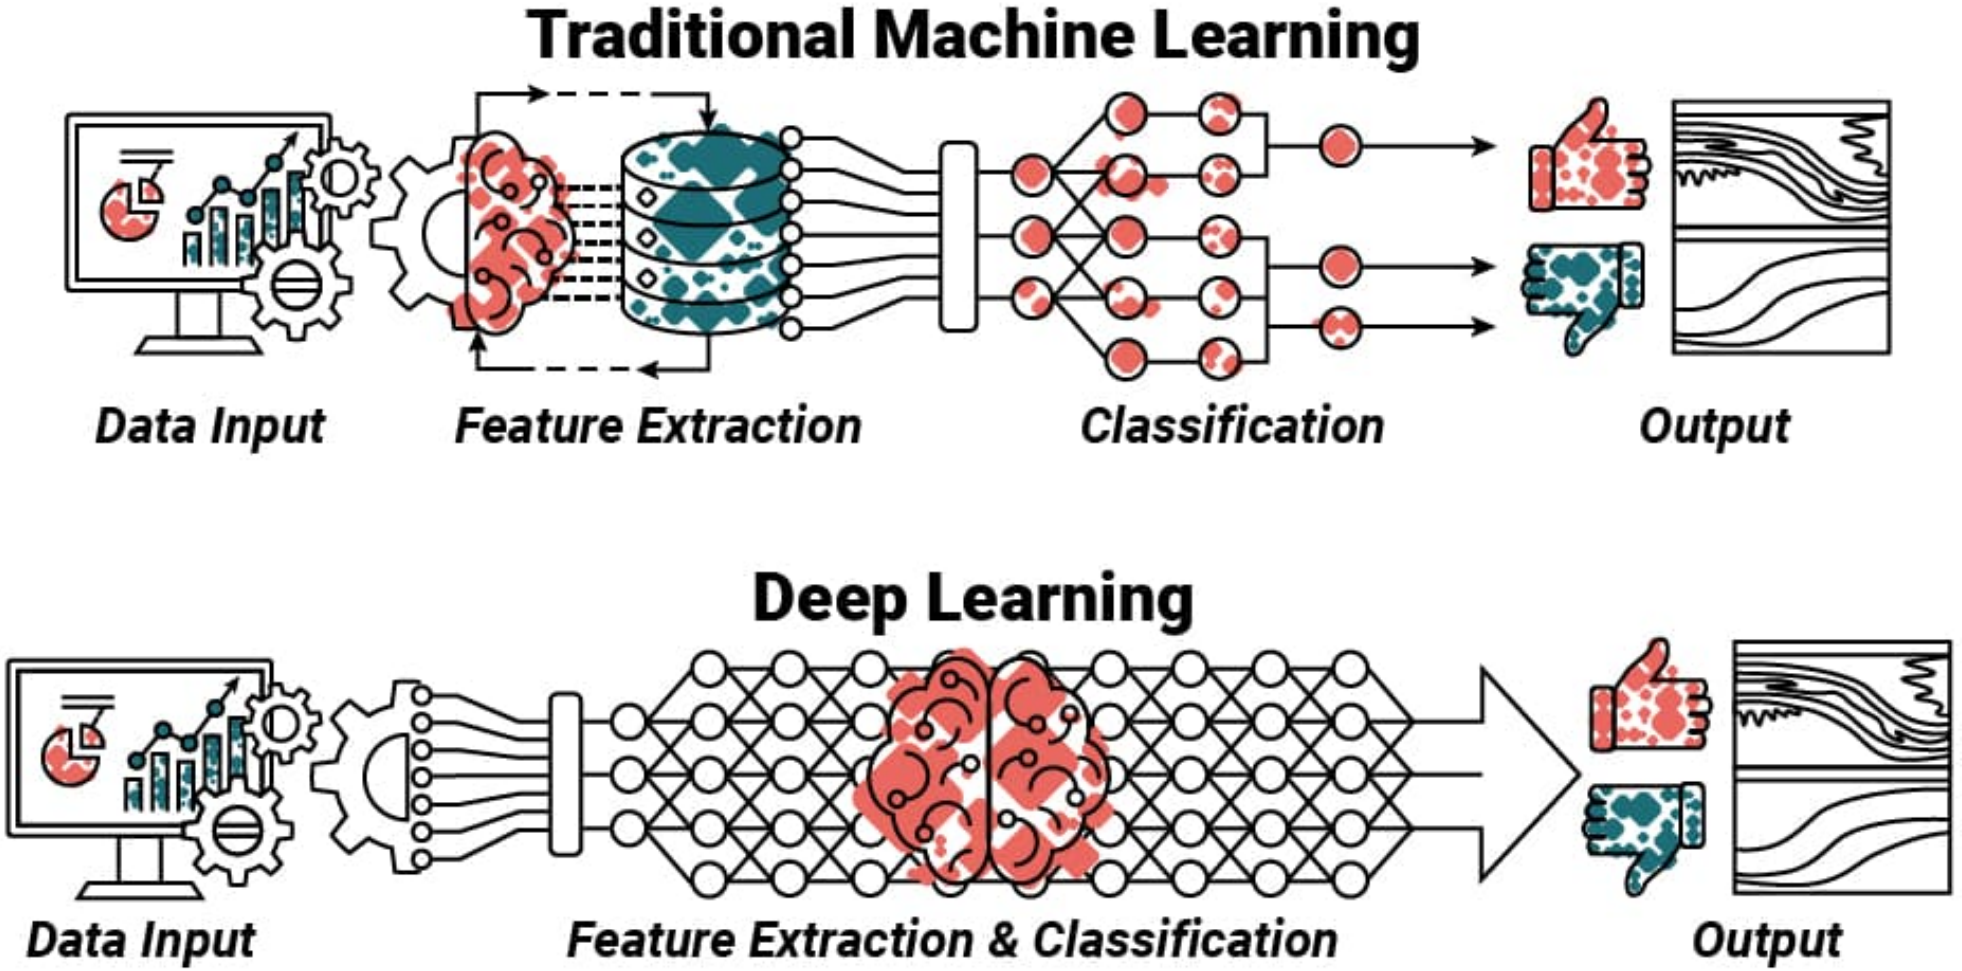
Traditional Machine Learning
Data Input
Feature Extraction
Classification
Output
Deep Learning
Data Input
Feature Extraction & Classification
Output

MACHINE LEARNING

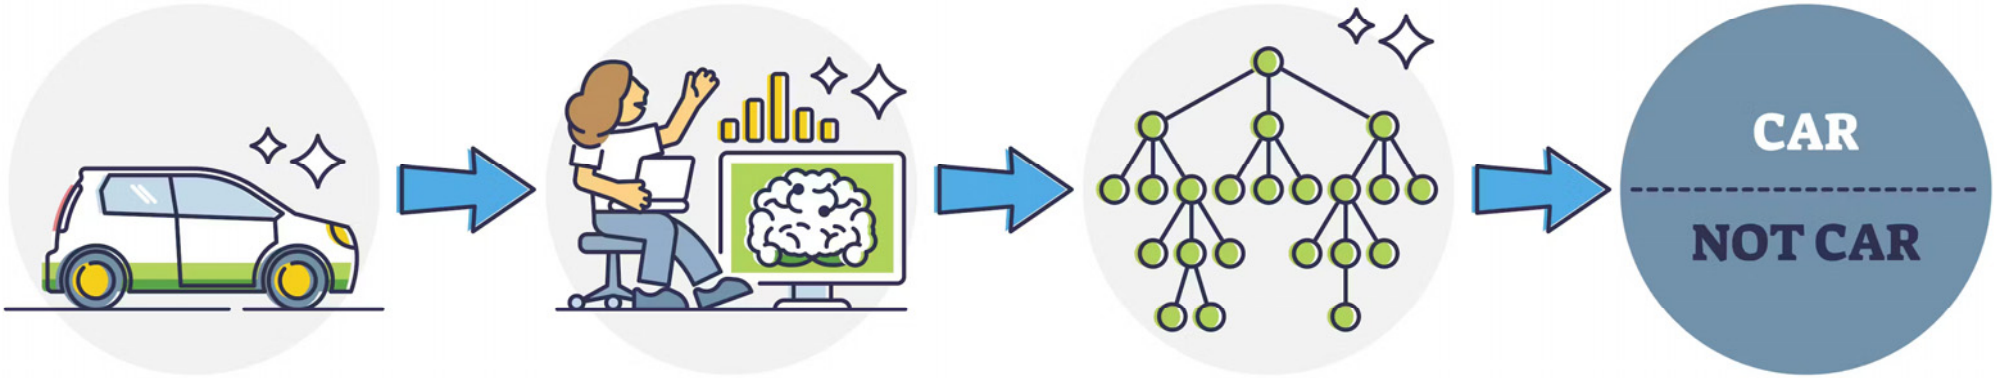

INPUT | FEATURE EXTRACTION | CLASSIFICATION | OUTPUT

DEEP LEARNING

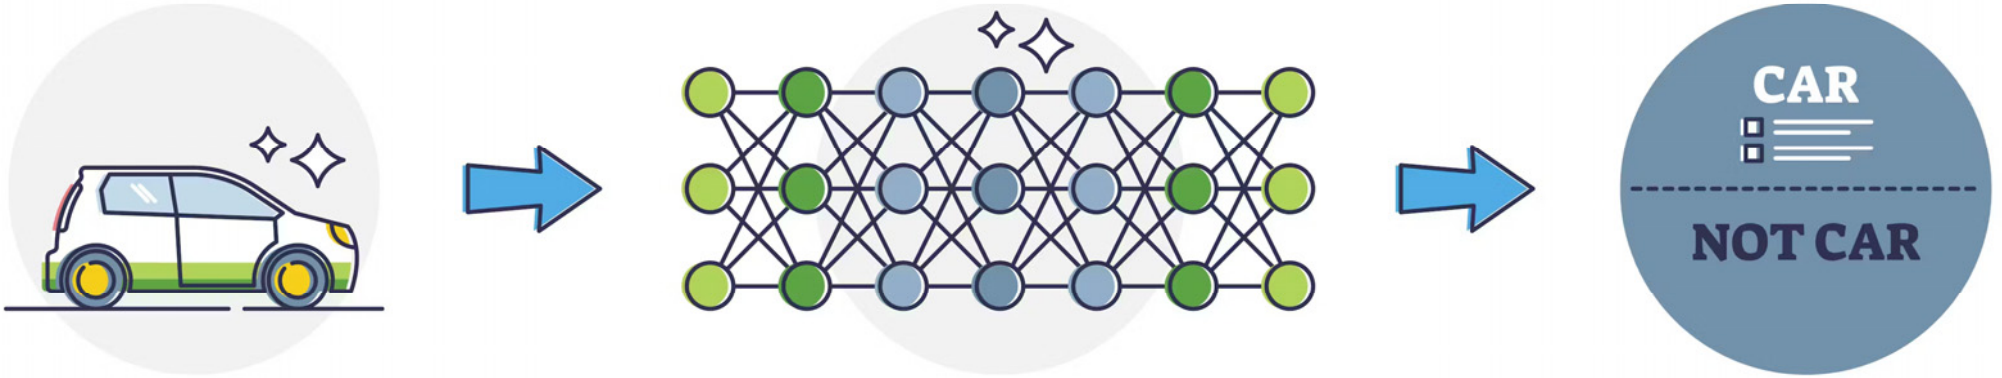

INPUT | FEATURE EXTRACTION + CLASSIFICATION | OUTPUT

MACHINE LEARNING

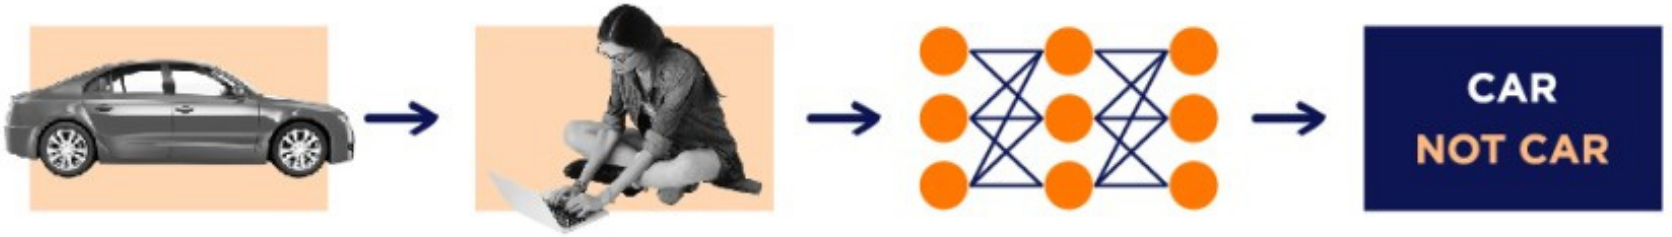

Input → Feature extraction → Classification → Output

DEEP LEARNING

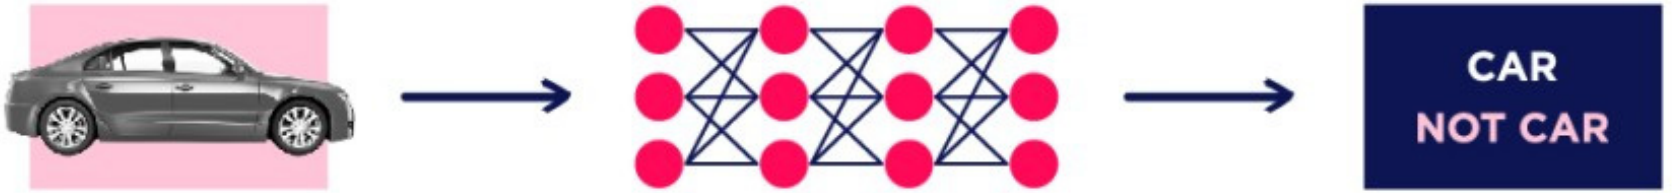

Input → Feature extraction + Classification → Output

ITERATORS

Source: www.iteratorshq.com

Machine Learning

Input → Feature extraction → Classification → Output

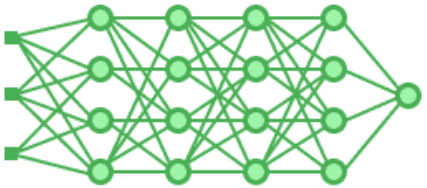

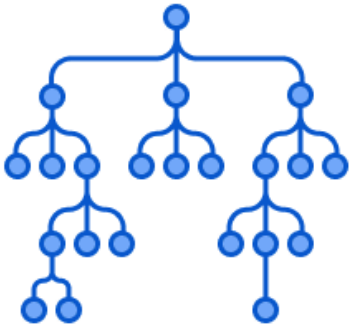

- Car
- Not a car

Deep Learning

Input → Feature extraction + Classification → Output

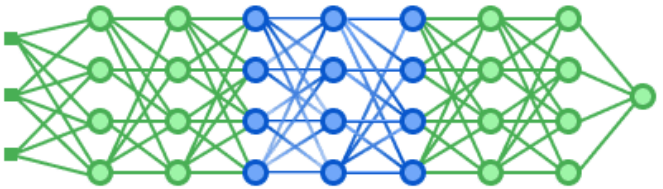

Hidden neurons

- Car
  - Color: Red
  - Made: Ford
  - Model: Mustang
- Not a car

mobidev

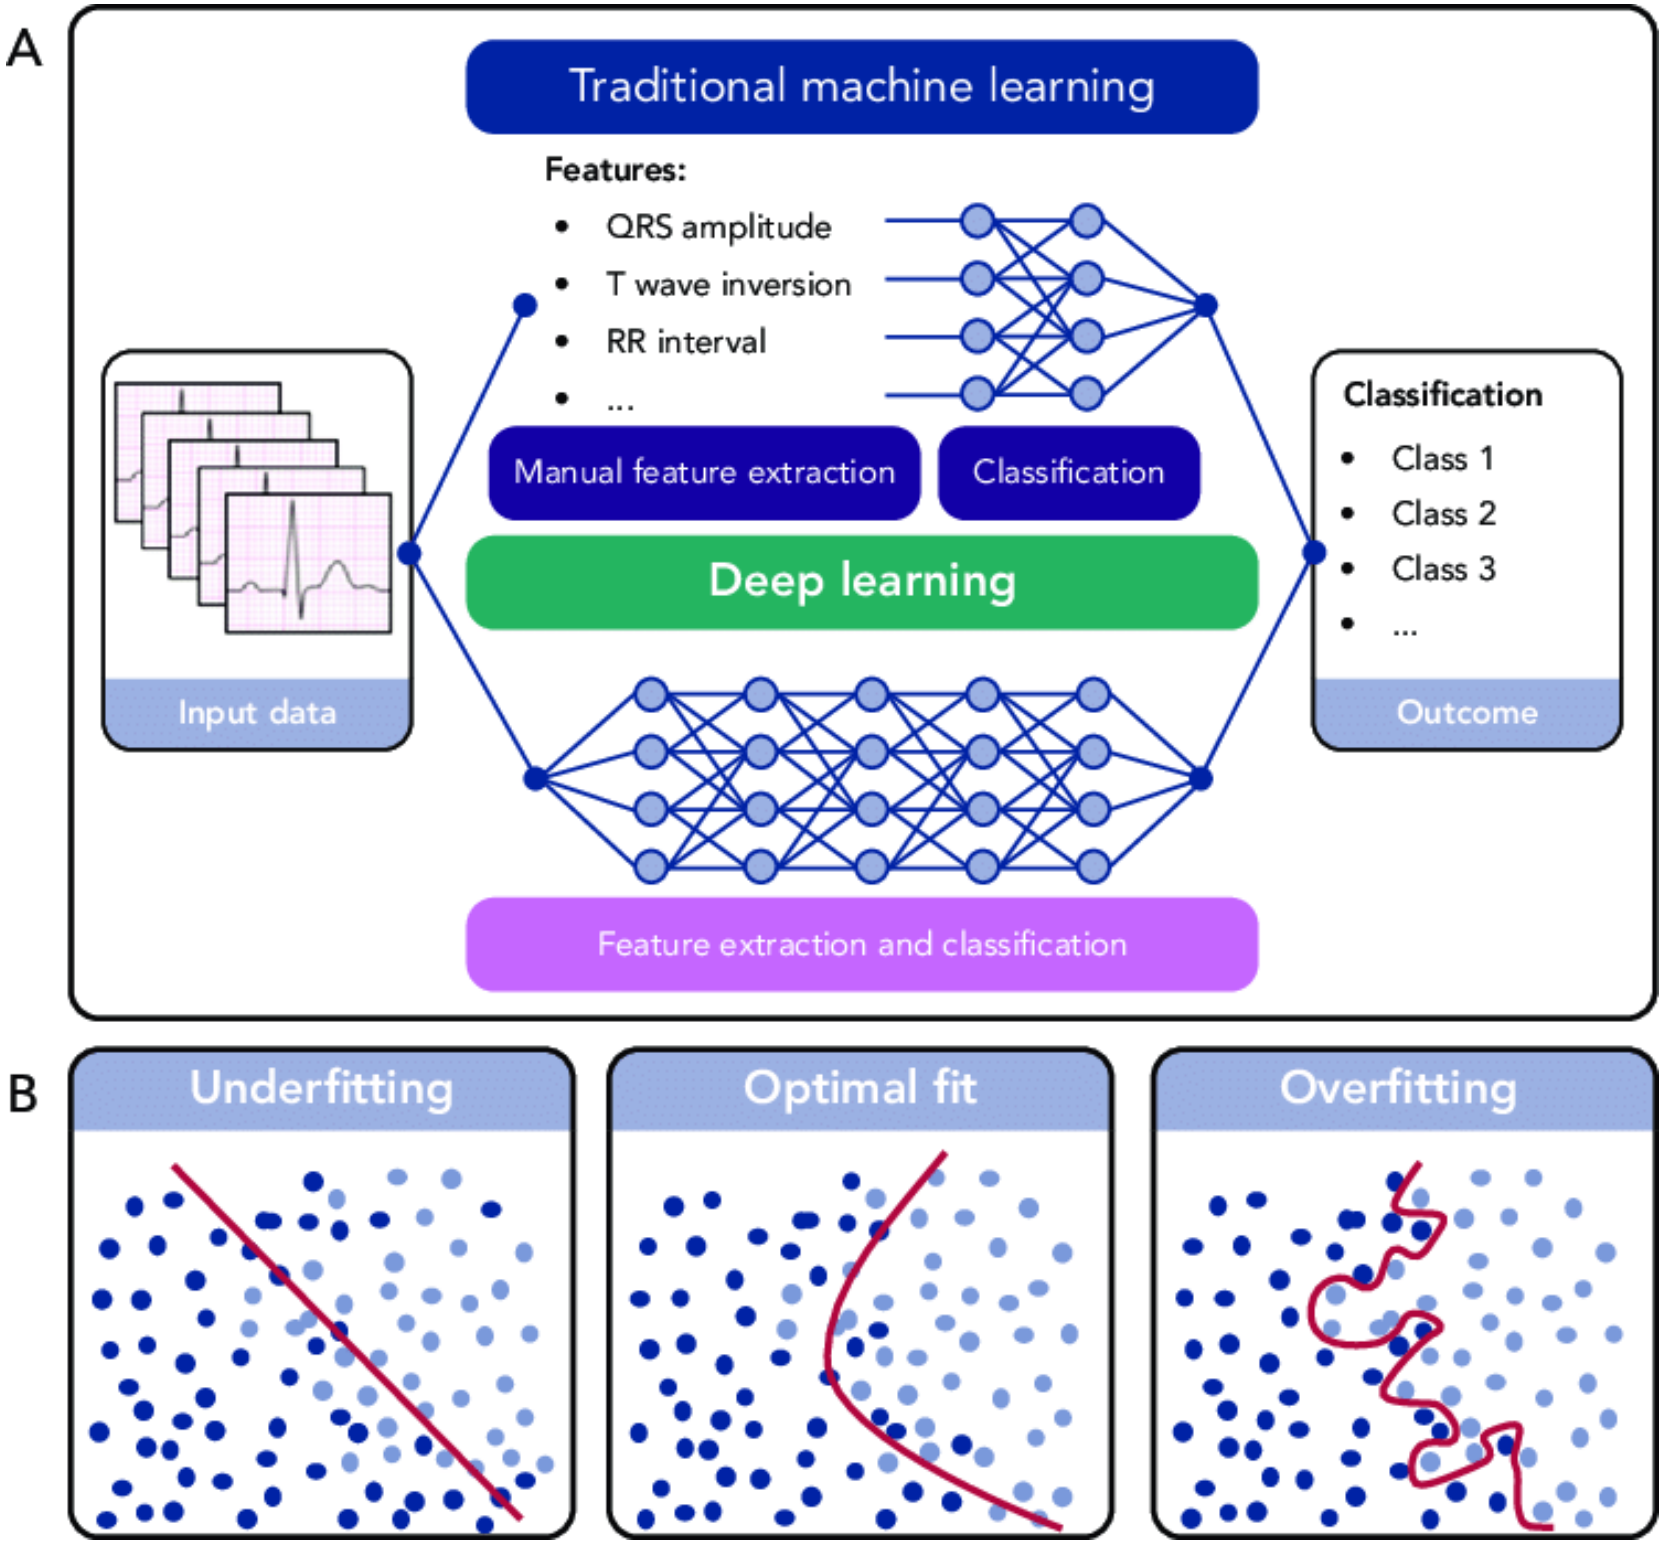
A
Traditional machine learning
Features:
QRS amplitude
T wave inversion
RR interval
...
Manual feature extraction
Classification
Deep learning
Input data
Classification
Class 1
Class 2
Class 3
...
Outcome
Feature extraction and classification
B
Underfitting
Optimal fit
Overfitting

