Technische Universität Dresden

UNIVERSITY OF GOTHENBURG

Ca' Foscari University of Venice

Faculty of Business and Economics | Business Information Systems

EUTOPIA Summer School: AI in Business and Healthcare

Lecture #2: AI Fundamentals

August 24, 2026 | Patrick Zschech | Intelligent Systems and Services

INTELLIGENT INFORMATION SYSTEMS

Motivation: AI is everywhere right now

Content of this lecture

- AI brings together a diverse set of computational approaches that enable machines to perform tasks typically associated with human intelligence, such as **perception**, **reasoning**, **prediction**, and **decision-making**.

- This lecture introduces the major **paradigms of AI** and explains the core ideas behind them, as well as their respective **strengths** and **limitations**.

- Participants will gain a conceptual understanding of how **AI capabilities** have **evolved** and how the characteristics of different paradigms influence when and how they can be applied in practice.

The four AI paradigms

From manually defined rules to autonomy

	Rule-Based Automation	Machine Learning & Deep Learning	Generative AI	Agentic AI
Primary Function	Executing predefined workflows and scripts	Learning from data to predict, classify, or optimize outcomes	Generating novel content (text, image, code)	Acting autonomously to achieve goals in dynamic contexts
Core Technology/ Advancement	Symbolic reasoning through inference engines	Machine learning, deep learning, NLP, computer vision	Generative pre-trained transformer (GPT), foundation models, LLMs	Tool learning, large action and reasoning models, multi-agent collaboration
Knowledge Source	Hard-coded rules by experts	Predefined (ML) or learned (DL) features and criteria	Learned patterns from large-scale data	Dynamically inferred by interacting with enterprise systems
Human Involvement	Fully manual setup and rule definition	Data labeling, feature engineering, model training and evaluation	Prompt engineering, fine-tuning, and human-in-the-loop validation	Prompt & context engineering, agent & workflow orchestration, oversight
Adaptivity	Operates only within predefined parameters	Generalizes from data but lacks real-time context awareness	Produces outputs conditioned on prompts and training context	Perceives, reasons, and acts based on evolving environmental states
Enterprise Integration	Limited and static workflows	Predefined enterprise processes and API-based integration	Ability to call functions and connect to enterprise data (RAG)	Autonomous discovery and coordination across enterprise tools, APIs, and agents
Example Applications	Invoice processing, compliance checks	Spam filter, credit scoring, image recognition	Marketing content generation, coding, document drafting	Sales agents, customer support agents, operations co-pilots

Rules & determinism ---- Learning from data ---- Autonomy & reasoning →

Adapted from Li MM, Reinhard P, Leimeister JM (2026) Agentic AI. *Business & Information Systems Engineering* 1–14. https://doi.org/10.1007/s12599-026-01009-w

1) **Symbolic AI and rule-based automation**
2) (Shallow) machine learning
3) Neural networks and deep learning
4) Foundation models and generative AI
5) Tool learning and agentic AI

Symbolic AI and rule-based automation

```
IF temperature > 38.5°C
AND white blood cell count > 12,000/mm³
AND CRP elevated
THEN consider bacterial infection
```

General scheme of a rule-based expert system

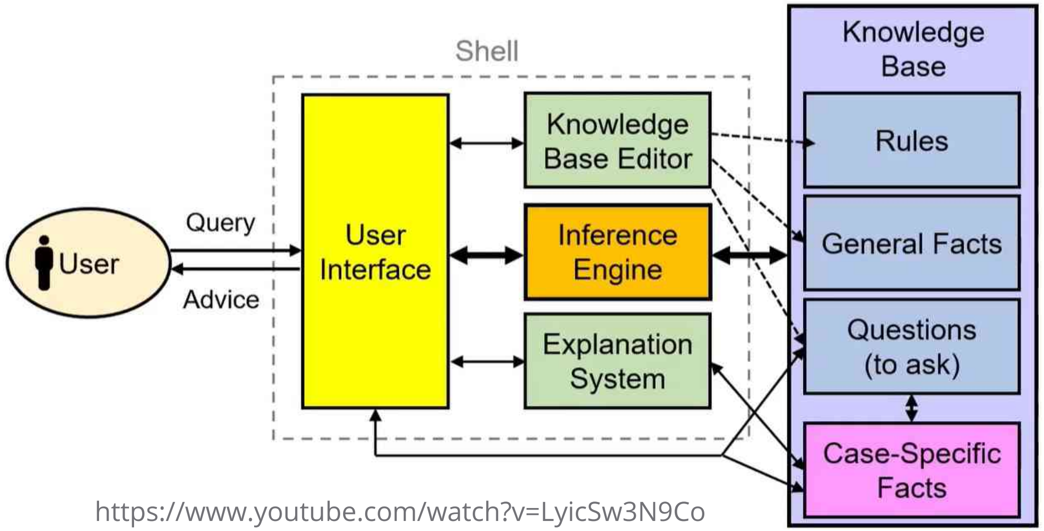

https://www.youtube.com/watch?v=LyicSw3N9Co

An early platform for expert systems (1984)

https://en.wikipedia.org/wiki/Expert_system

1) Symbolic AI and rule-based automation
2) **(Shallow) machine learning**
3) Neural networks and deep learning
4) Foundation models and generative AI
5) Tool learning and agentic AI

Rule-based AI vs. machine learning

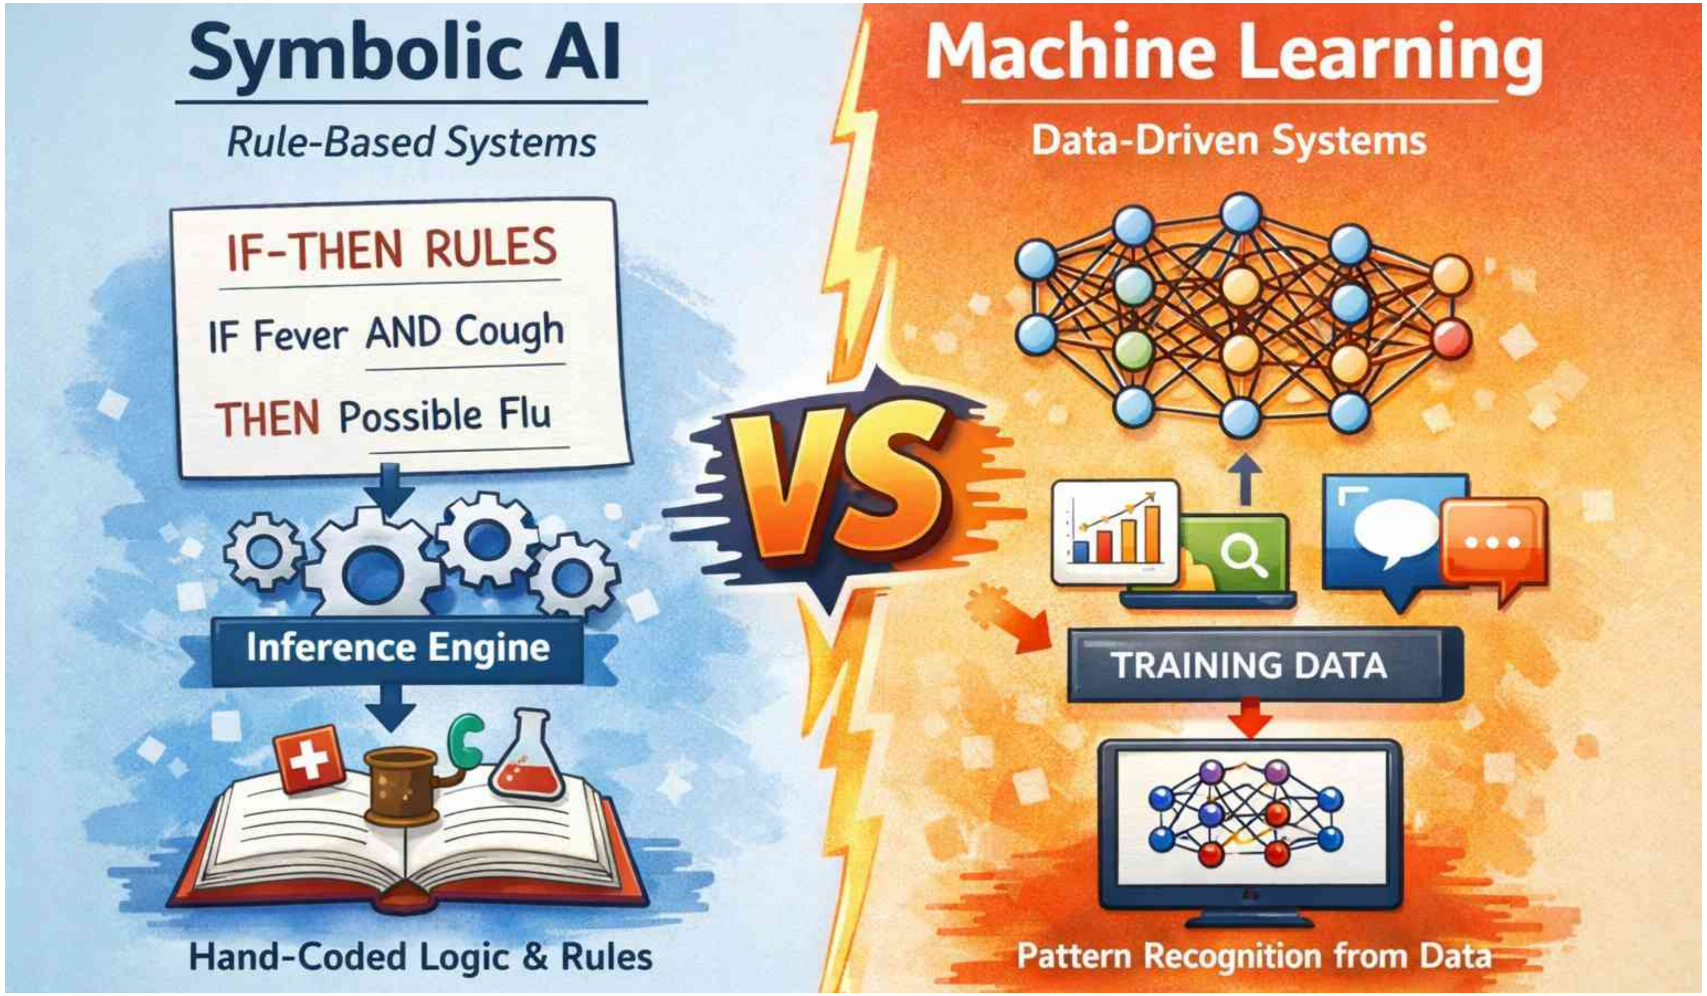

Machine learning – different types of learning

Supervised vs. unsupervised vs. reinforcement learning

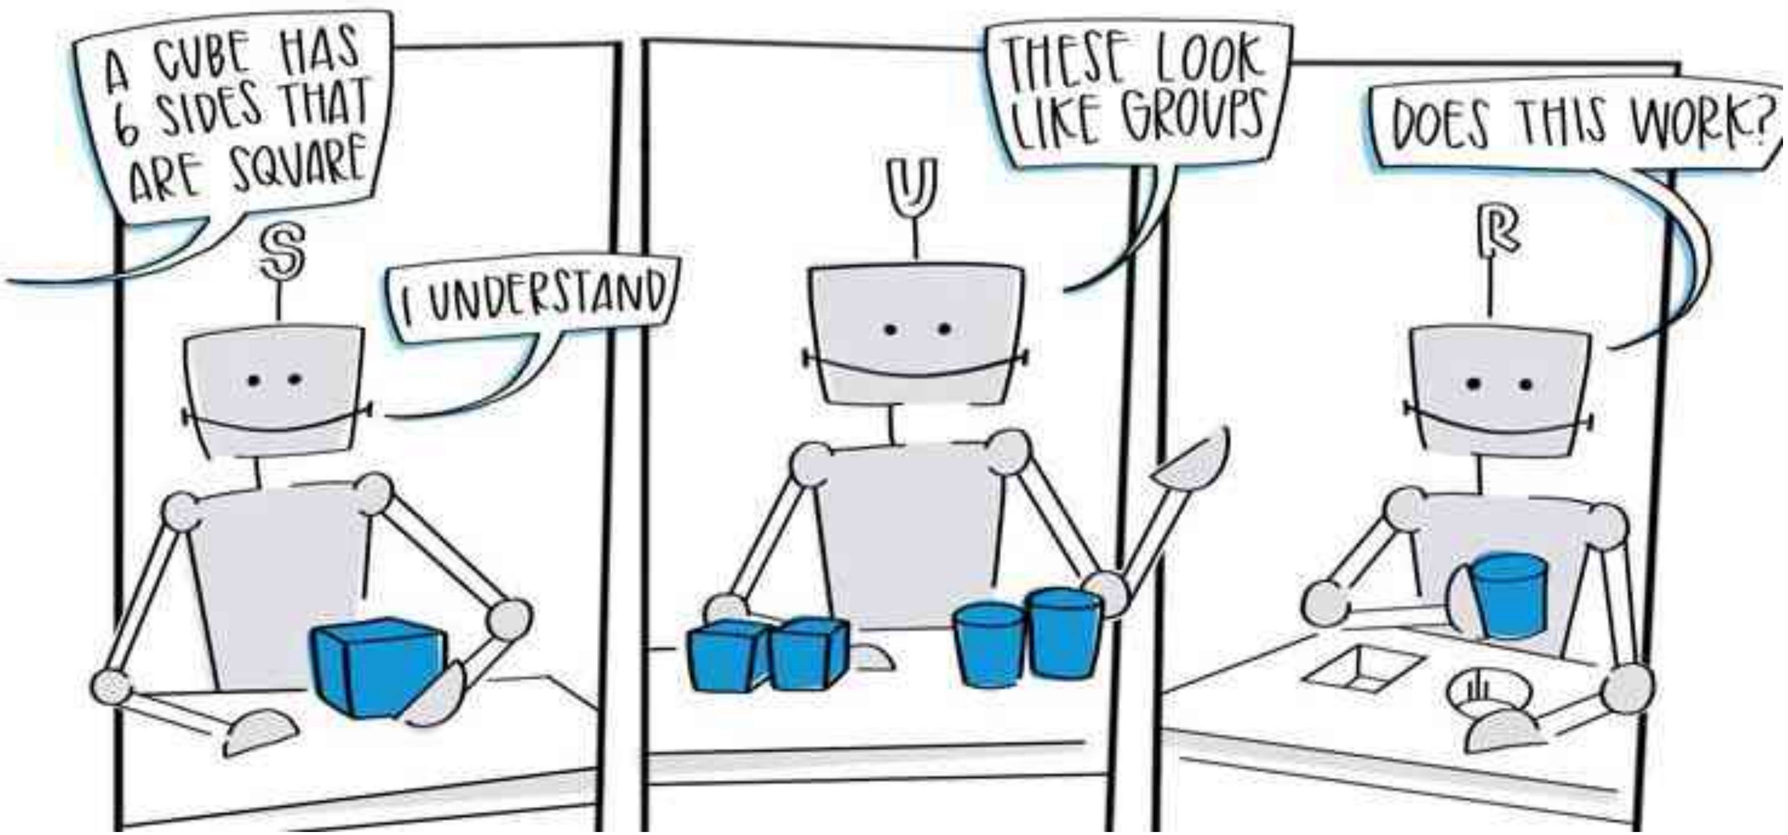

supervised **unsupervised** **reinforcement**

https://www.ceralytics.com/3-types-of-machine-learning/

Machine learning – different types of learning

Supervised vs. unsupervised vs. reinforcement learning

Supervised learning:

- **Principle:** The algorithm is given training data which contains examples of **input variables / features (= X)** along with their corresponding **output variables / target (= y)** as the "correct answer" and the goal is to learn a function that maps inputs to outputs.
- **Examples:** Classification models, regression models

Unsupervised learning:

- **Principle**: No labels (i.e., target values) are given to the algorithm, leaving it on its own to **find structure in the input data (= X)**. Unsupervised learning can be a goal by itself (discovering hidden patterns) or a means towards a supervised approach (feature learning).
- **Examples:** Cluster analysis, dimensionality reduction, density estimation

Reinforcement learning:

- **Principle:** "(…) problem of finding suitable actions to take in a given situation in order to **maximize a reward**. Here the learning algorithm is not given examples of optimal outputs, in contrast to supervised learning, but must instead discover them by a **process of trial and error**. Typically, there is a sequence of states and actions in which the learning algorithm is interacting with its environment." (Bishop, 2006)
- **Example:** Trading agents, gaming

Machine learning with tabular data

Features (X) and targets (y) are organized in tables with columns and rows

- **Tabular data** refers to data that is organized in a table with rows and columns.
  - The **rows** represent observations (data points / **samples**).
  - The **columns** represent attributes for those observations (variables / **features**).

Tabular Data

Columns

Rows

age	weight	gender	smoking	target: heart attack
high (65)	low (50)	male	yes	no
high (60)	low (53)	male	yes	no
high (65)	low (59)	male	yes	no
high (60)	medium (63)	male	yes	no
high (58)	high (77)	female	no	no
high (60)	high (81)	female	yes	**yes**
high (63)	high (85)	male	yes	**yes**
high (65)	high (89)	female	no	**yes**
low (25)	medium (63)	male	yes	no
low (30)	medium (70)	female	yes	**yes**
low (20)	medium (70)	male	yes	no
low (30)	high (77)	female	no	**yes**
low (20)	high (77)	male	no	no
low (25)	high (85)	male	yes	**yes**

Machine learning with tabular data

A simple regression example for house price prediction

General model structure:

$$y = a + b \times x$$

Application-specific model:

$$price = a + b \times square\ feet$$

Specific model with trained weights:

$$price = 50{,}000 + 300 \times square\ feet$$

Calculation example:

$$200{,}000 = 50{,}000 + 300 \times 500$$
$$200{,}000 = 50{,}000 + 150{,}000$$

Square feet (=X)	Actual price (=y) (blue points)	Predicted price (=ŷ) (green line)
500	210,000	200,000
800	290,000	290,000
1200	410,000	410,000
1800	600,000	590,000
2000	670,000	650,000
…	…	…