Why Deep Learning?

Applications

Speech/Audio Processing

Computer Vision

Natural Language Processing

مروری بر یادگیری عمیق

۳

یادگیری عمیق: وضعیت کنونی

A brief History

A long time ago...

1958 Perceptron

1969
Perceptron criticized

1974 Backpropagation

awkward silence (AI Winter)

Convolution Neural Networks for Handwritten Recognition

1998

1995
SVM reigns

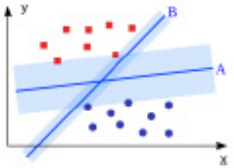

2006
Restricted Boltzmann Machine

Google Brain Project on 16k Cores

2012

2012
AlexNet wins ImageNet

IMAGENET

https://medium.com/global-silicon-valley/machine-learning-yesterday-today-tomorrow-3d3023c7b519#.wd0gd11rl

A brief History

The Big Bang aka "One net to rule them all"

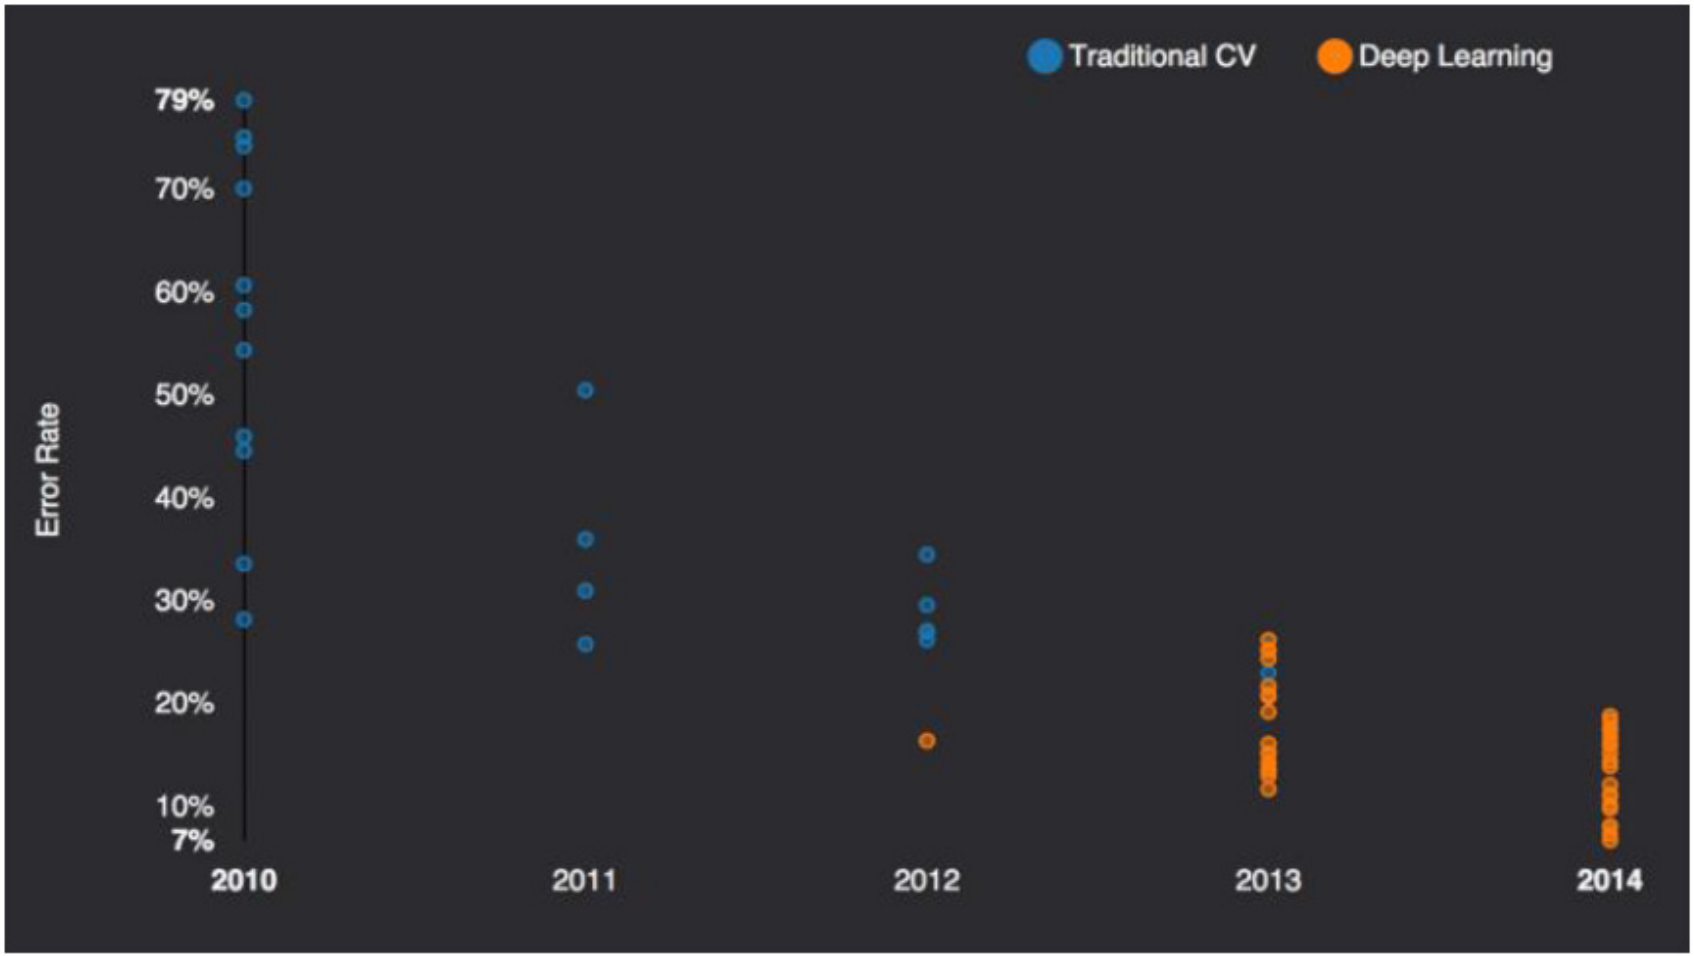

ImageNet: The "computer vision World Cup"

Source

https://www.quora.com/What-is-the-state-of-the-art-today-on-ImageNet-classification-In-other-words-has-anybody-beaten-Deep-Residual-Learning

A brief History

The Big Bang aka “One net to rule them all”

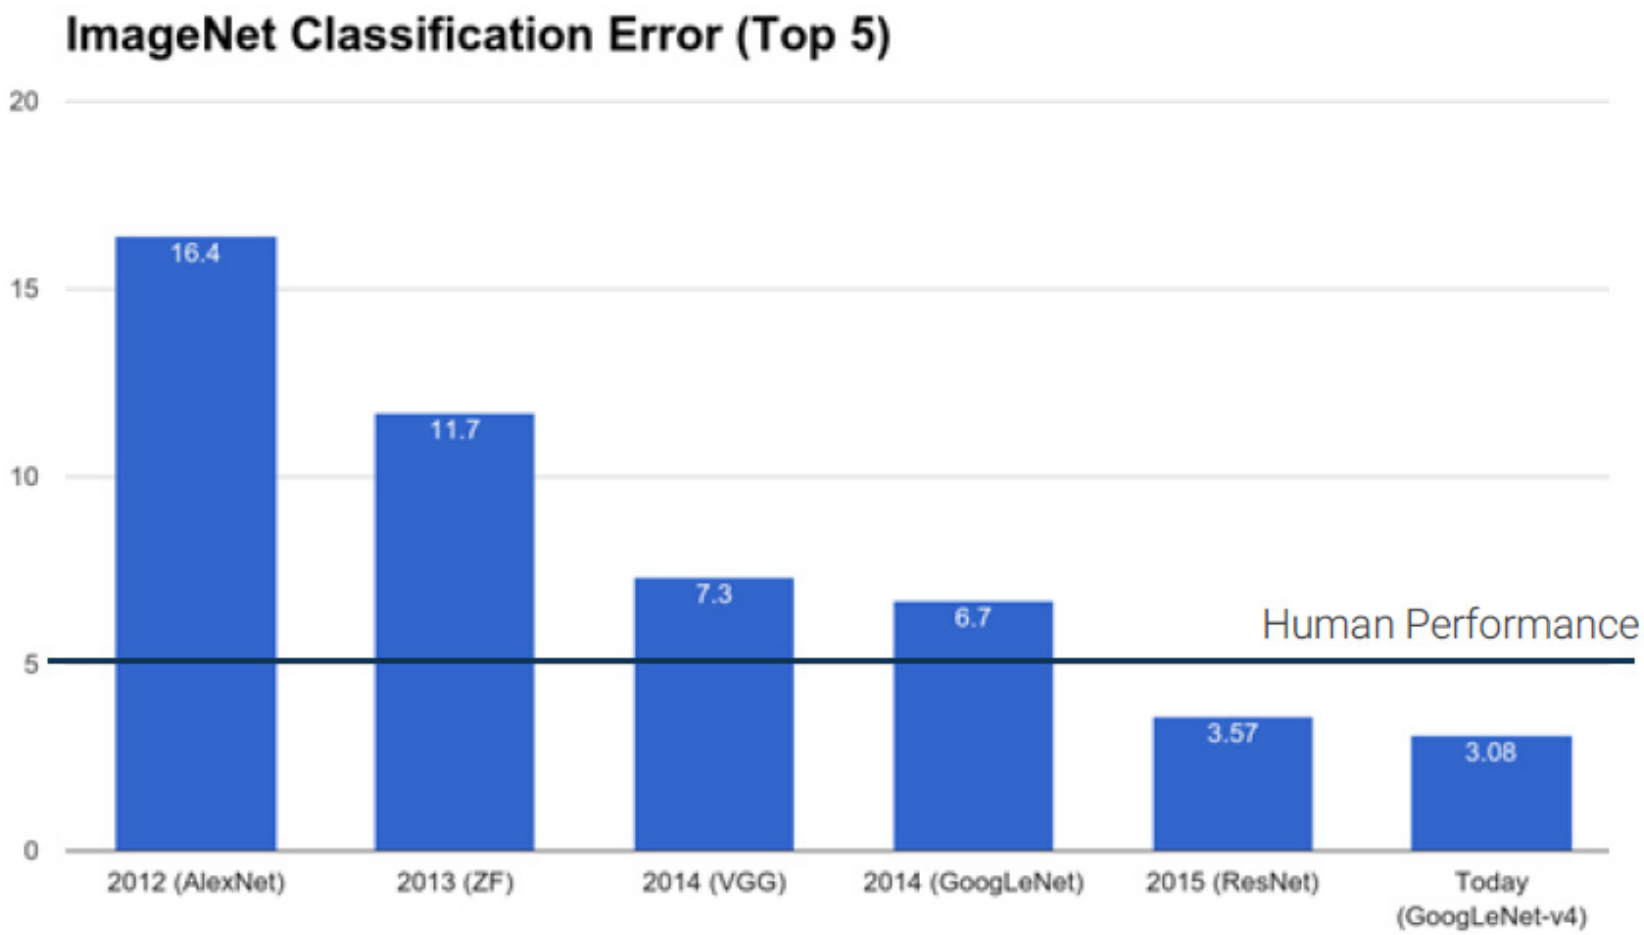

ImageNet: The “computer vision World Cup”

Source

https://arxiv.org/abs/1610.05256

A brief History

The Big Bang aka “One net to rule them all”

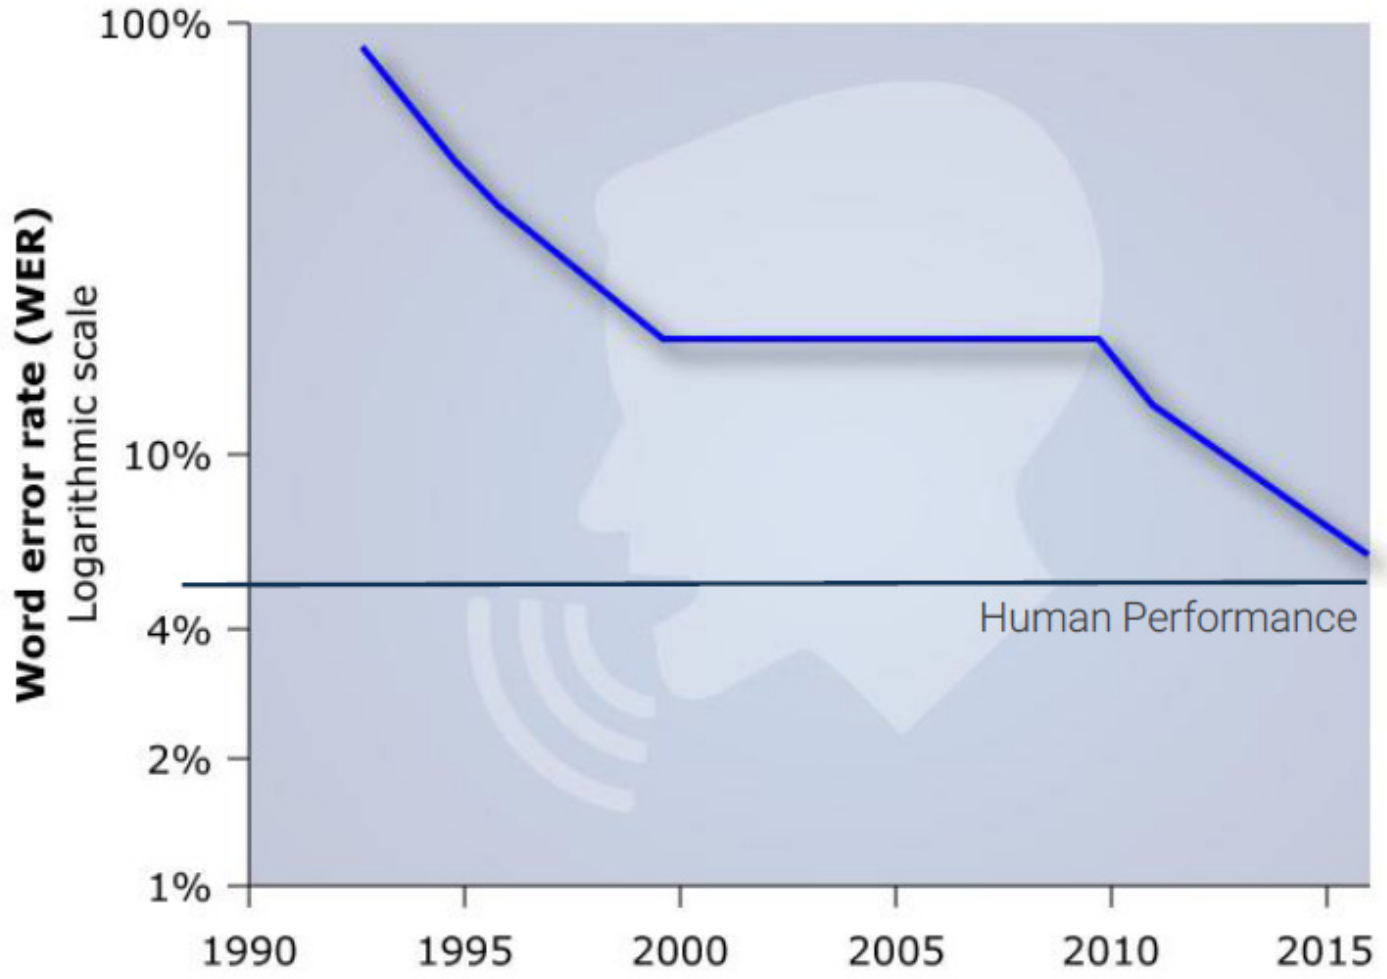

Deep Learning in Speech Recognition

Source

What changed?