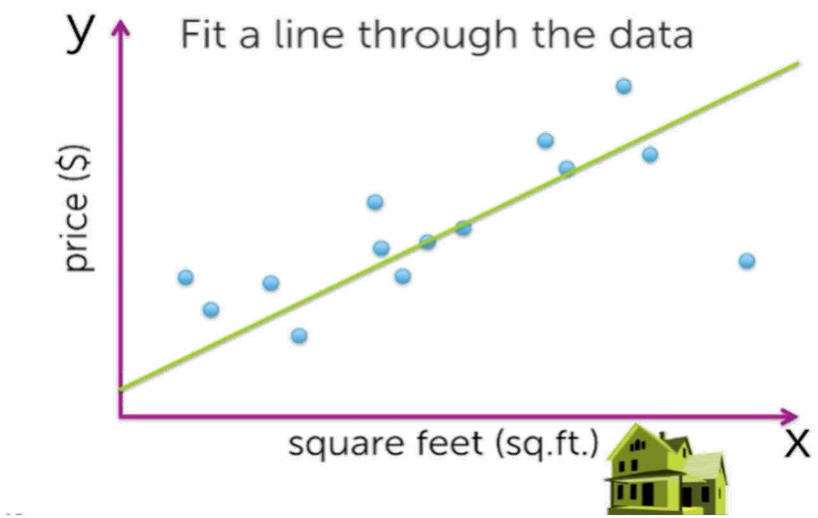

Tabular data vs. images vs. text

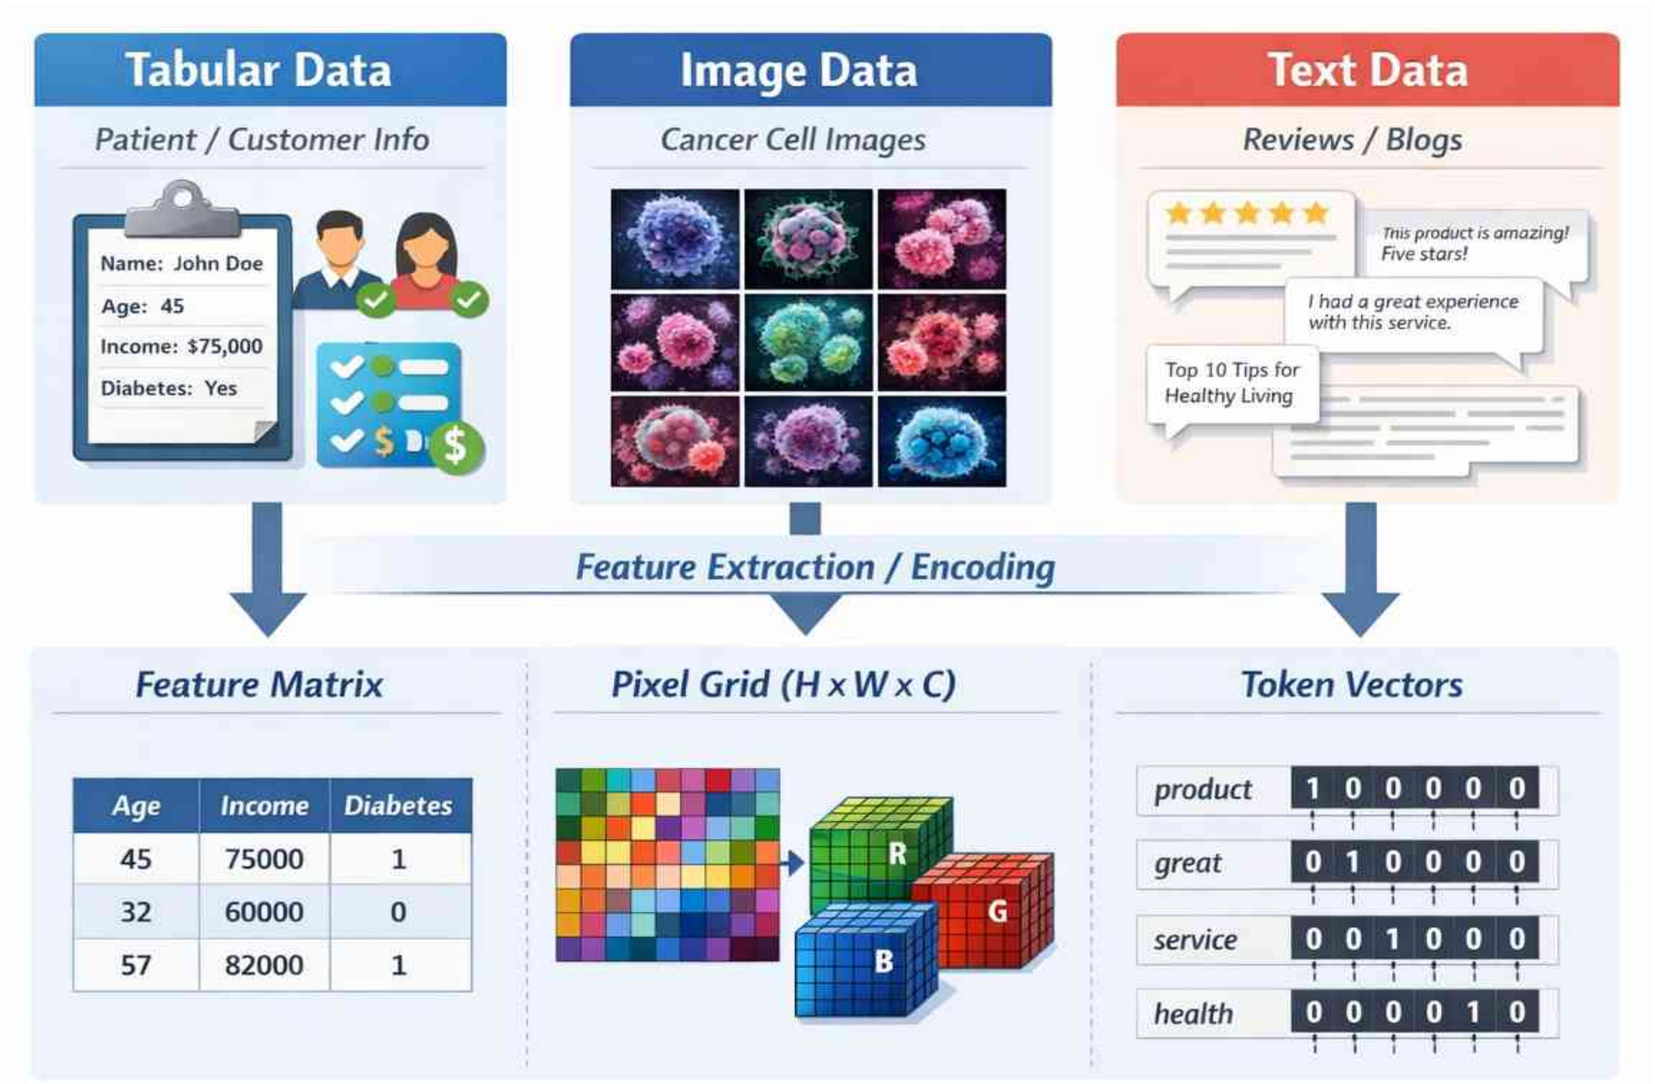

Examples of fascinating applications

Some examples of supervised learning tasks

Features / input (= x)	Target / output (= y)	Application
Voice recording	Transcript	Speech recognition
Historical market data	Future market data	Trading bots
Photograph	Caption	Image tagging
Drug chemical properties	Treatment efficacy	Pharma R&D
Store transaction details	Is the transaction fraudulent?	Fraud detection
Recipe ingredients	Customer reviews	Food recommendations
Purchase histories	Future purchase behavior	Customer retention
Car locations and speed	Traffic flow	Traffic lights
Faces	Names	Face recognition