Old wine in new bottles

Big Data
(Digitalization)

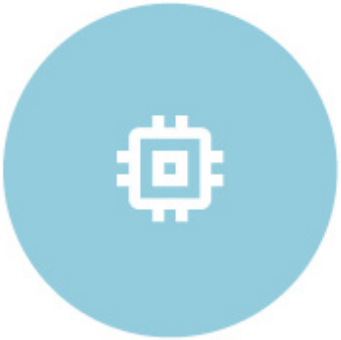

Computation
(Moore's Law, GPUs)

Algorithmic
Progress

Generative AI

V/S

Discriminative AI

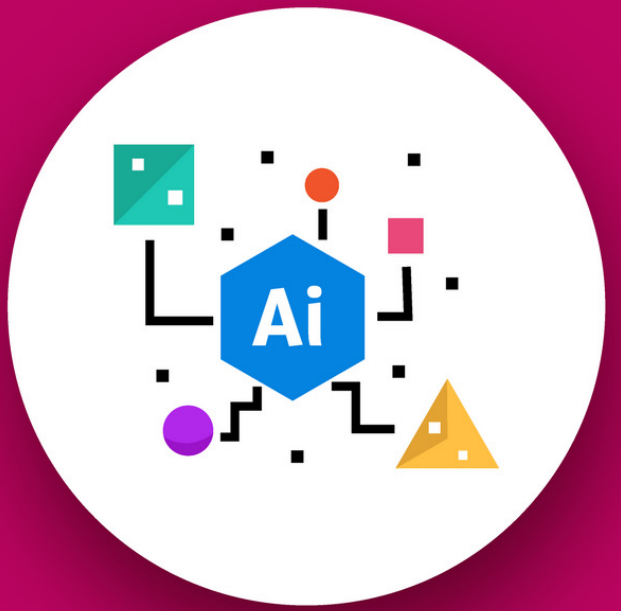

Discriminative

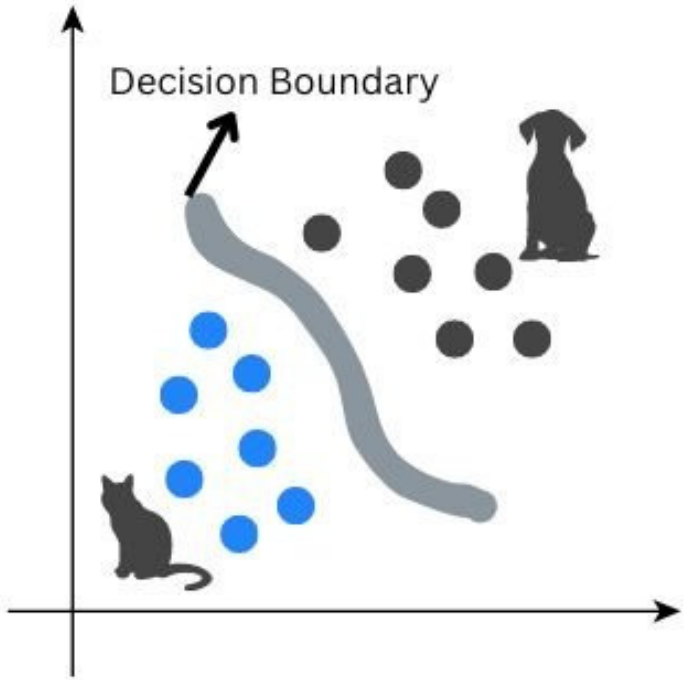

Classify or Label data point as cat or dog

Generative

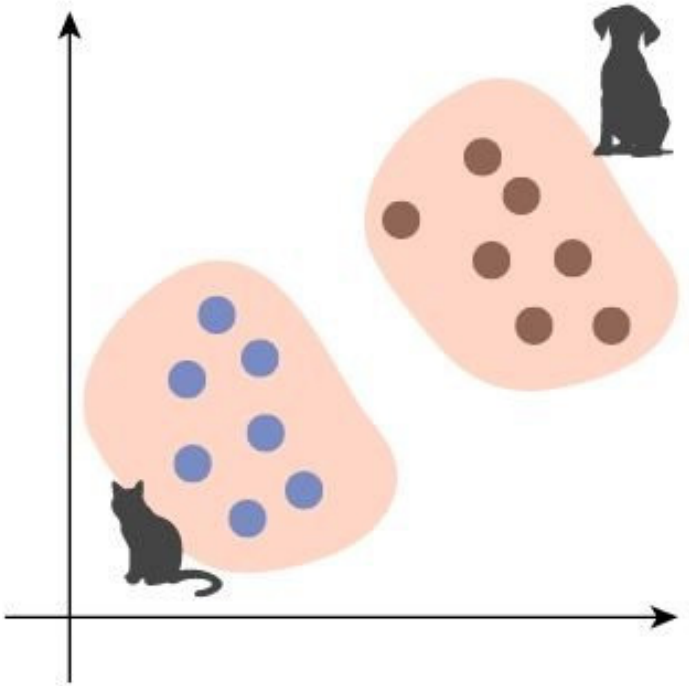

Produce a new data point that looks like cats or dogs

viso.ai

DISCRIMINATIVE AI VS GENERATIVE AI

Mention your slide subtitle here to provide additional context

FACTORS	Discriminative AI	Generative AI
Focus	Primarily focuses on boundaries and differences between classes	Focuses on providing insights by going through the underlying data distribution & structure.
Mathematical Intuition	By maximizing the joint probability of P(X, Y)	It maximizes the conditional probability P(Y|X)
Objective	The conditional probability P(y)	The joint probability distribution P(x, y) or likelihood of P(x)
Computational Cost	With less data, it can perform well with classification tasks	It requires more data to better understand the data and extract meaningful insights.
Parameter Count	Since it only focuses on the decision boundary, it takes only a few parameters.	It has a larger parameter count as it takes the entire data distribution.
Sample Generation	Excels in image classification, sentiment analysis speech recognition, and object detection.	It can generate new data samples based on the trained data.
Application	Primarily used for classifying existing data such as image classification and spam filtering.	It is used for data generation filling up missing data, density estimation and many more.

What Is Generative AI?

1 **Generative AI**	AI techniques that learn from data about existing artifacts and use this to generate new artifacts.
2 **Foundation Models**	Large machine learning models that are trained on a broad set of unlabeled data and are fine-tuned to a wide range of applications.
3 **Large Language Models (LLM)**	AI that is trained on vast amounts of text to interpret and generate human-like textual output.
4 **ChatGPT**	An OpenAI service that incorporates a conversational chatbot with LLM to create content.

Source: Gartner

Gartner.

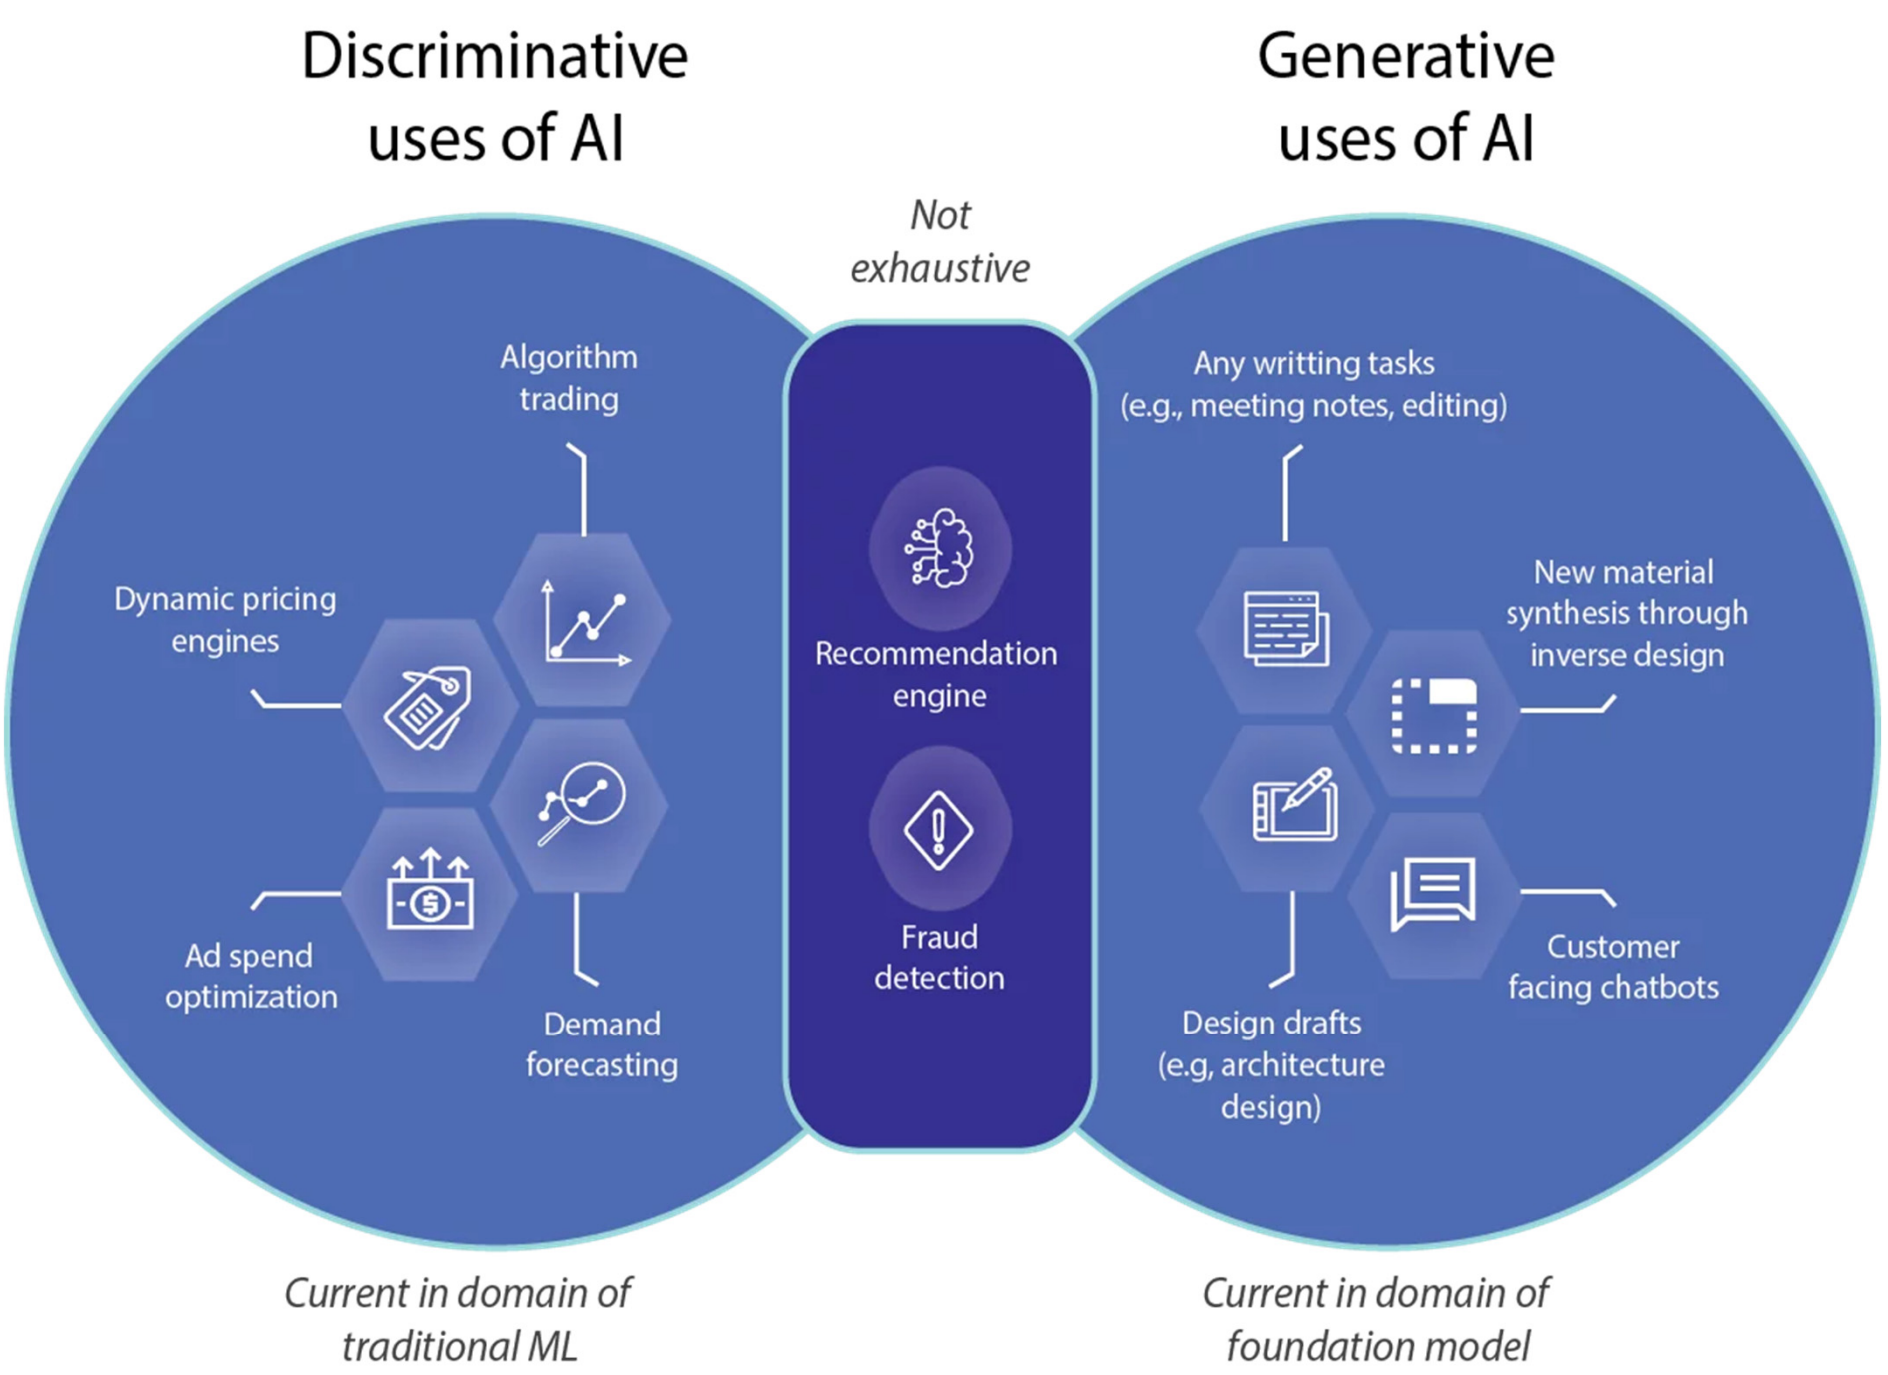
Discriminative uses of AI
Generative uses of AI
Not exhaustive
Algorithm trading
Dynamic pricing engines
Ad spend optimization
Demand forecasting
Recommendation engine
Fraud detection
Any writting tasks (e.g., meeting notes, editing)
New material synthesis through inverse design
Customer facing chatbots
Design drafts (e.g, architecture design)
Current in domain of traditional ML
Current in domain of foundation model

Generative AI Reshaping the future

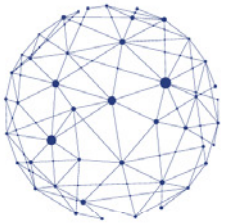

external

AI foundation models of global knowledge

internal

AI foundation models of enterprise knowledge

Other AI models and frameworks, including near-AGI

Types of AI enablement

- Generative Info collection
- Generative research
- Generative insights
- Generative automation
- Generative innovation
- Generative decisions

AI-Enabled knowledge worker

- Safety
- Ethics
- Compliance

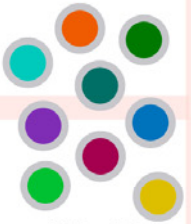

AI-Enabled Work Apps

How Artificial Intelligence will reshape the workplace and employee experience

Strategic Work

- Pursuing new markets
- Enhancing products and services
- Designing CX/EX experiences
- Imagining new products + processes
- Creating + using unique knowledge
- Communicating + collaborating
- Making better decisions, faster

Tactical Work

- Optimizing business operations
- Business administration
- Gathering needed information

datasciencedojo
data science for everyone


Text-to-Text
Image-to-Text
Foundation Models
Text-to-Video
Text-to-Image

A New Era of AI: Foundation Models

Step function improvements over legacy AI technologies

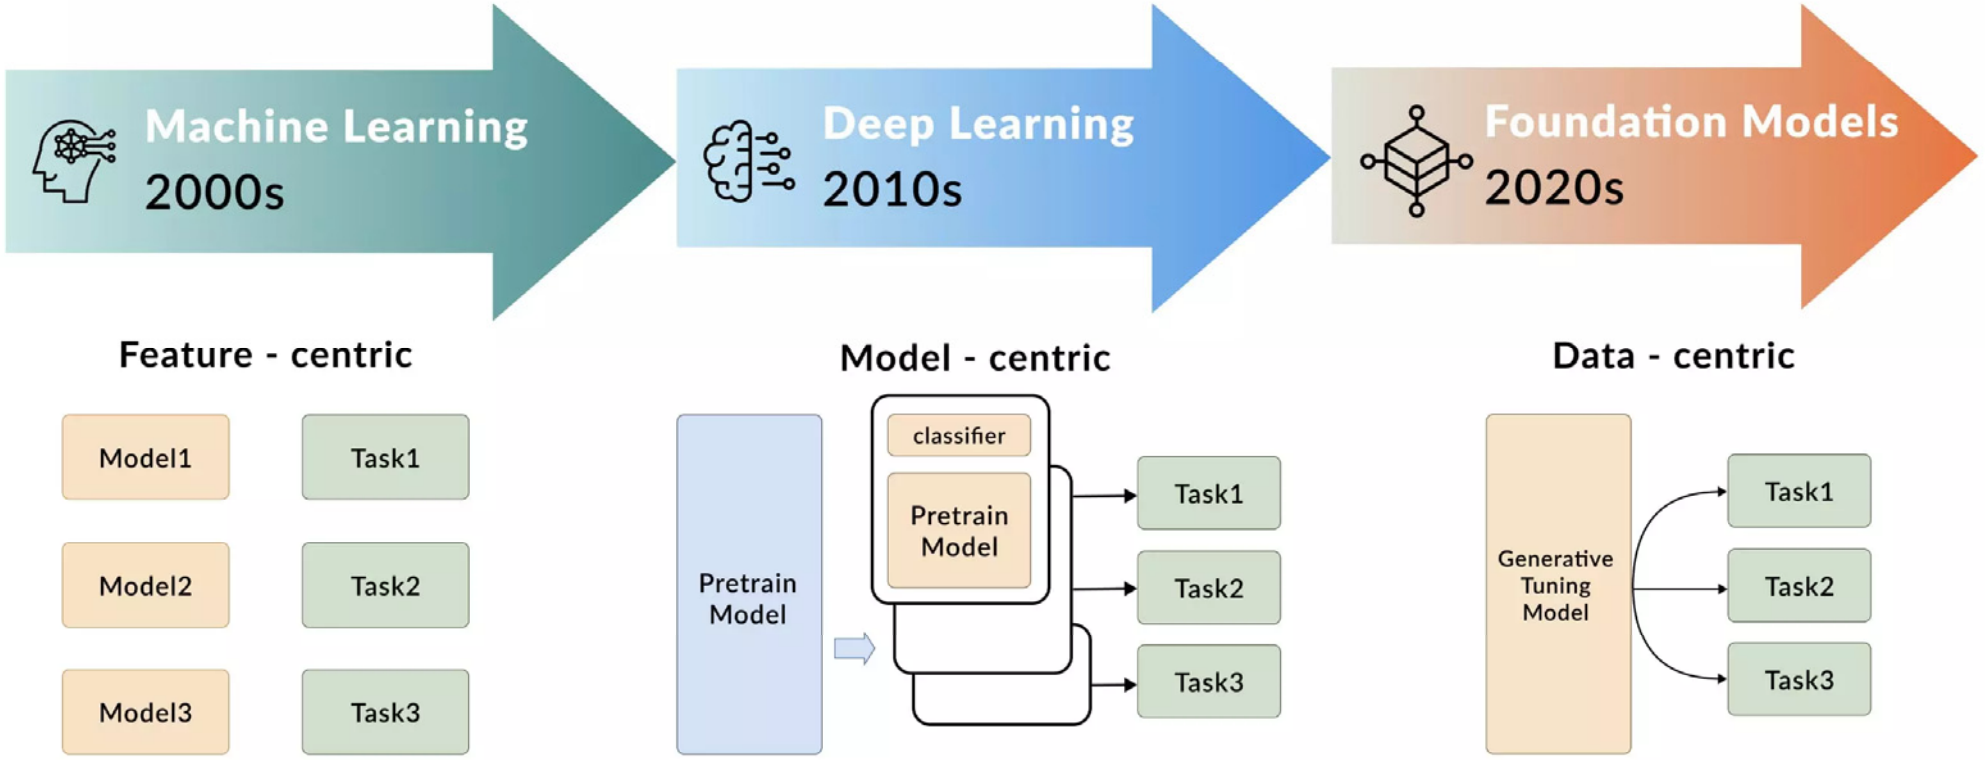

DATAFOREST

Traditional ML

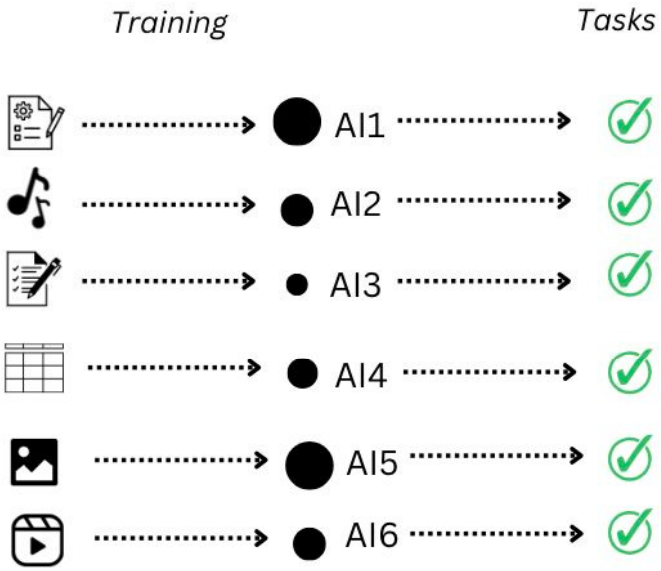

- Individual siloed models
- Require task-specific training
- Lots of human supervised training

Foundation Models

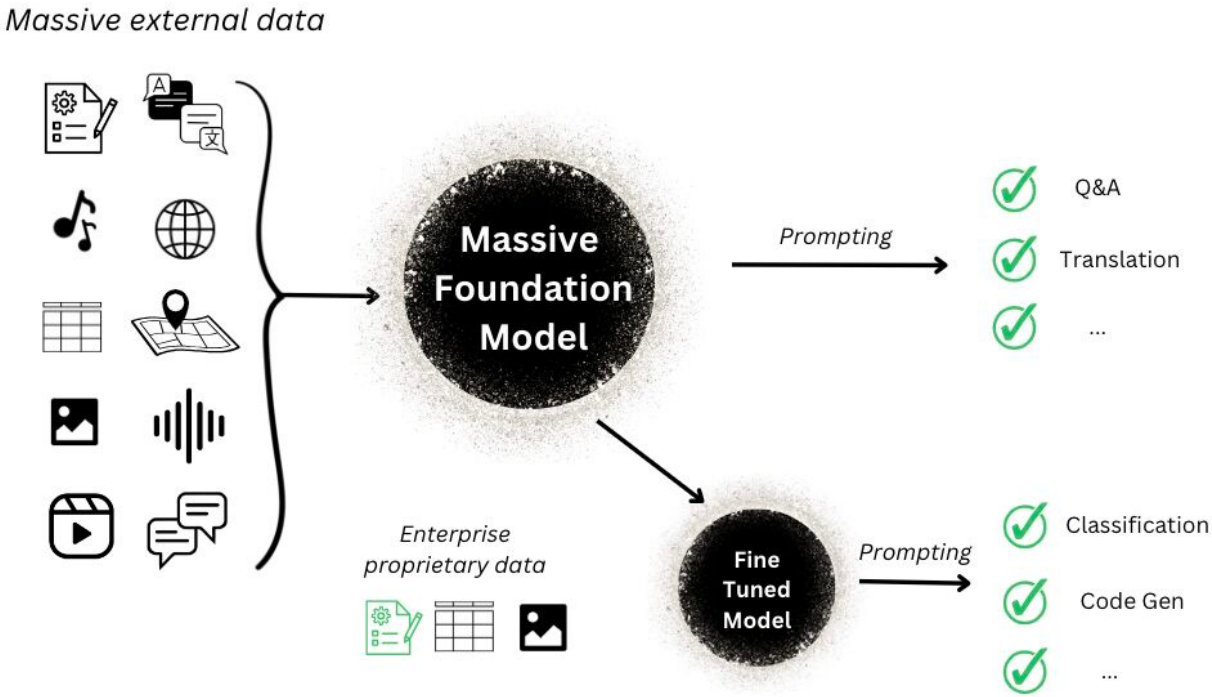

- Massive multi-tasking model
- Adaptable with little or no training
- Pre-trained unsupervised learning

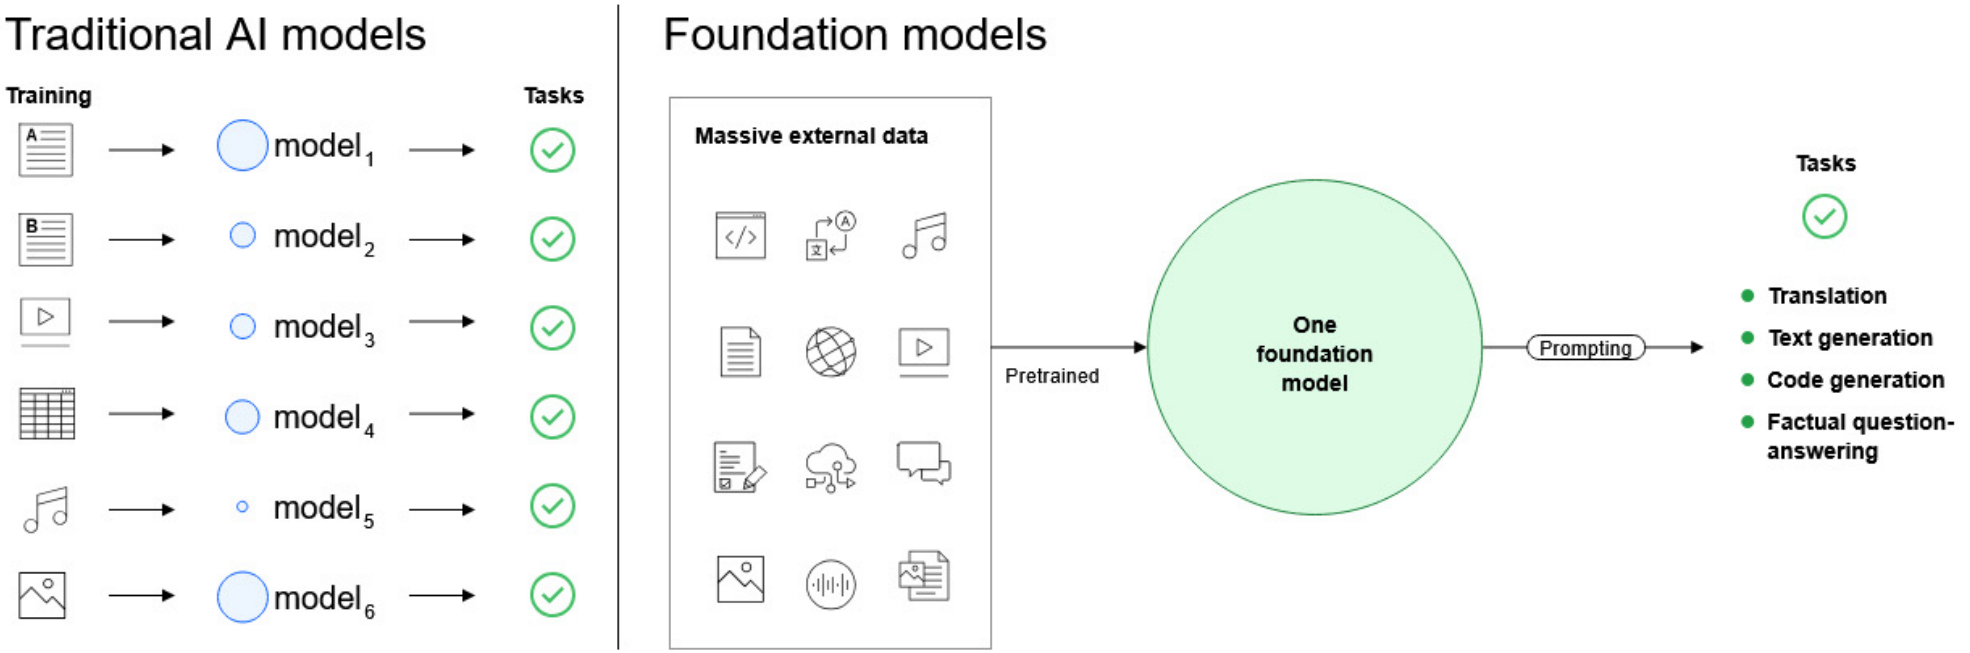
Traditional AI models
Training
Tasks
model 1
model 2
model 3
model 4
model 5
model 6
Foundation models
Massive external data
Pretrained
One foundation model
Prompting
Tasks
Translation
Text generation
Code generation
Factual question-answering

FOUNDATION MODEL

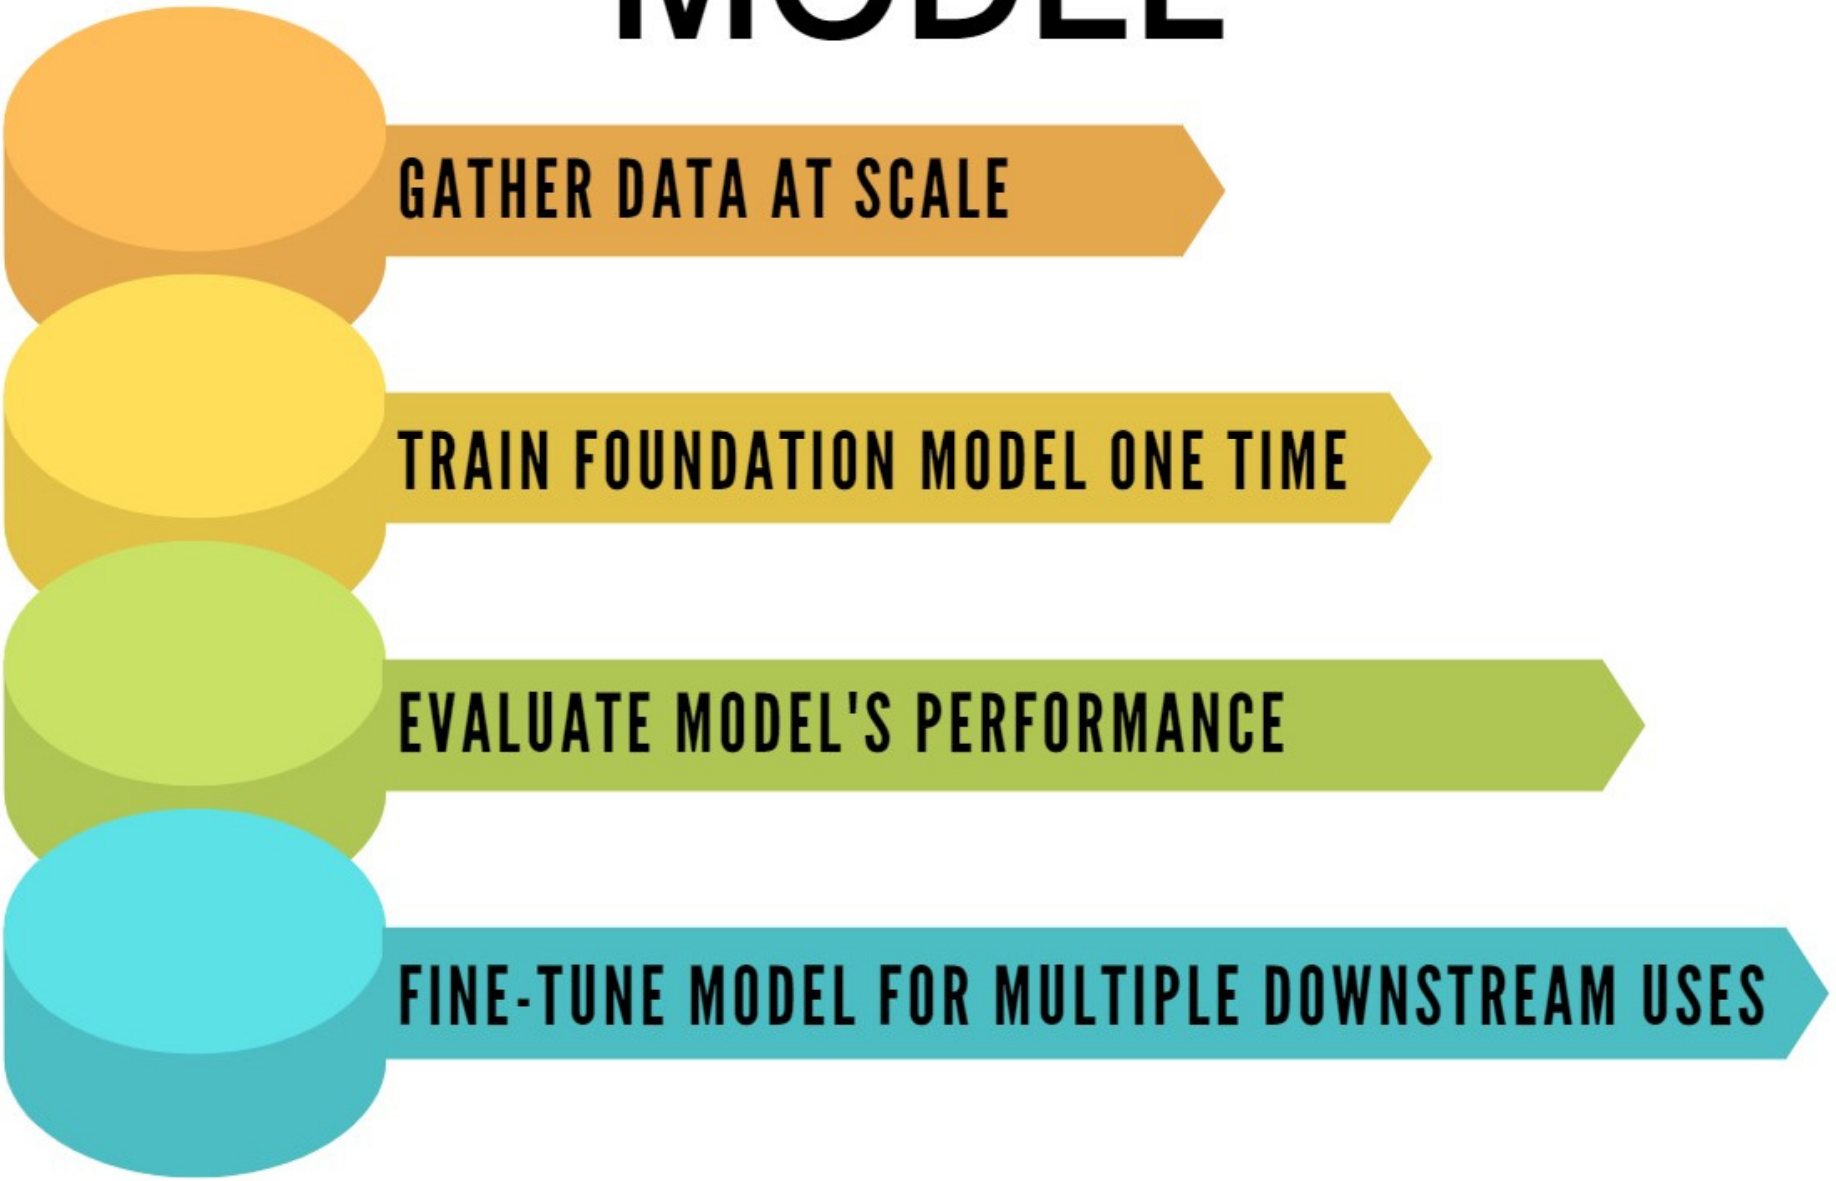

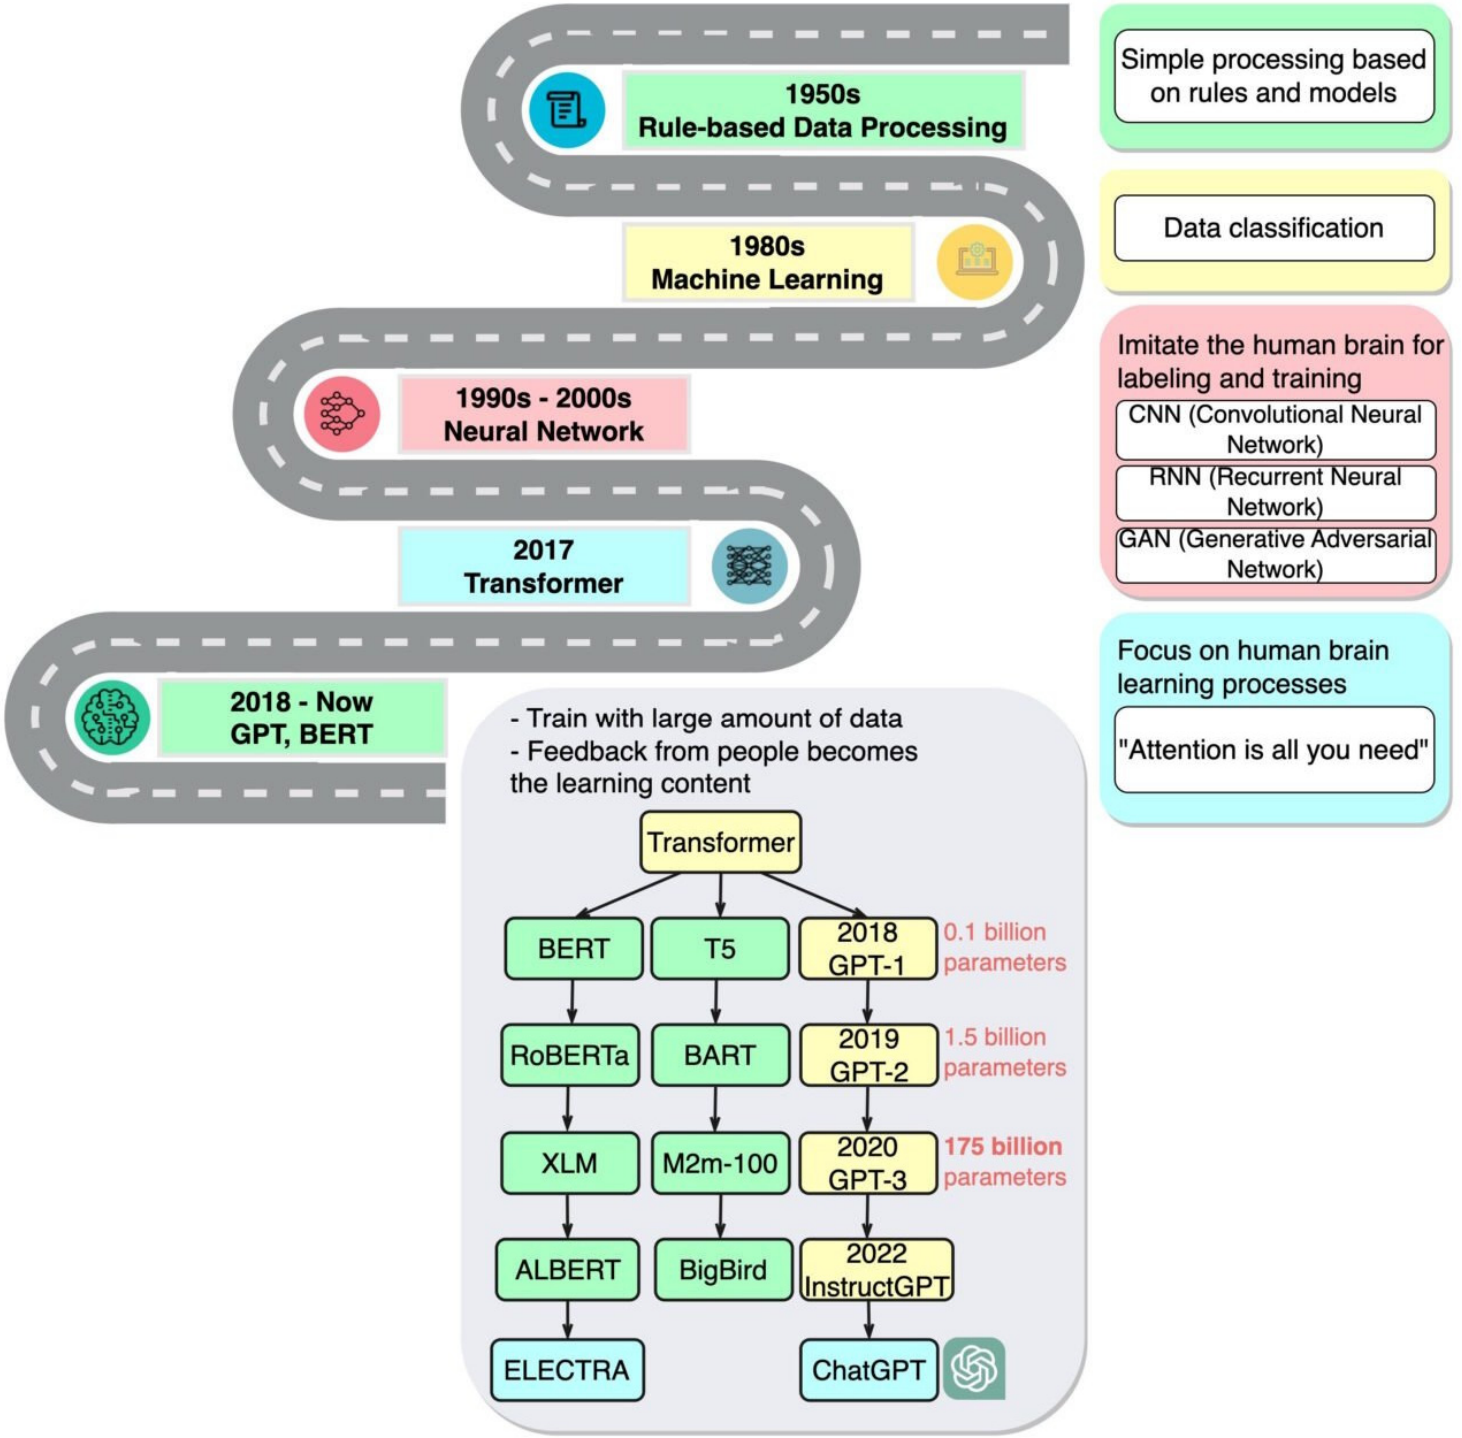
1950s
Rule-based Data Processing
Simple processing based on rules and models
1980s
Machine Learning
Data classification
1990s - 2000s
Neural Network
Imitate the human brain for labeling and training
CNN (Convolutional Neural Network)
RNN (Recurrent Neural Network)
GAN (Generative Adversarial Network)
2017
Transformer
Focus on human brain learning processes
"Attention is all you need"
2018 - Now
GPT, BERT
- Train with large amount of data
- Feedback from people becomes the learning content
Transformer
BERT
T5
2018 GPT-1
0.1 billion parameters
RoBERTa
BART
2019 GPT-2
1.5 billion parameters
XLM
M2m-100
2020 GPT-3
175 billion parameters
ALBERT
BigBird
2022 InstructGPT
ELECTRA
ChatGPT