The universe of machine learning

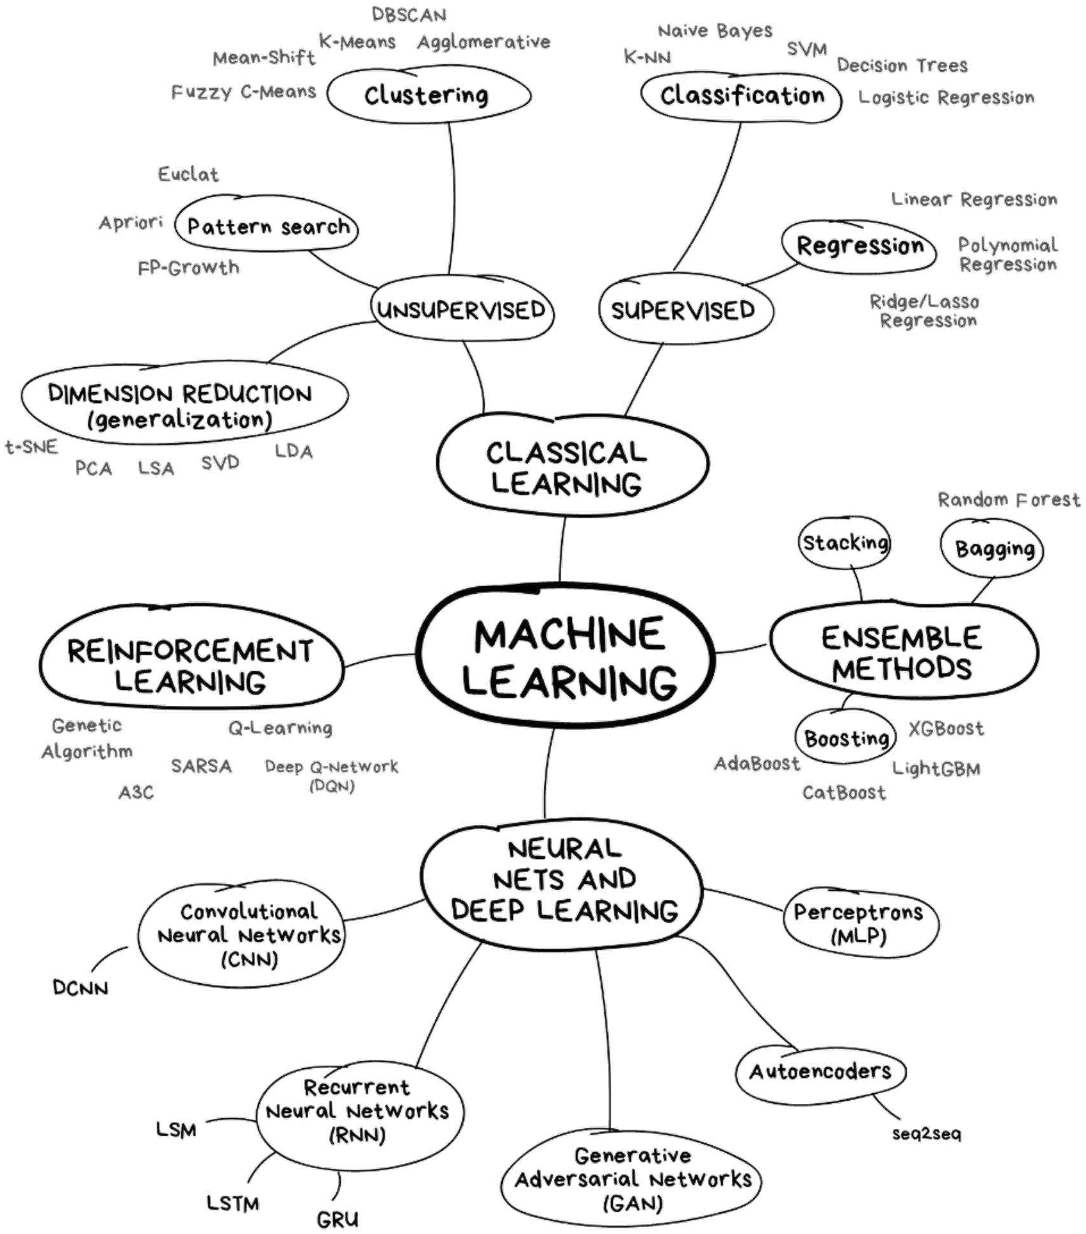

Broad variety of machine learning models

From simple to complex models

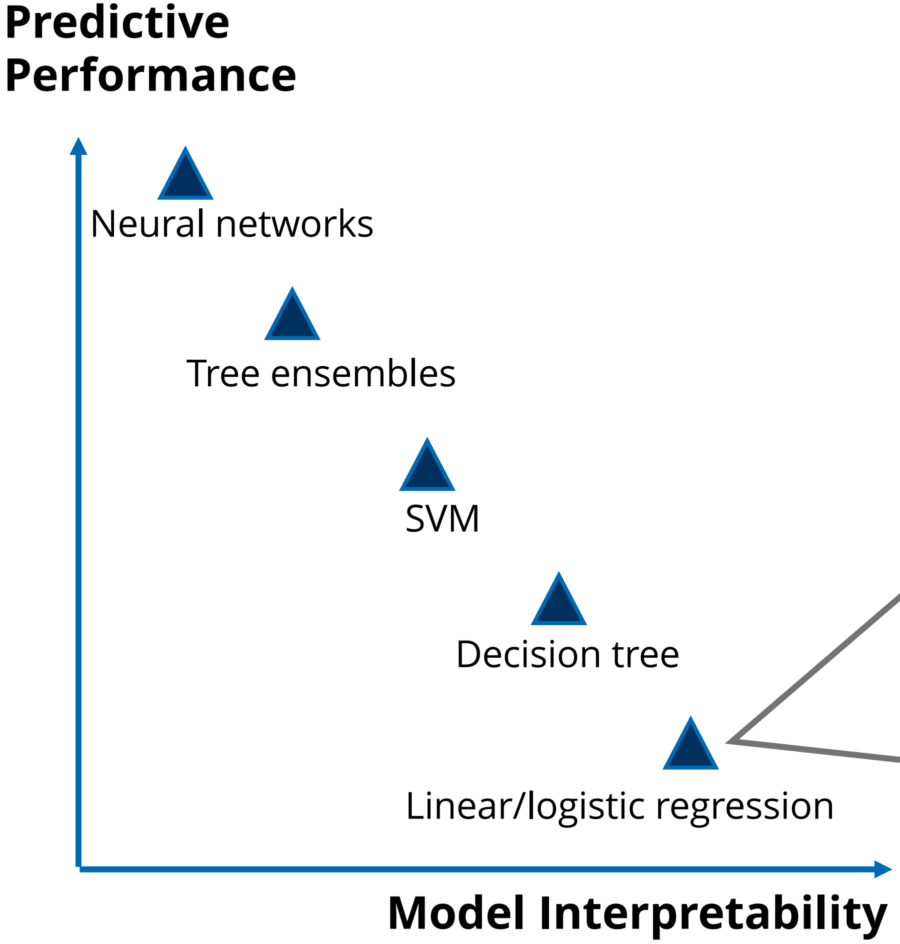

Linear regression (LR): $$\boldsymbol{y = a + bx}$$

$$price = a + b \times square\ feet$$

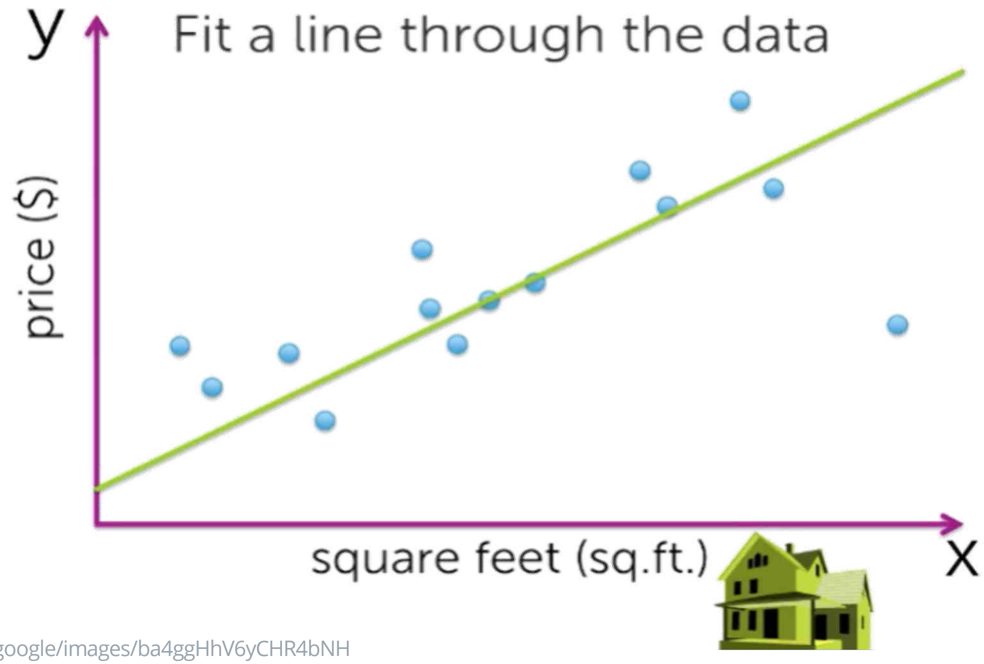

https://share.google/images/ba4ggHhV6yCHR4bNH

Broad variety of machine learning models
From simple to complex models

Predictive Performance