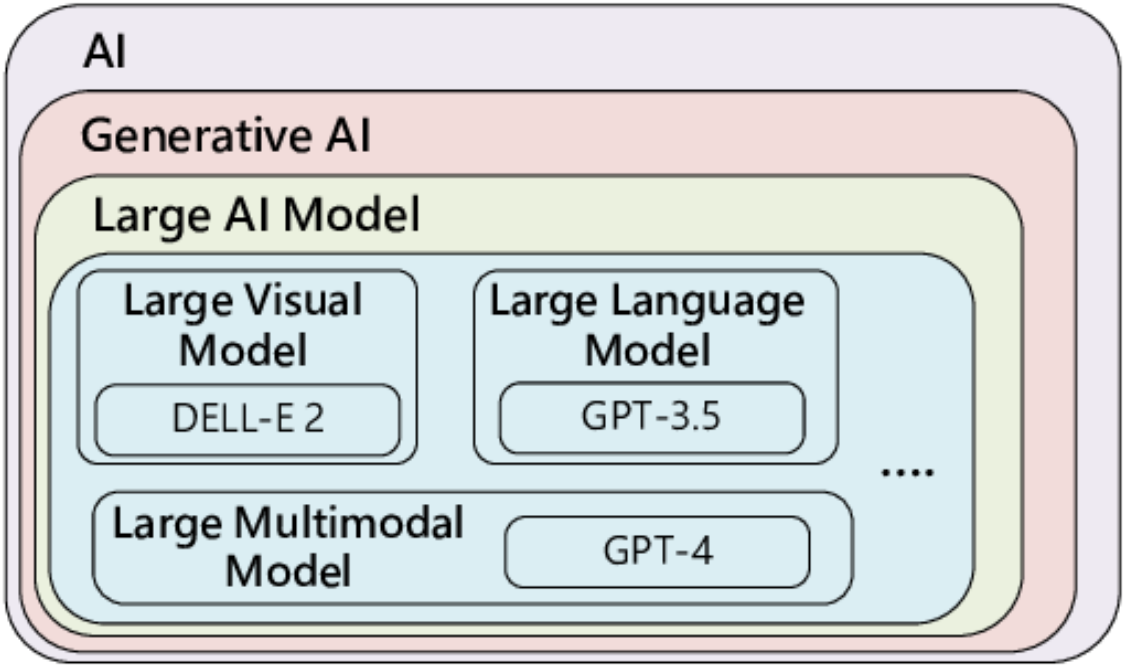
AI
Generative AI
Large AI Model
Large Visual Model
DELL-E 2
Large Language Model
GPT-3.5
....
Large Multimodal Model
GPT-4

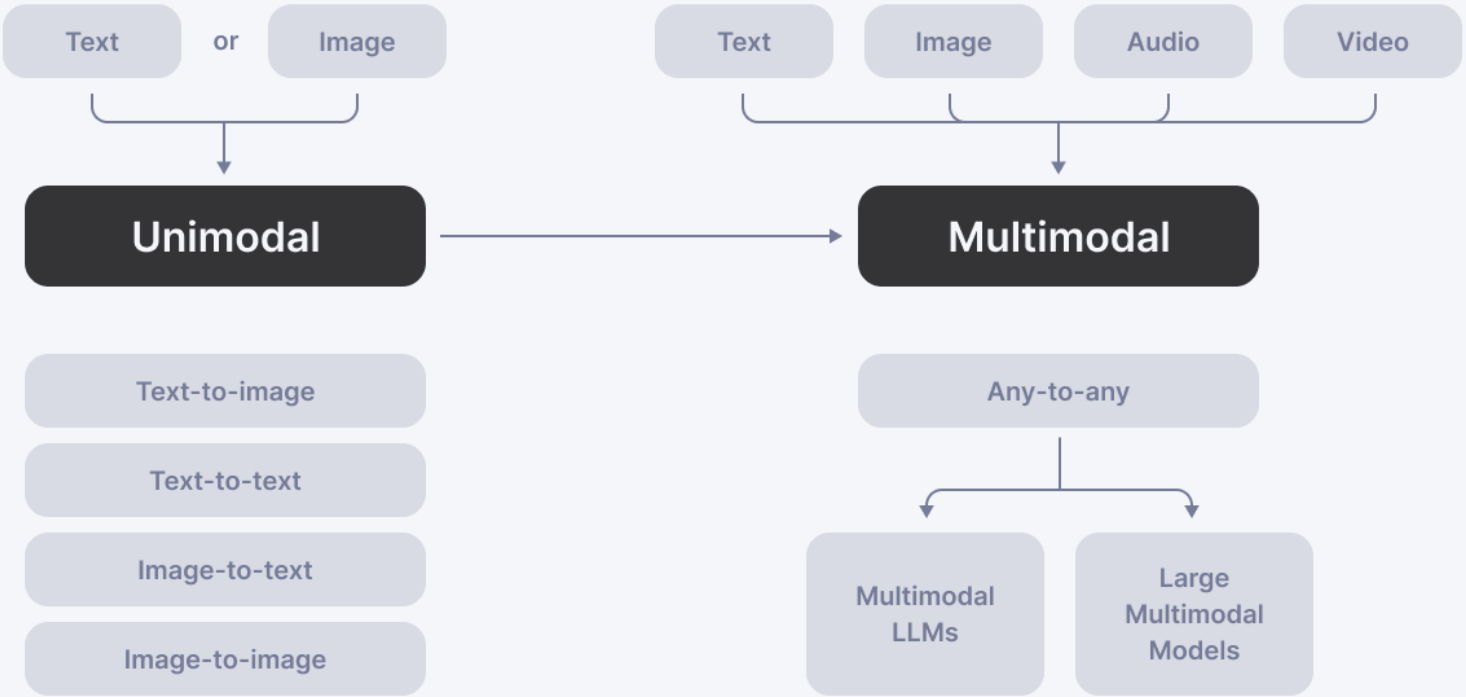
Text
or
Image
Text
Image
Audio
Video
Unimodal
Multimodal
Text-to-image
Text-to-text
Image-to-text
Image-to-image
Any-to-any
Multimodal LLMs
Large Multimodal Models

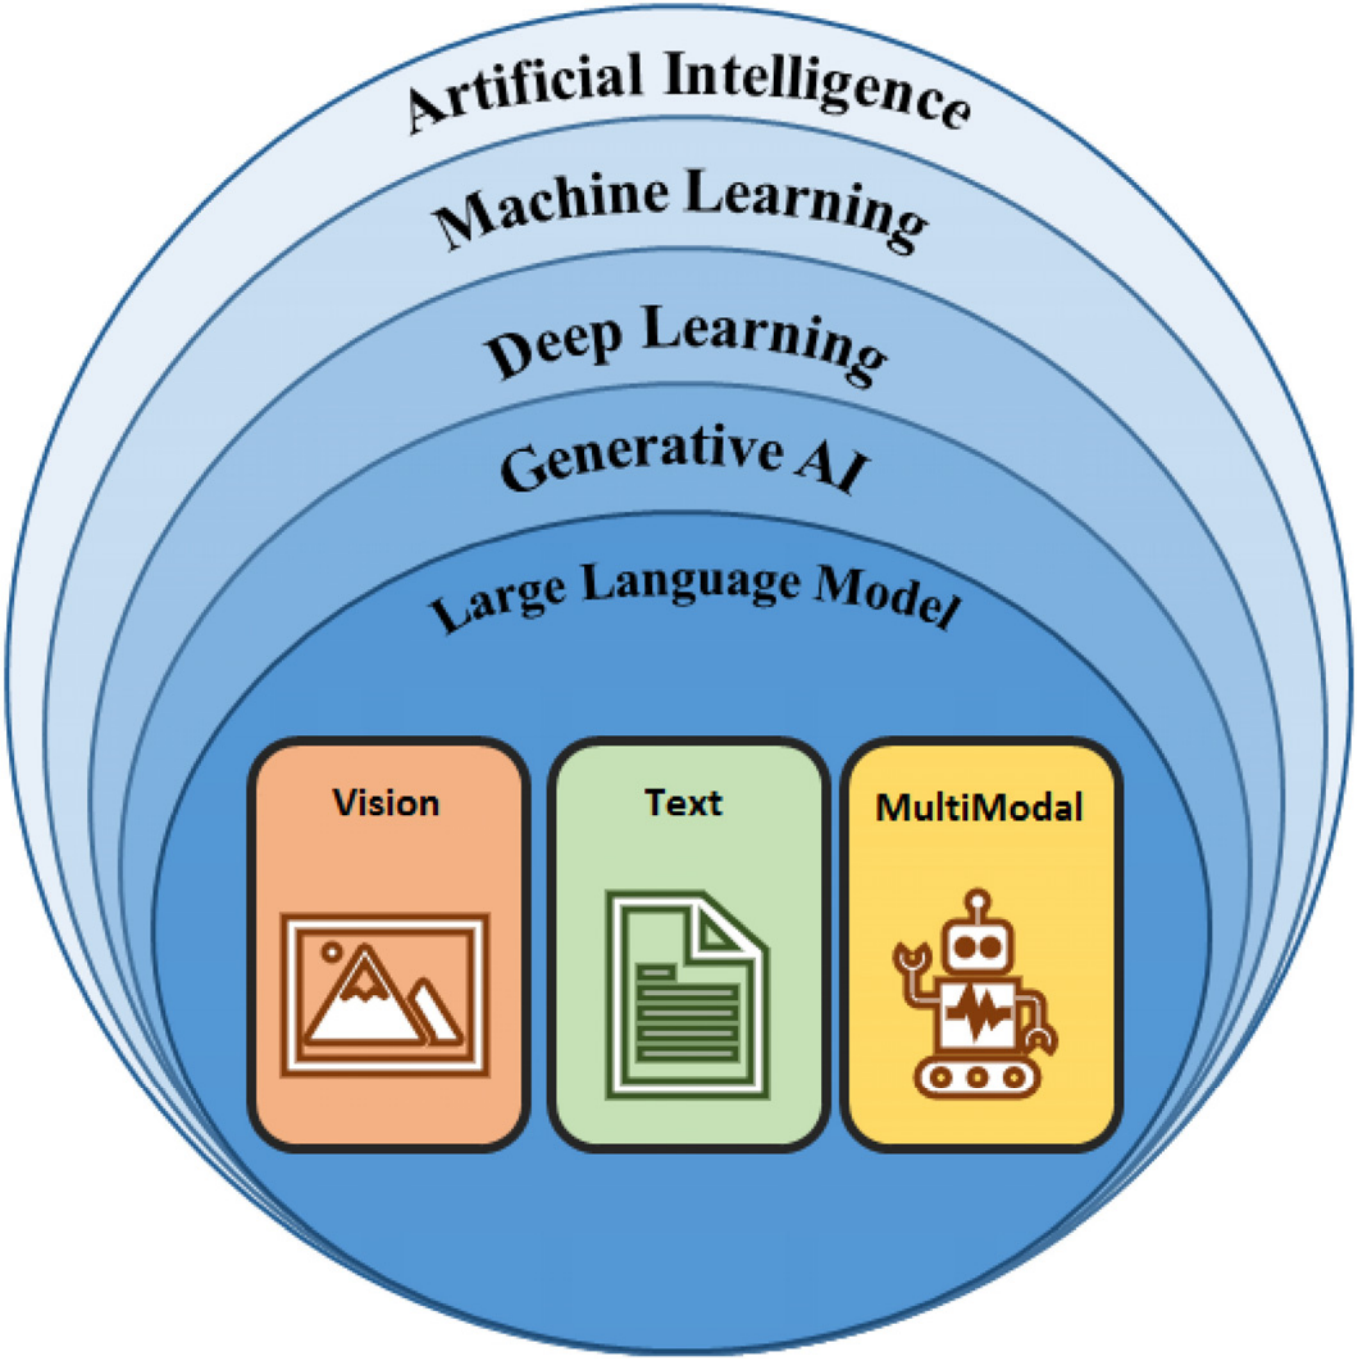
Artificial Intelligence
Machine Learning
Deep Learning
Generative AI
Large Language Model
Vision
Text
MultiModal

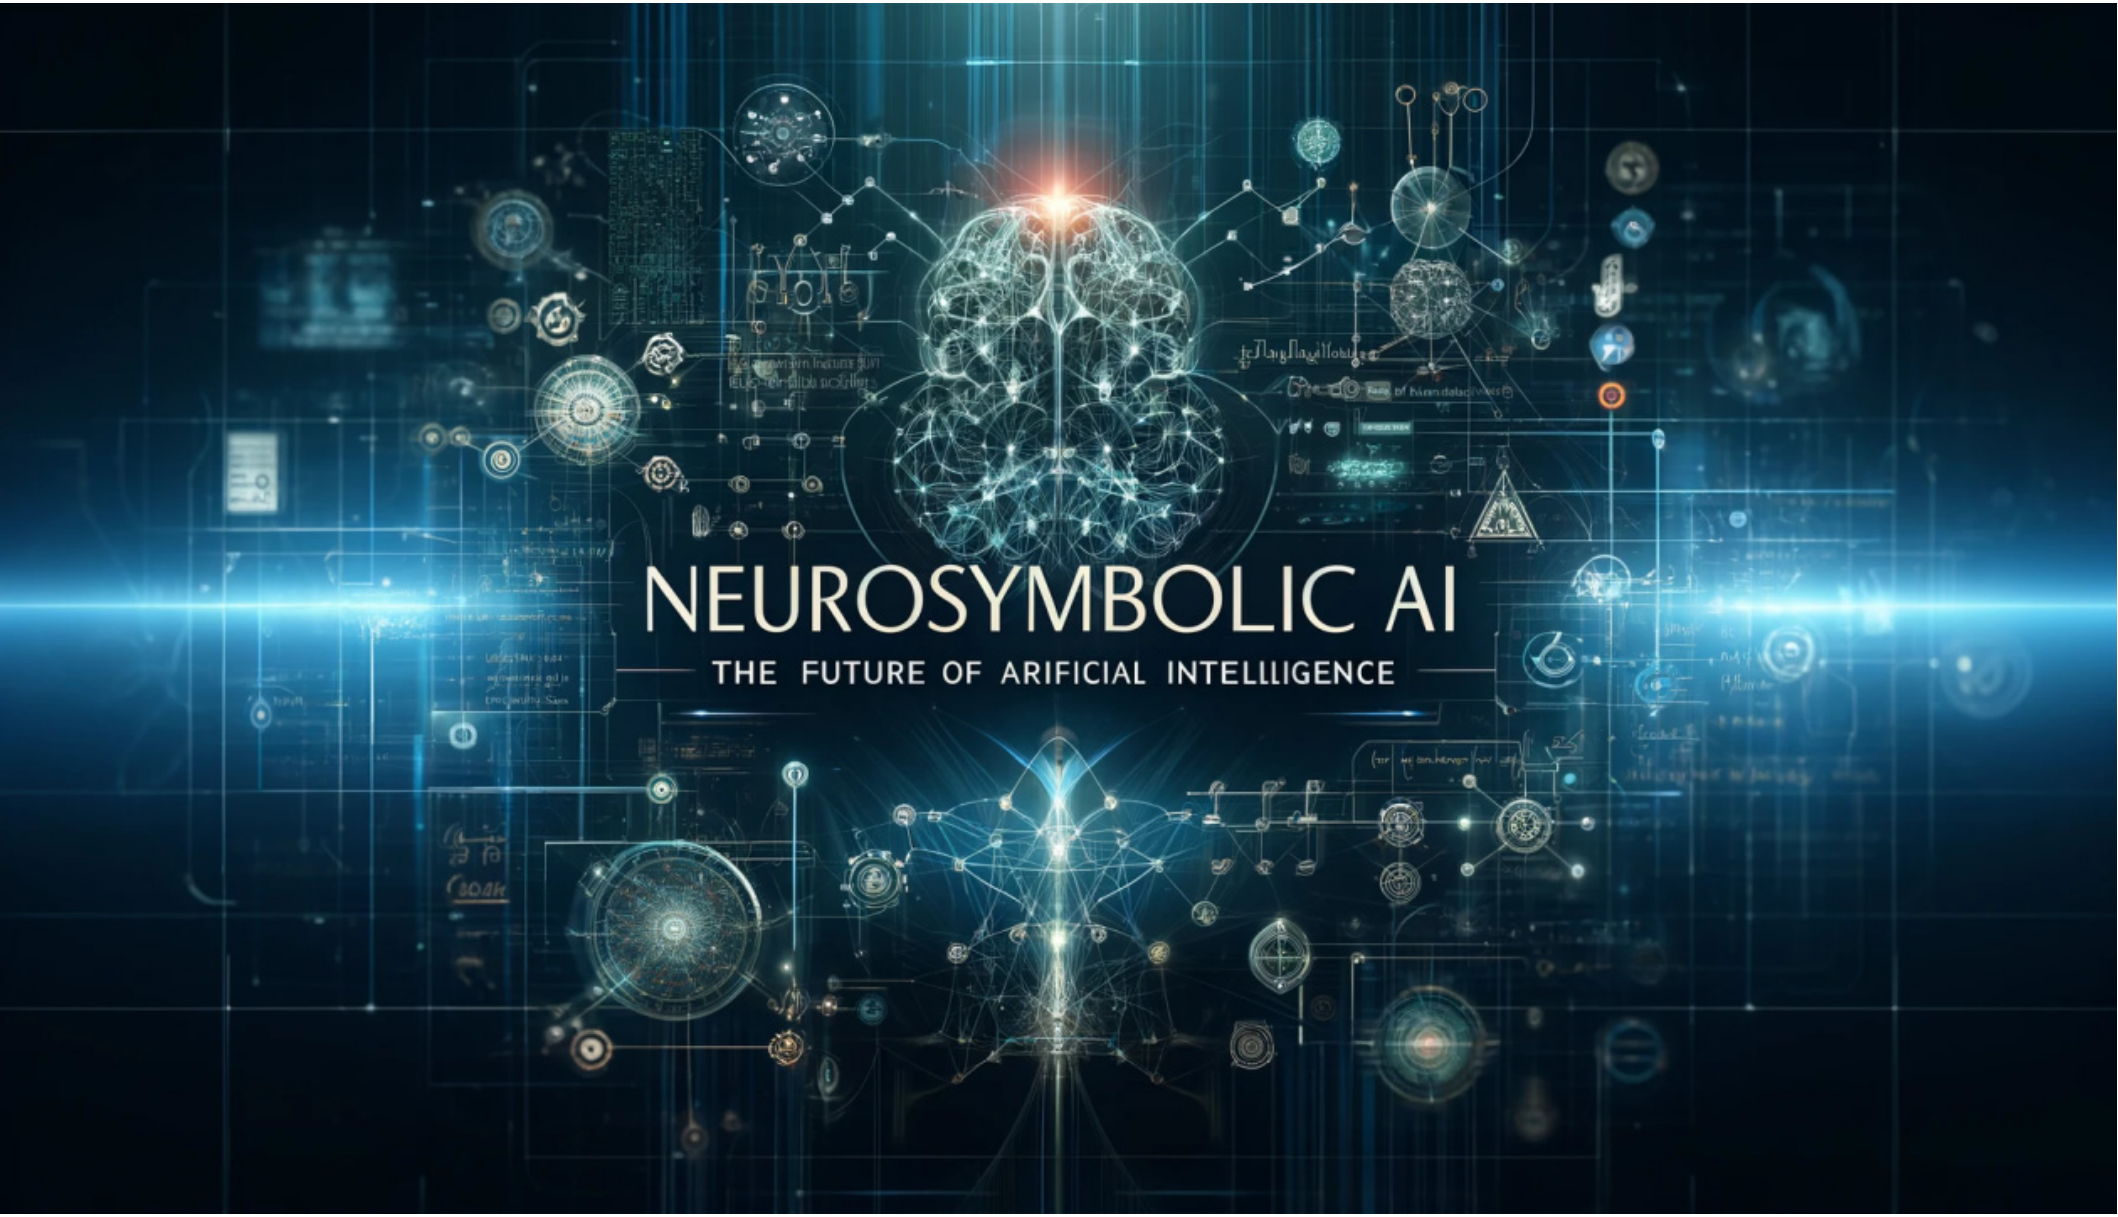
NEUROSYMBOLIC AI
THE FUTURE OF ARIFICIAL INTELLIGENCE

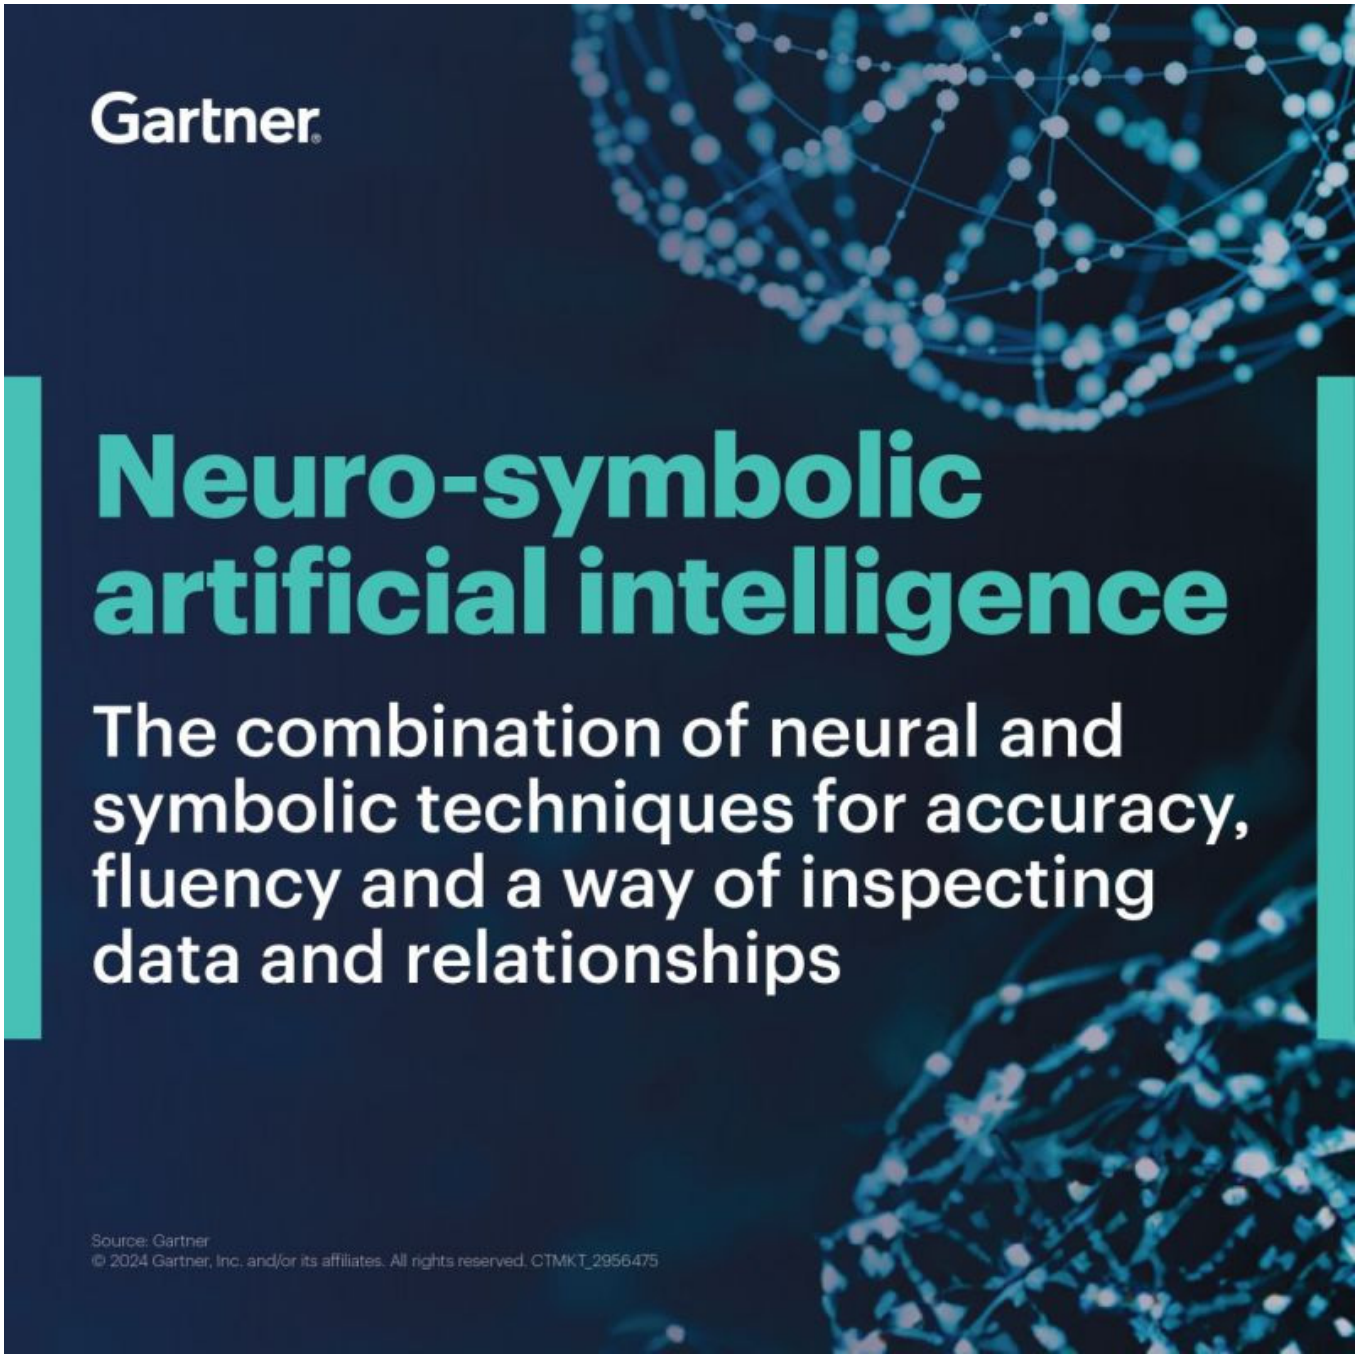
Gartner
Neuro-symbolic artificial intelligence
The combination of neural and symbolic techniques for accuracy, fluency and a way of inspecting data and relationships
Source: Gartner
© 2024 Gartner, Inc. and/or its affiliates. All rights reserved. CTMKT_2956475

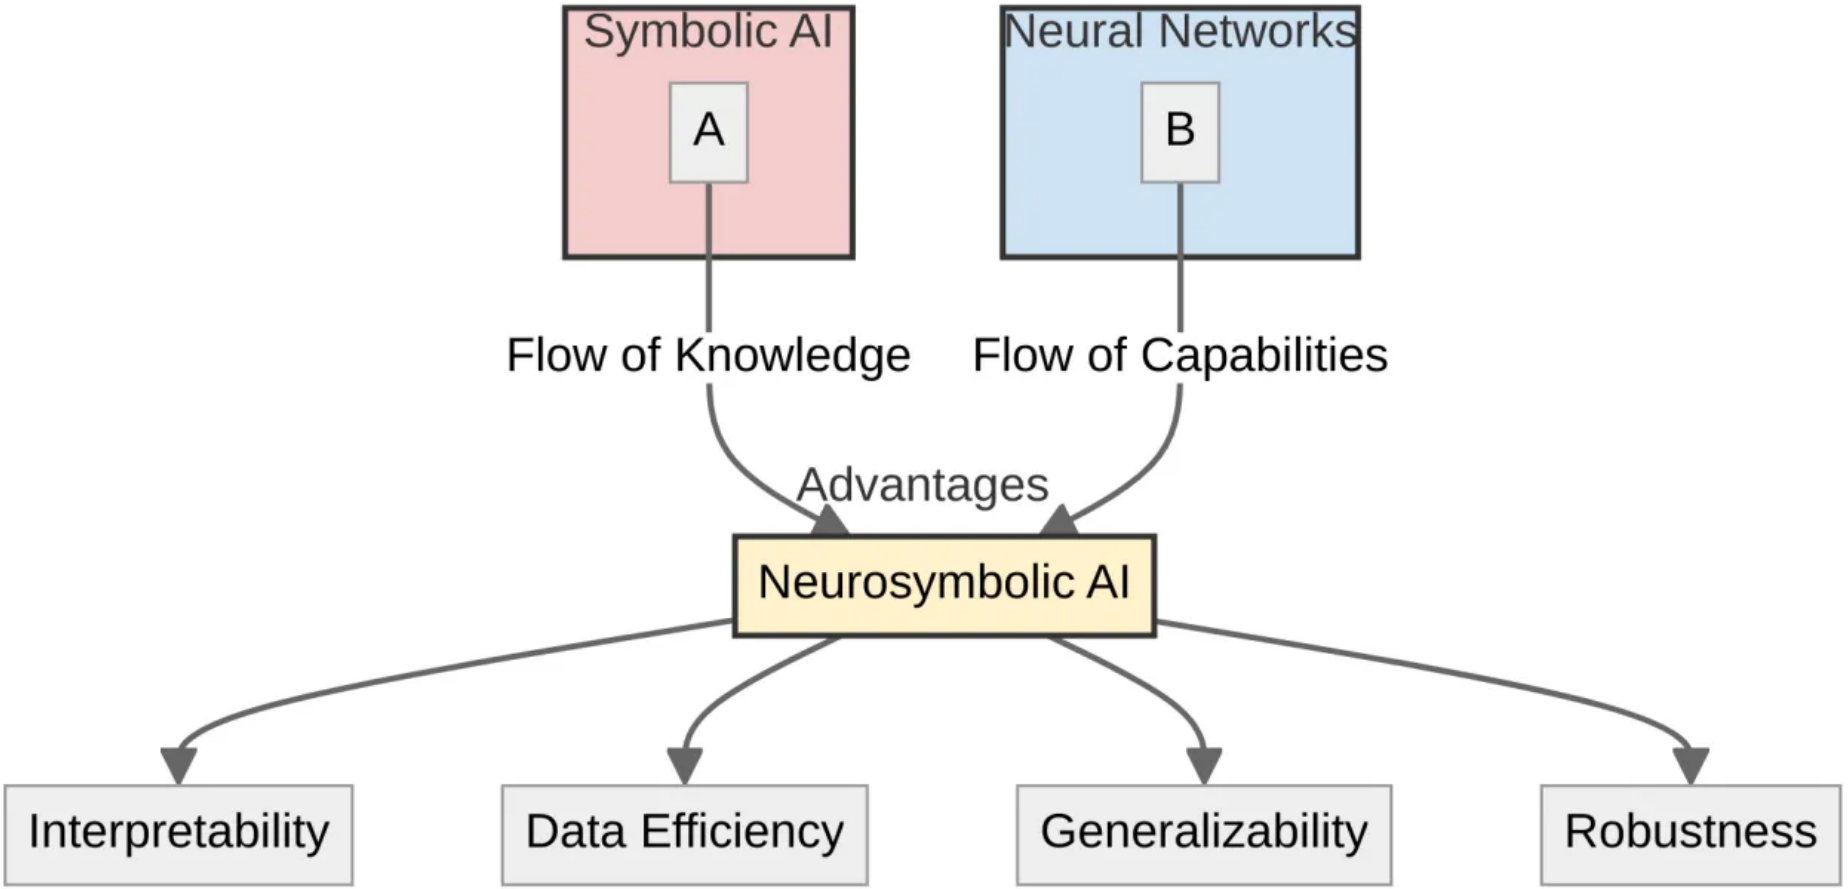
Symbolic AI
A
Neural Networks
B
Flow of Knowledge
Flow of Capabilities
Advantages
Neurosymbolic AI
Interpretability
Data Efficiency
Generalizability
Robustness

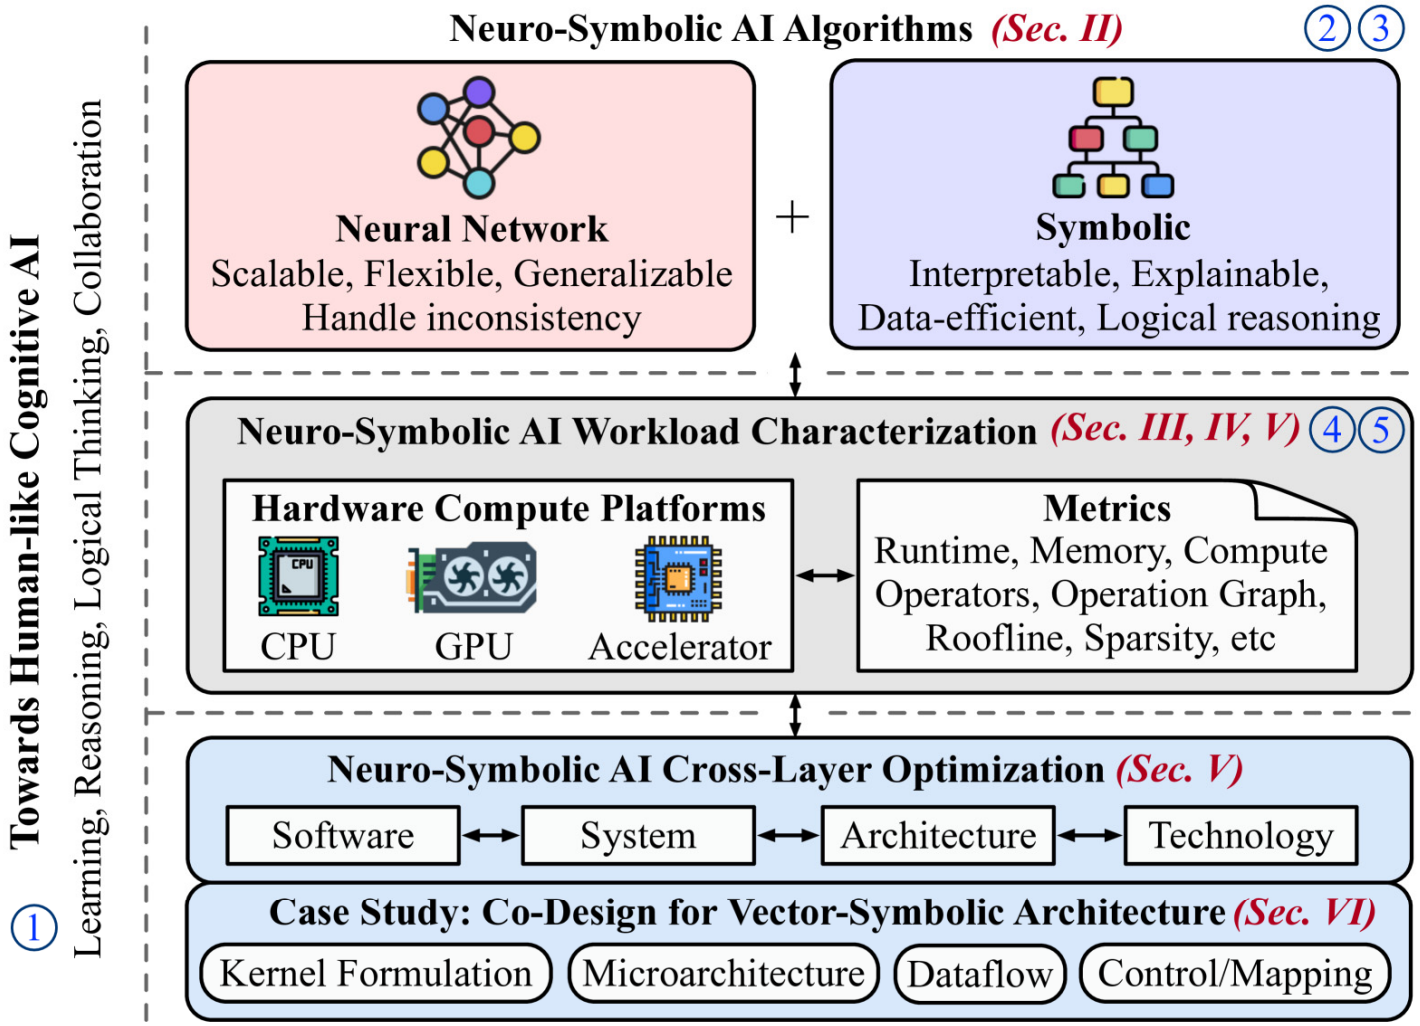

Challenges	Research Opportunities *(Sec. VII)*
Limited evaluation tasks	① Building more cognitive datasets/testbeds
Ad-hoc model design	② Unifying neuro-symbolic-probabilistic models
Limited extensibility	③ Developing efficient software frameworks
Large #algos and #HW	④ Benchmarking diverse neuro-symbolic workloads
Heterog. compute kernels	⑤ Designing cognitive architectures

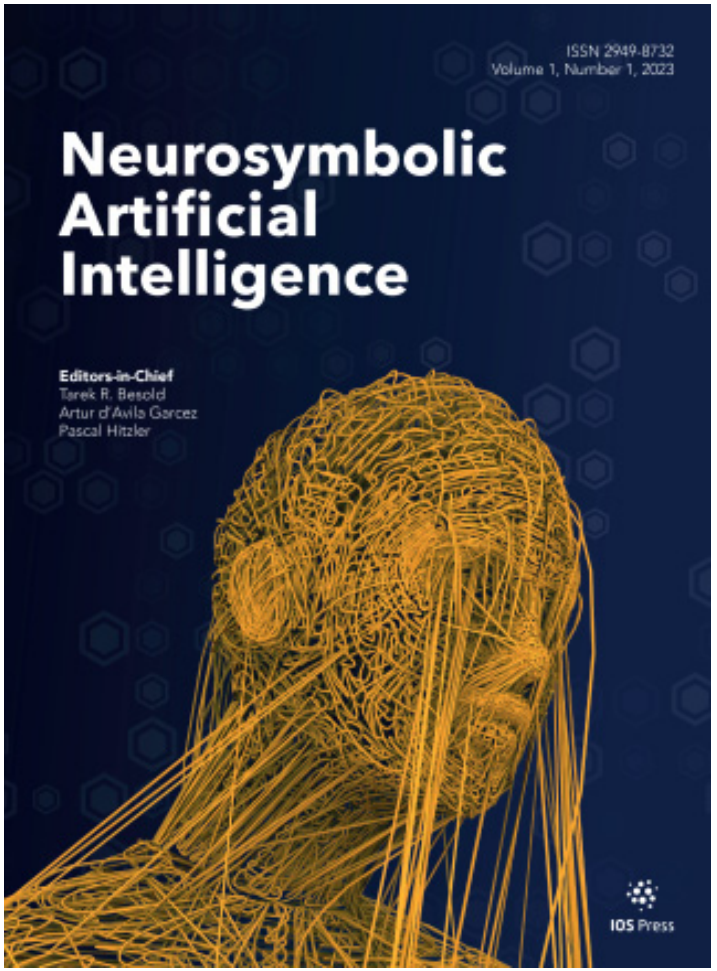
ISSN 2949-8732
Volume 1, Number 1, 2023
Neurosymbolic Artificial Intelligence
Editors-in-Chief
Tarek R. Besold
Artur d'Avila Garcez
Pascal Hitzler
IOS Press