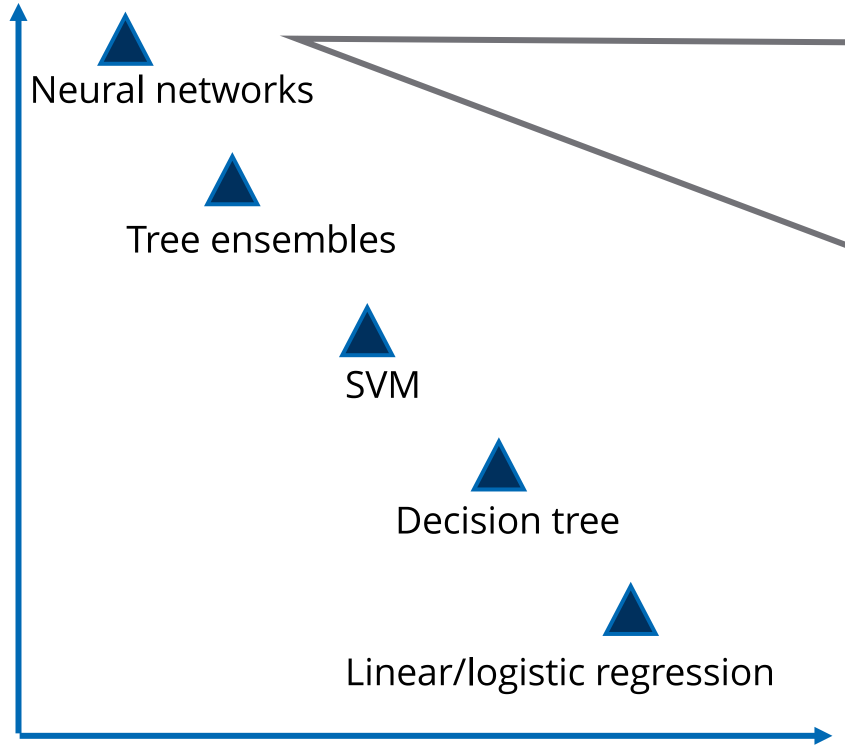

Model Interpretability

Artificial neural network

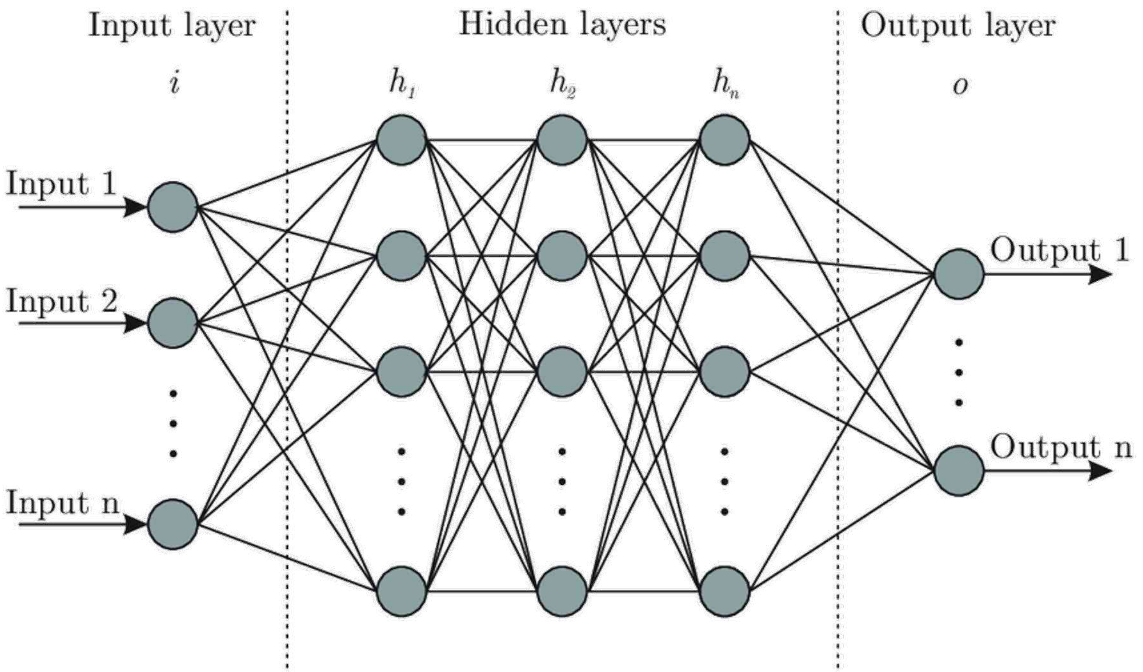

https://images.app.goo.gl/JGUGviUAbcY7henZ6

1) Symbolic AI and rule-based automation
2) (Shallow) machine learning
3) **Neural networks and deep learning**
4) Foundation models and generative AI
5) Tool learning and agentic AI

Artificial neural networks (ANNs)

The general structure

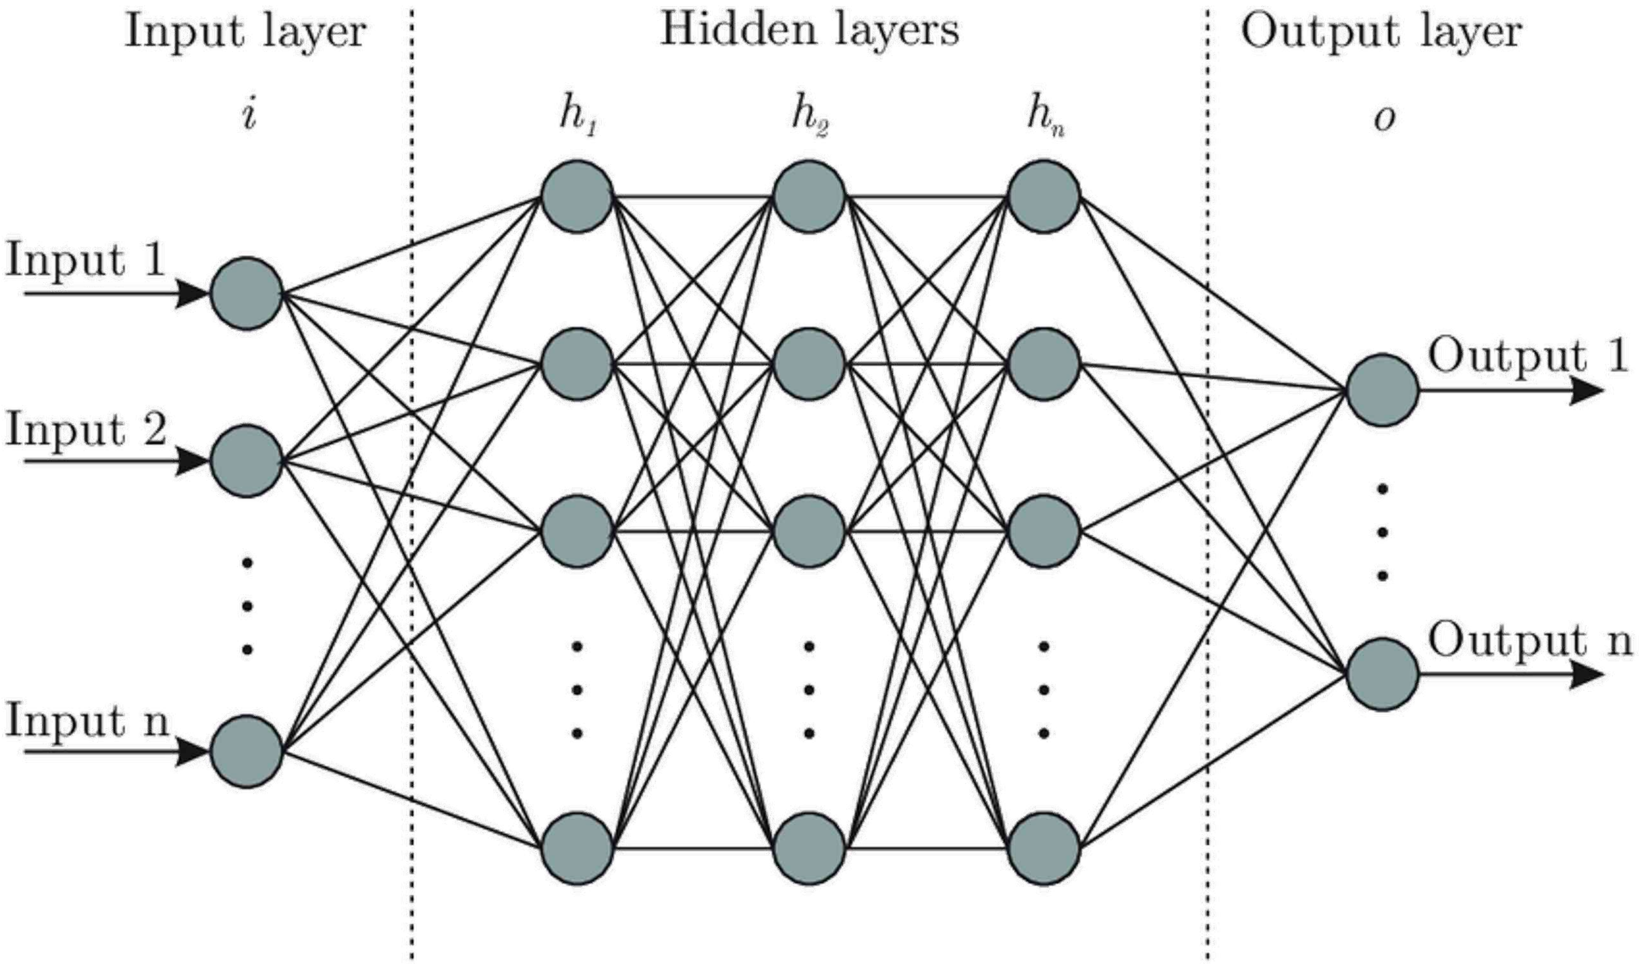

Artificial neural networks (ANNs)