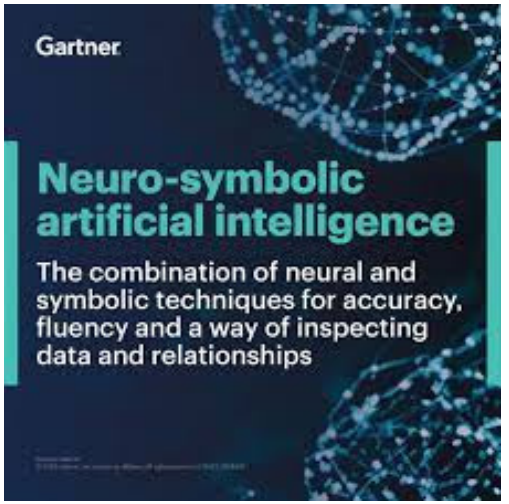
Gartner
Neuro-symbolic artificial intelligence
The combination of neural and symbolic techniques for accuracy, fluency and a way of inspecting data and relationships

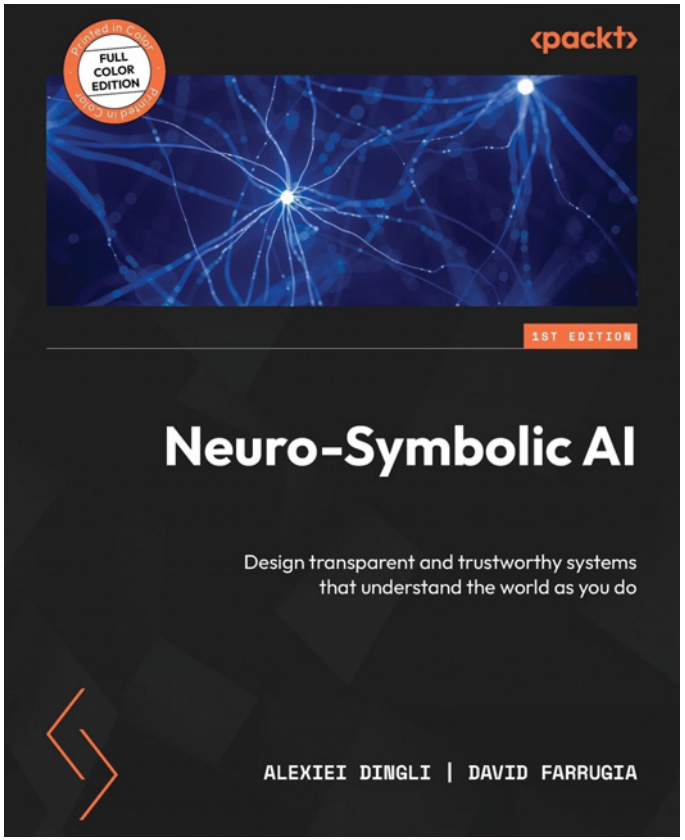
packt
FULL COLOR EDITION
1ST EDITION
Neuro-Symbolic AI
Design transparent and trustworthy systems that understand the world as you do
ALEXIEI DINGLI | DAVID FARRUGIA

The Big Players

Companies

The Big Players

Startups

Hype or Reality?

Academic Publications about Deep Learning

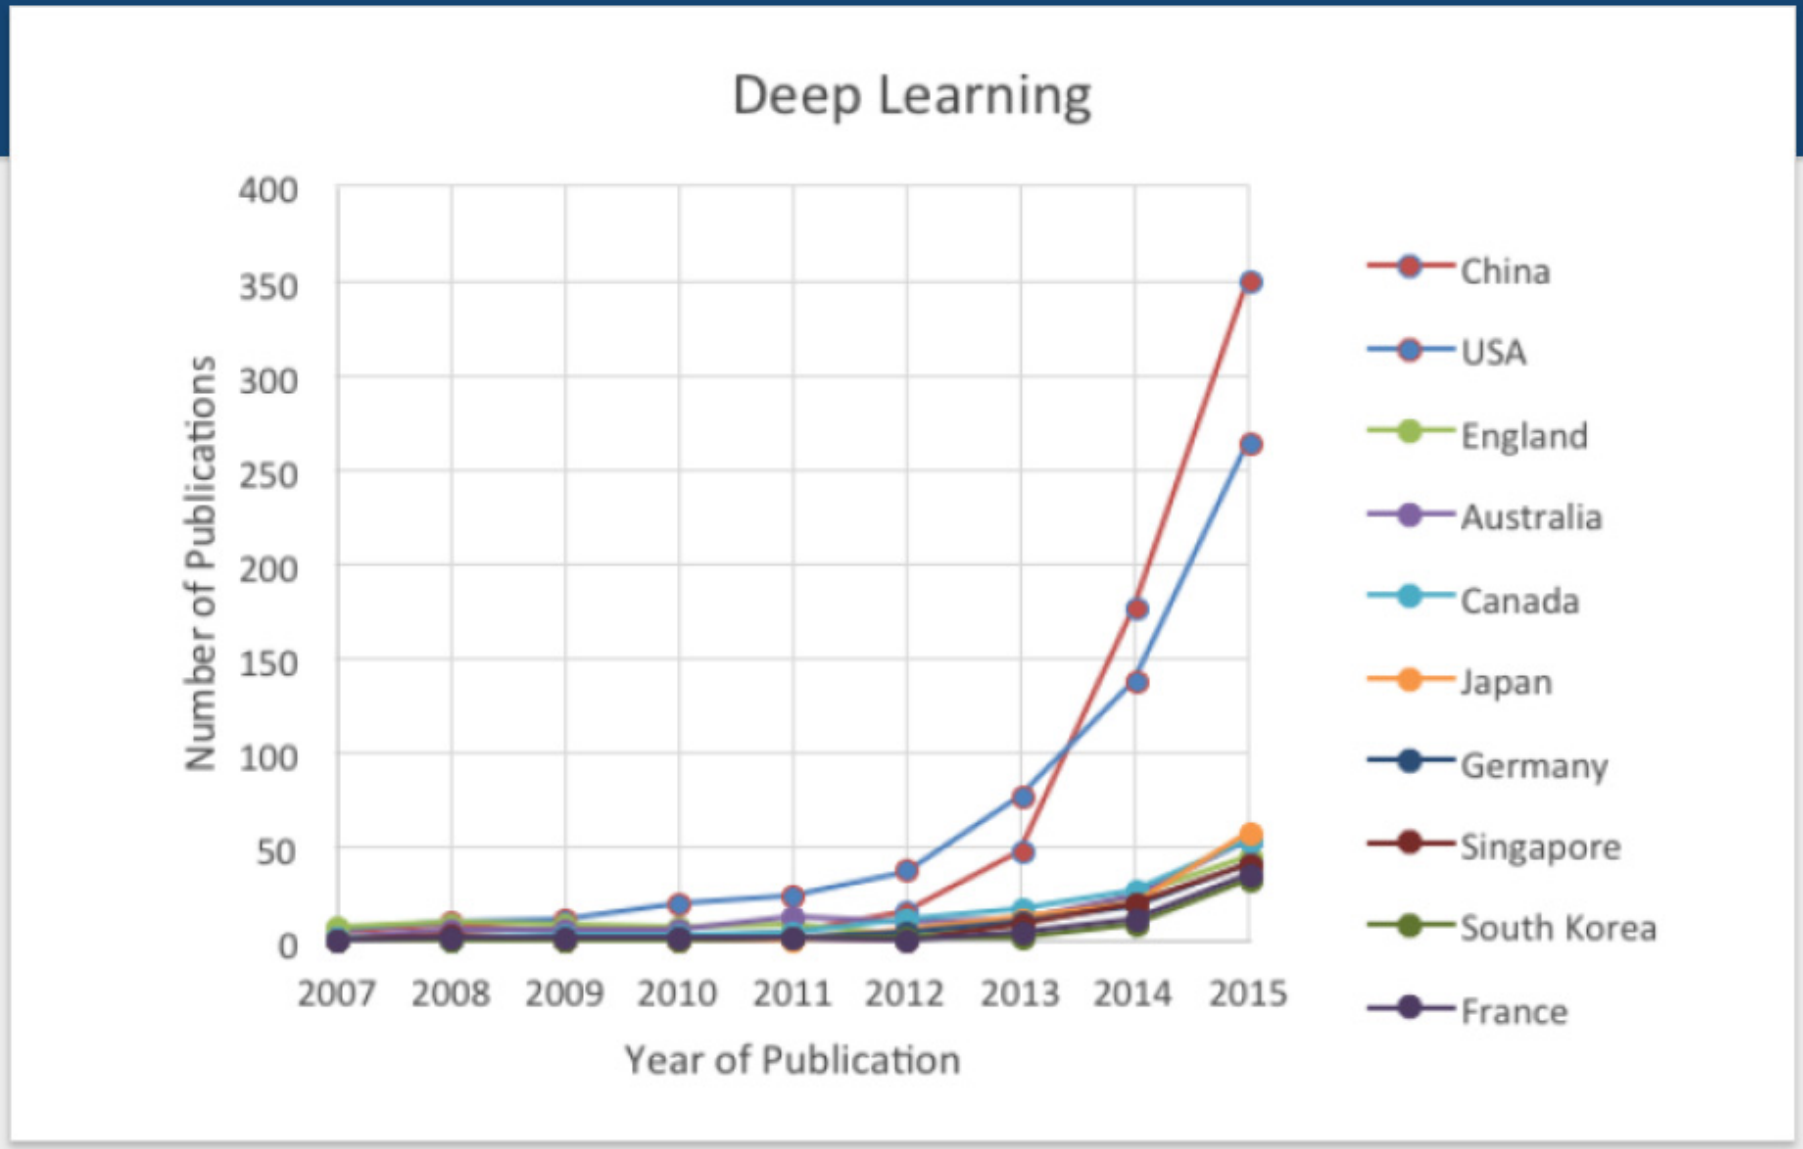

https://www.facebook.com/neil.d.lawrence/posts/10153894875151042?fref=nf&pnref=story

Hype or Reality?

NIPS (Computational Neuroscience Conference) Growth

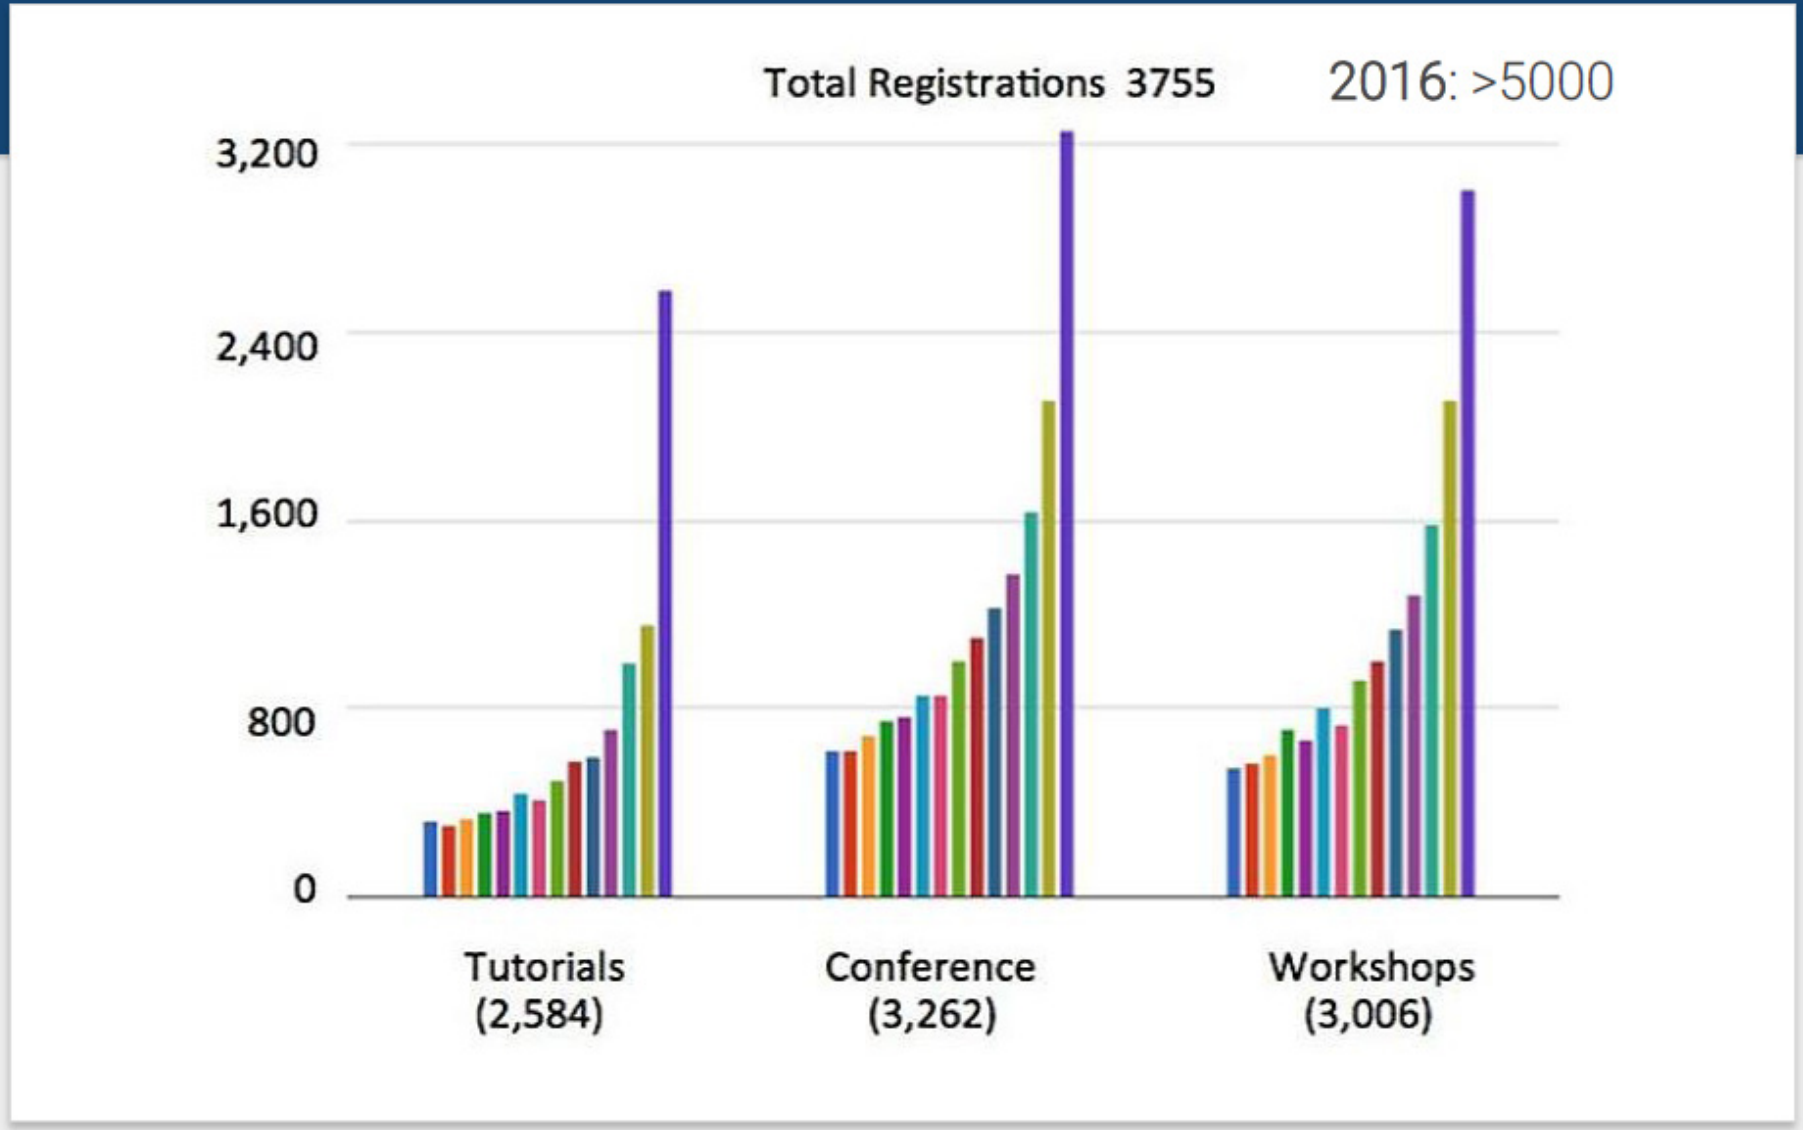

Source

https://www.wired.com/2016/06/how-google-is-remaking-itself-as-a-machine-learning-first-company/#.ekxlwoxis

Hype or Reality?

Google

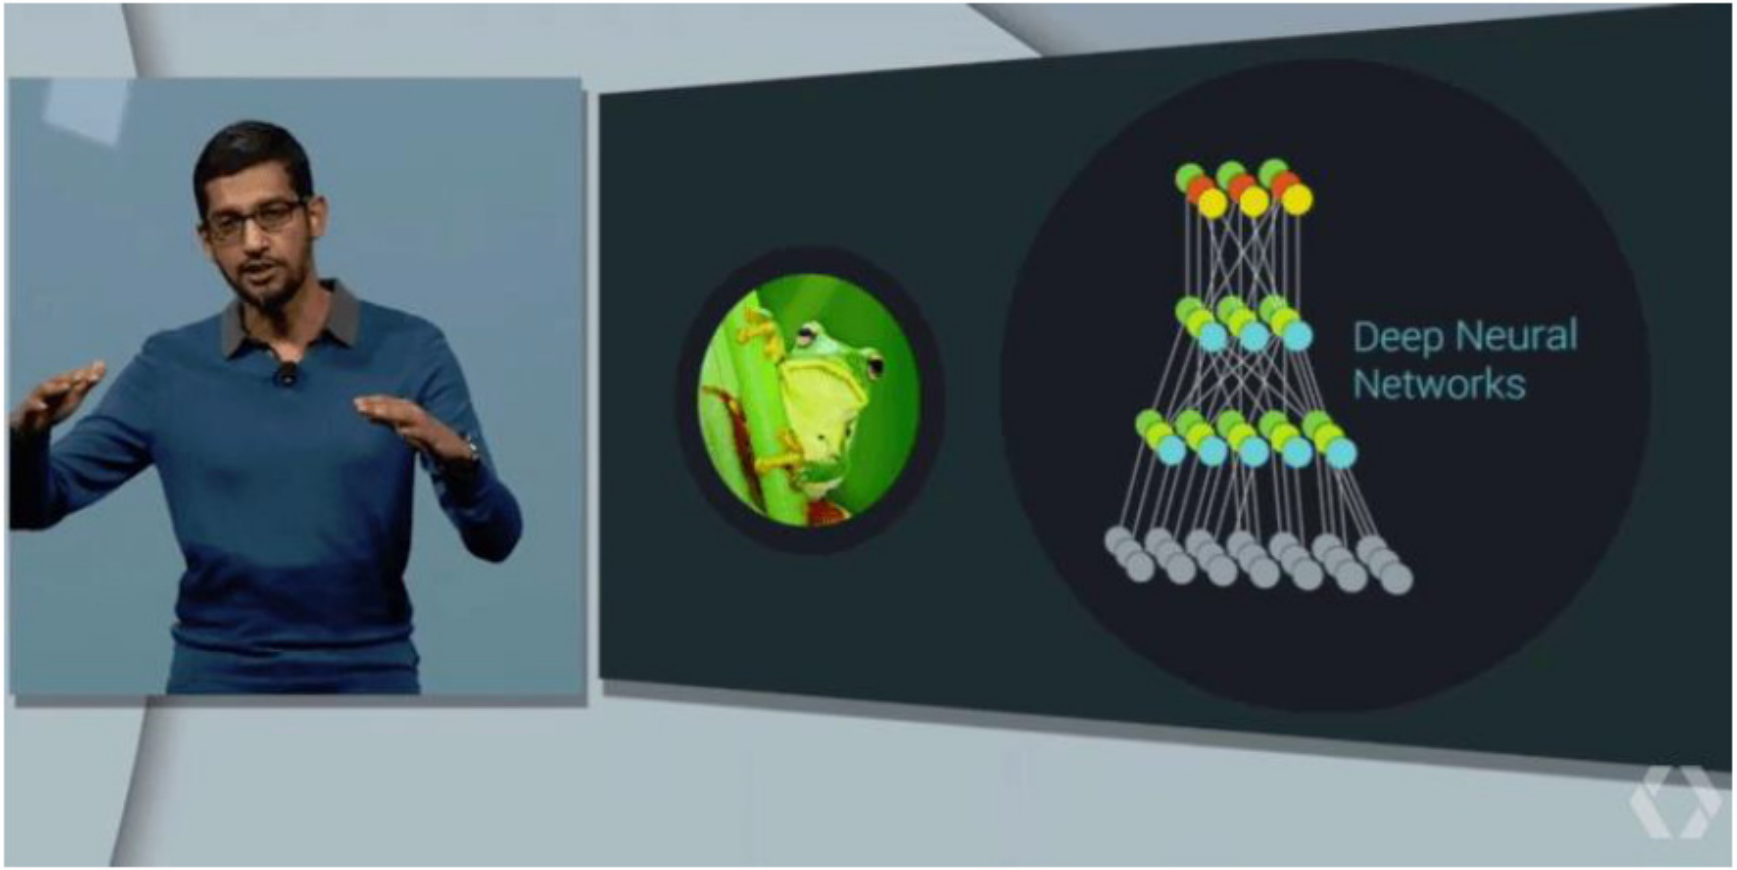

Machine learning is a core transformative way by which we are rethinking everything we are doing – *Sundar Pichai (CEO Google)*

Source

https://www.bloomberg.com/news/articles/2015-12-08/why-2015-was-a-breakthrough-year-in-artificial-intelligence

Hype or Reality?