Mathematics of a single neuron

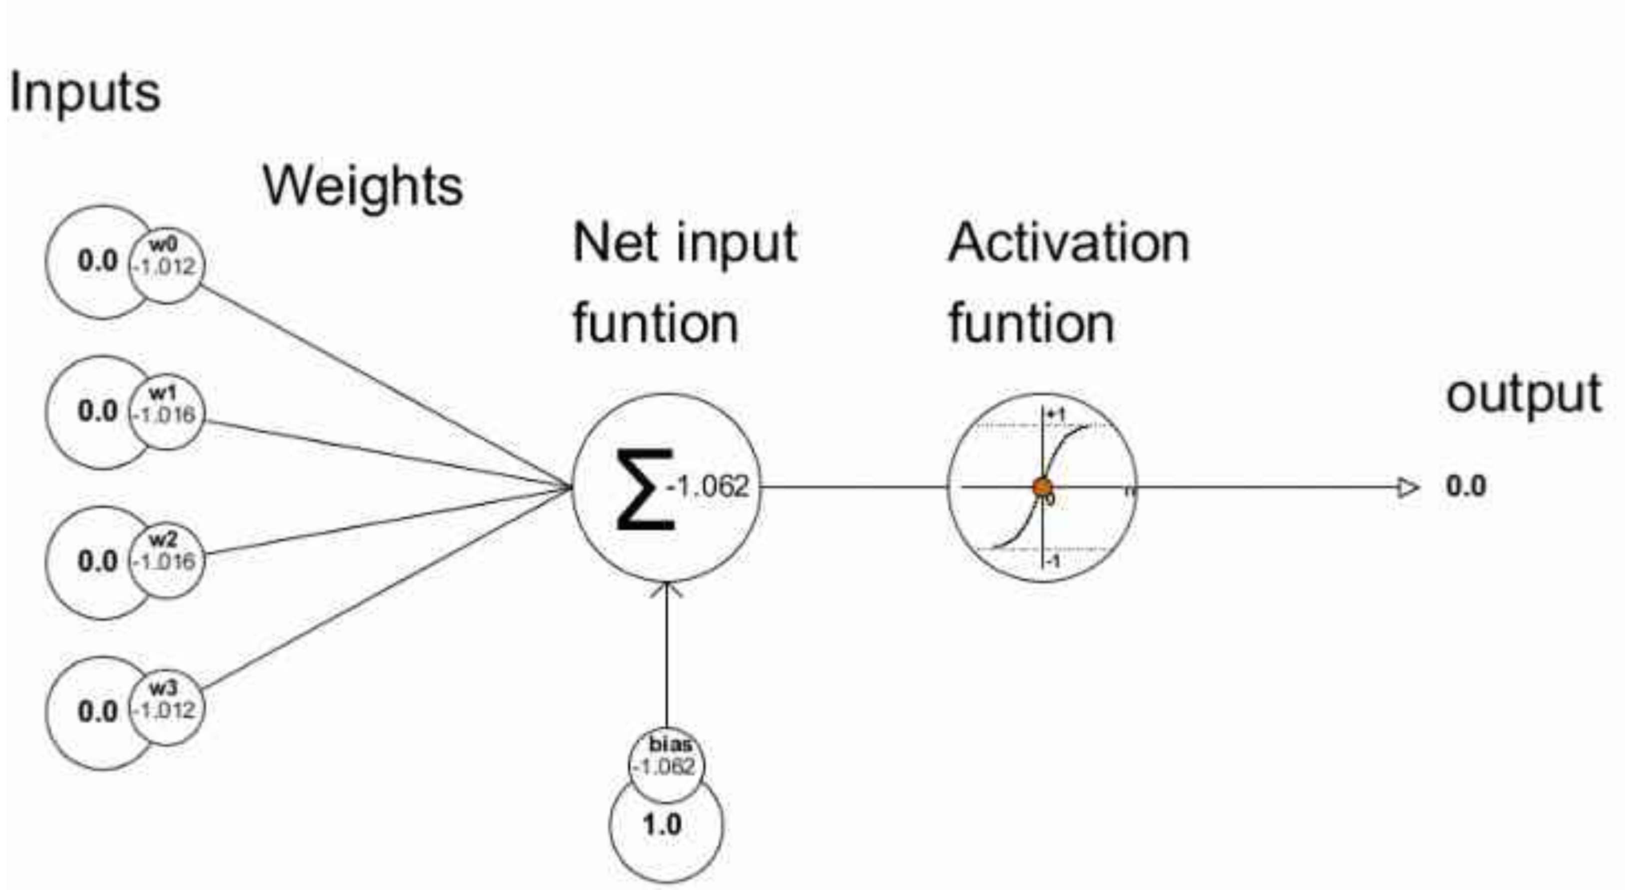

https://www.linkedin.com/feed/update/urn:li:activity:7004139140158009344?utm_source=share&utm_medium=member_android

Artificial neural networks (ANNs)