Google

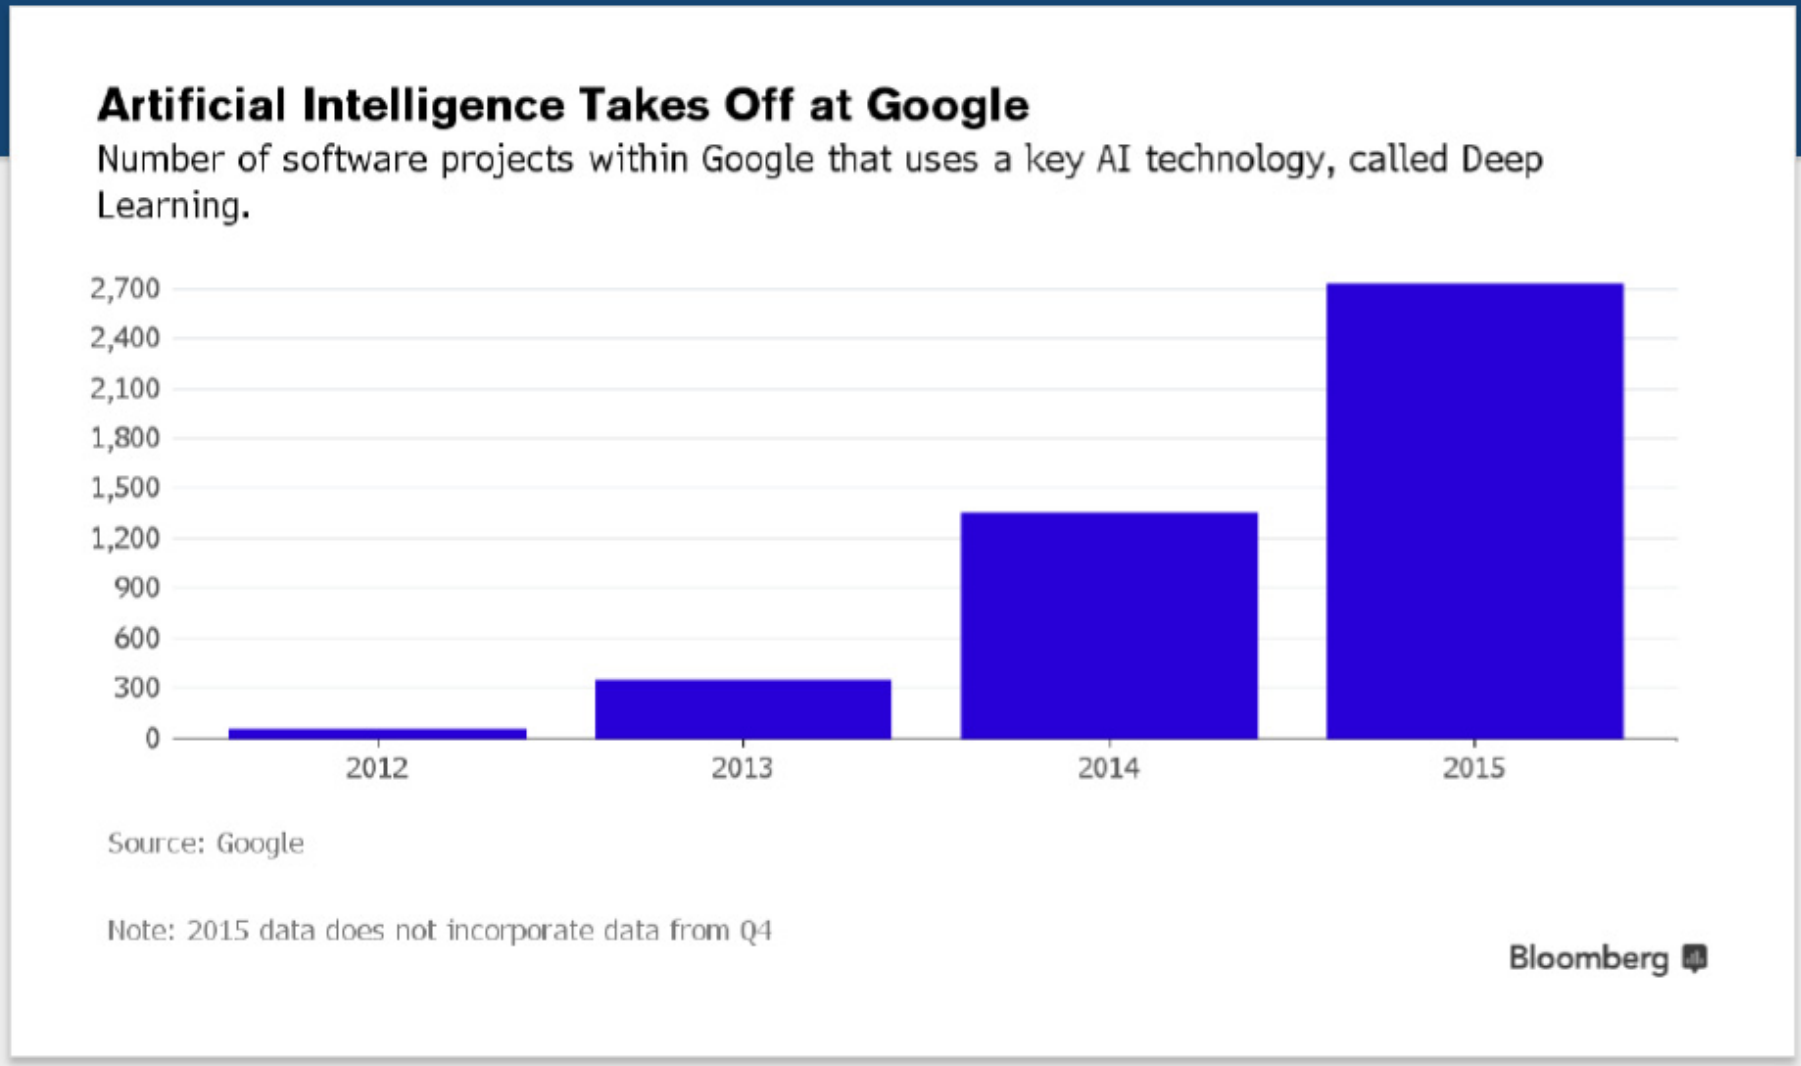

Source

http://www.nature.com/news/there-is-a-blind-spot-in-ai-research-1.20805

Hype or Reality?

Investments in AI technologies

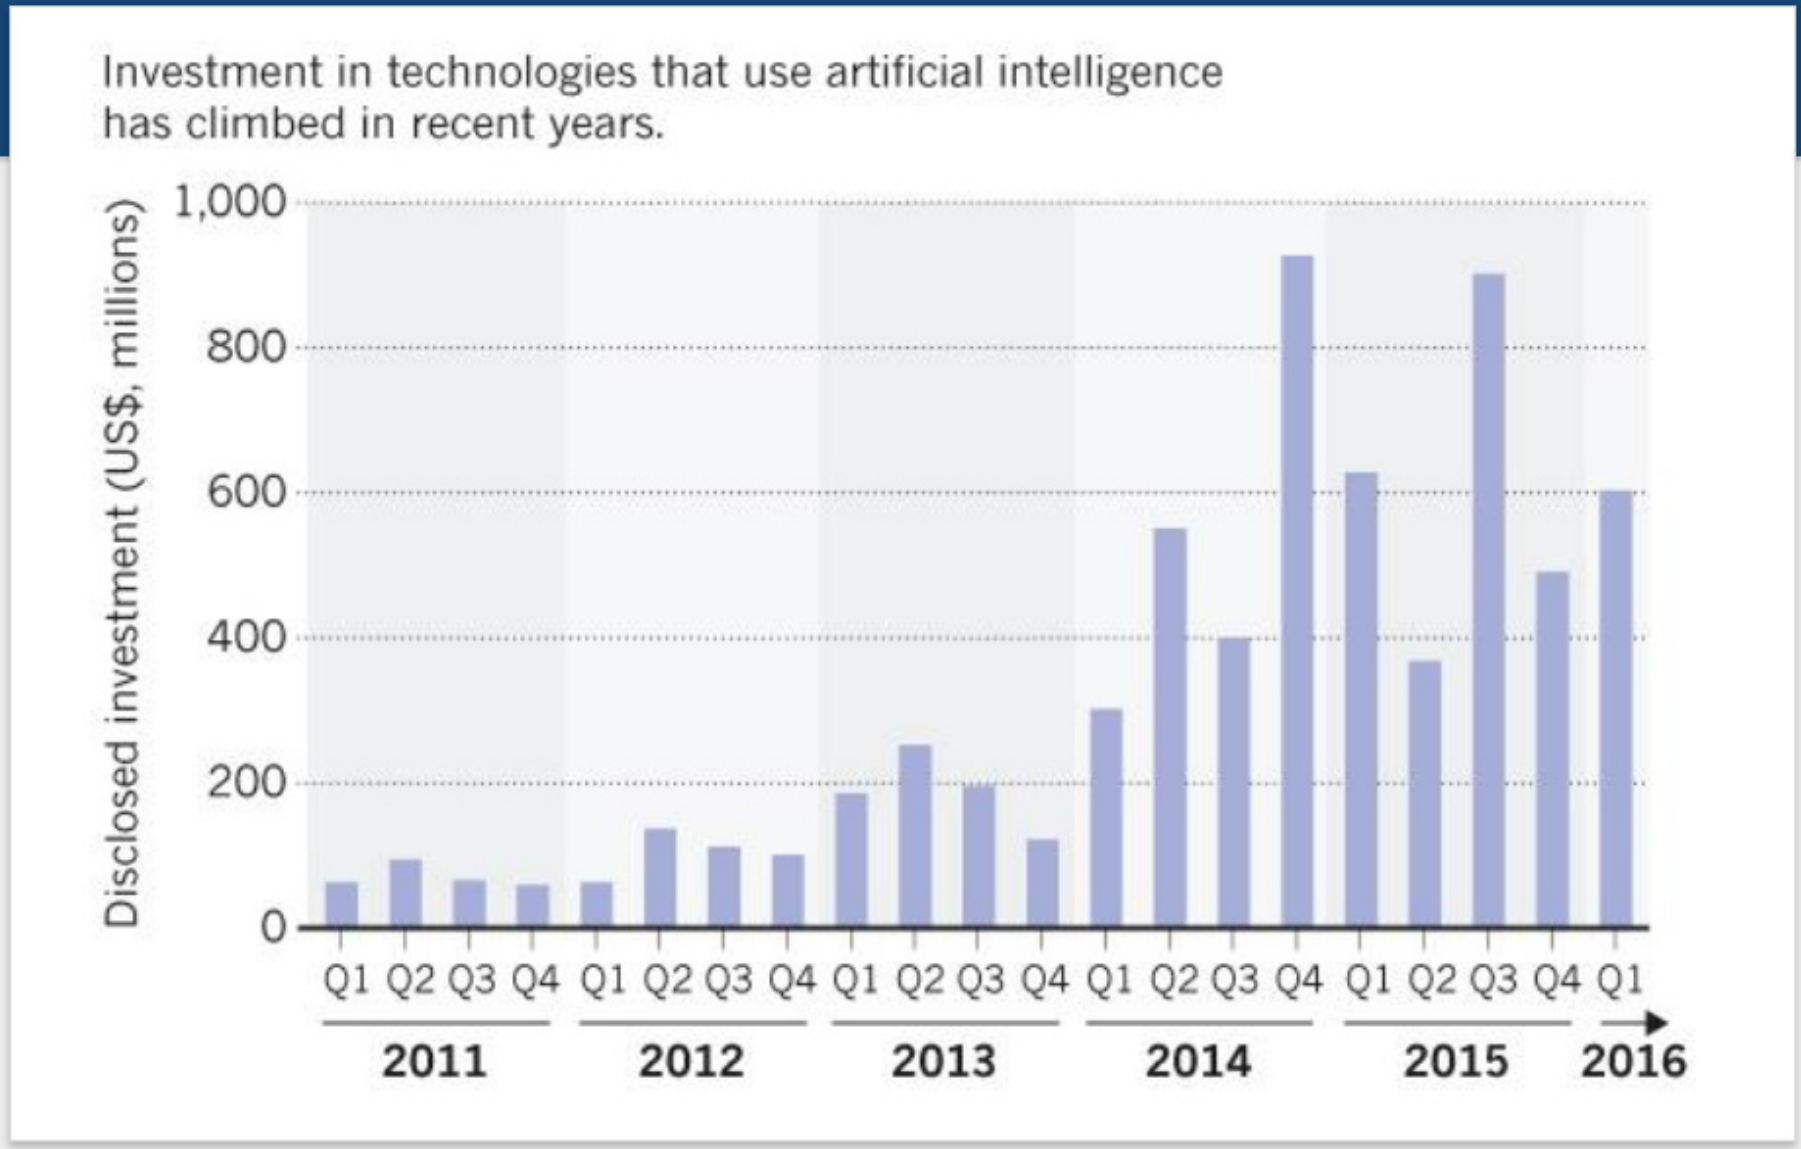

Source

https://blogs.nvidia.com/blog/2016/01/12/accelerating-ai-artificial-intelligence-gpus/

Hype or Reality?

Growing Interest from Organizations

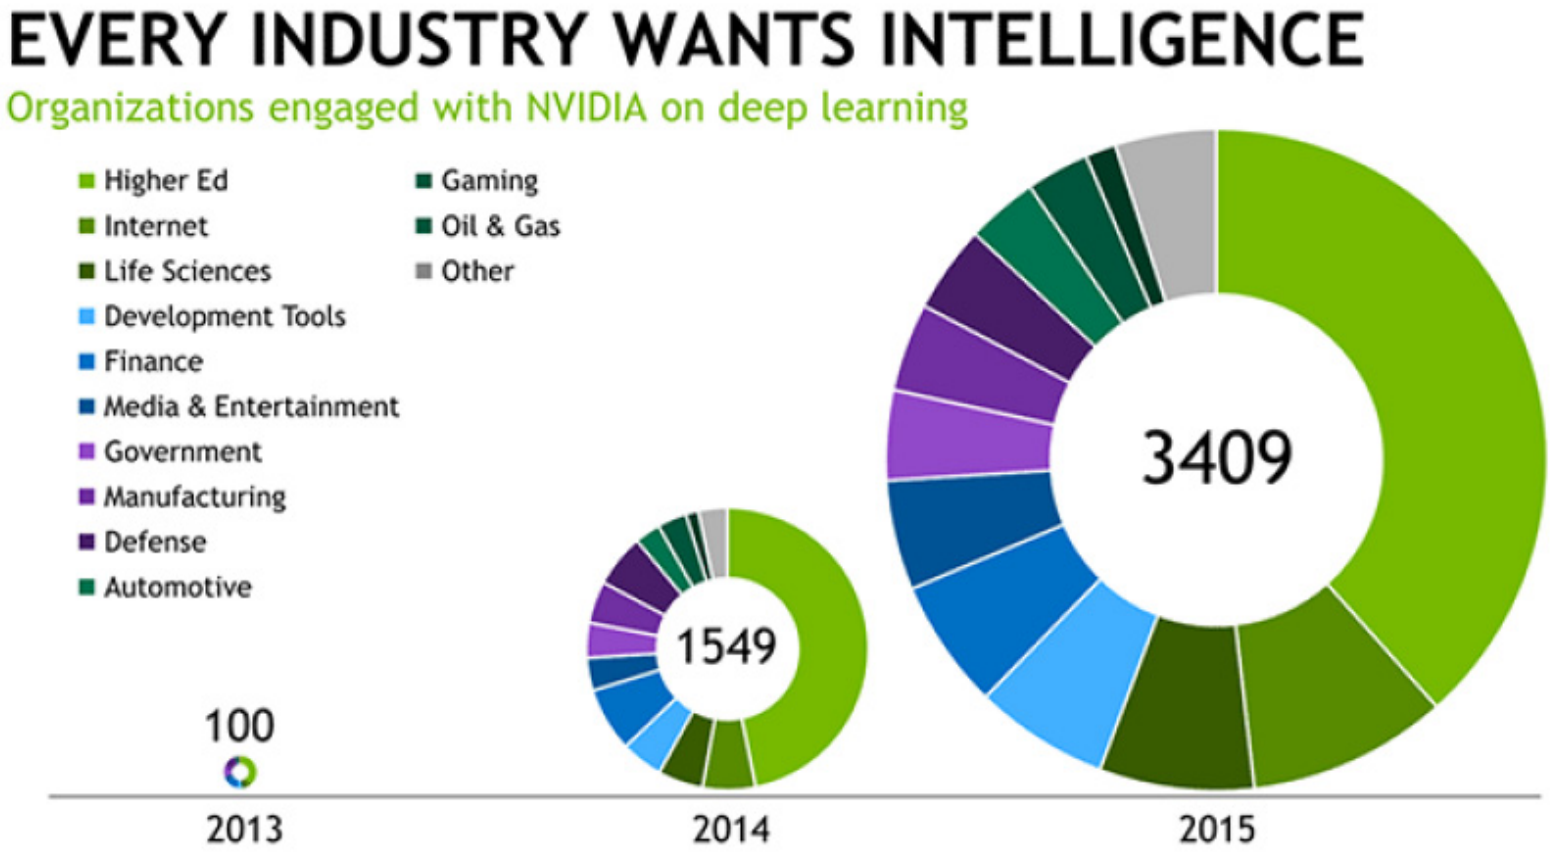

Source

The Big Players

Superstar Researchers

Geoffrey Hinton: University of Toronto & Google

Yann LeCun: New York University & Facebook

Andrew Ng: Stanford & Baidu

Yoshua Bengio: University of Montreal

Jürgen Schmidhuber: Swiss AI Lab & NNAISENSE

Hype or Reality?

Quotes

I have worked all my life in Machine Learning, and I've never seen one algorithm knock over benchmarks like Deep Learning
– *Andrew Ng (Stanford & Baidu)*

Deep Learning is an algorithm which has no theoretical limitations of what it can learn; the more data you give and the more computational time you provide, the better it is – *Geoffrey Hinton (Google)*

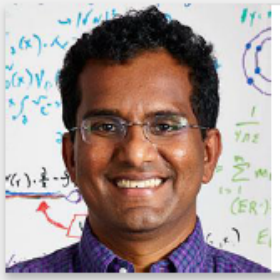

Human-level artificial intelligence has the potential to help humanity thrive more than any invention that has come before it – *Dileep George (Co-Founder Vicarious)*

For a very long time it will be a complementary tool that human scientists and human experts can use to help them with the things that humans are not naturally good – *Demis Hassabis (Co-Founder* DeepMind)

مروری بر یادگیری عمیق

۴

منابع

منبع اصلی


Deep Learning
The Past, Present and Future of
Artificial Intelligence
Lukas Masuch
@lukasmasuch
+lukasmasuch