The neural network zoo

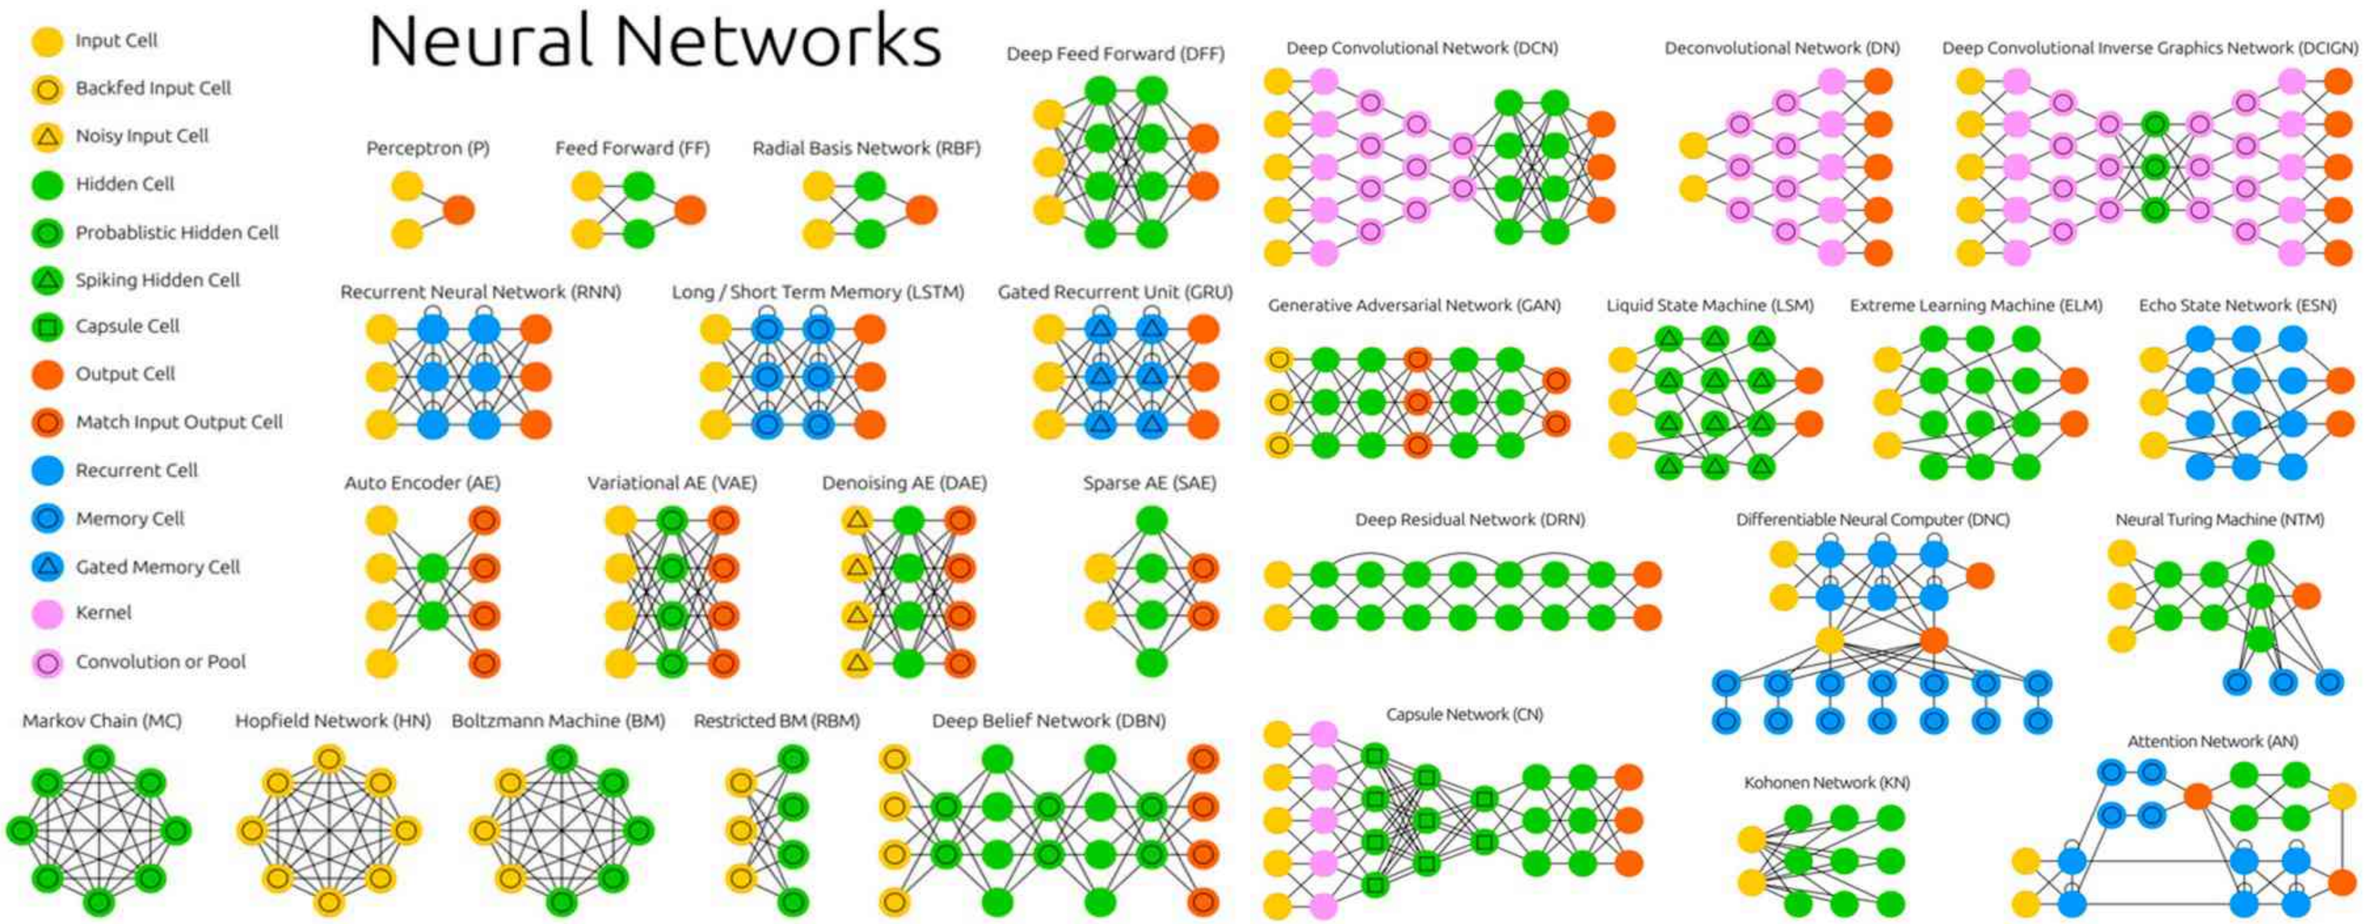

https://doi.org/10.3390/proceedings47010009

Flexibility of neural networks for different learning tasks

Application examples

(Deep) supervised learning

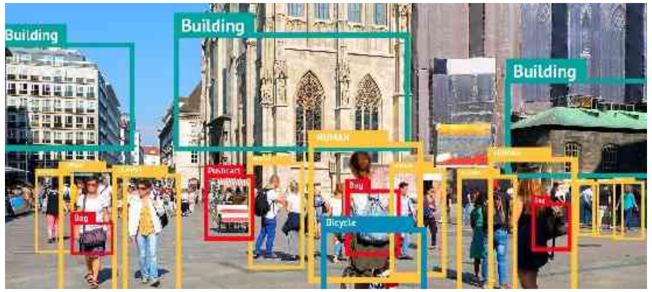

https://images.app.goo.gl/JJzgipy7CWMBNxwx6

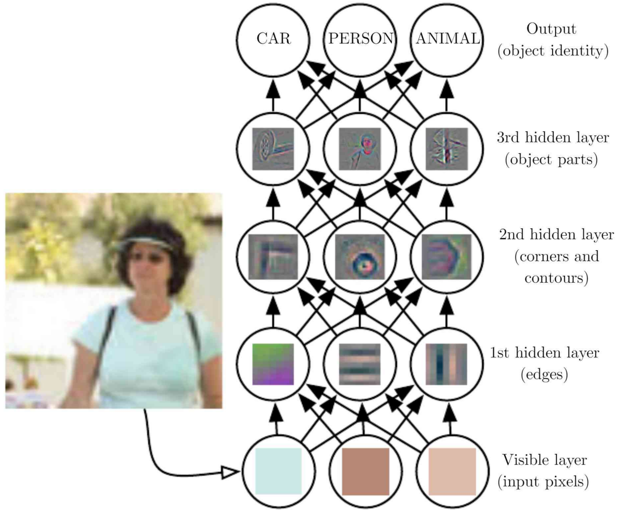

(Deep) unsupervised learning

https://images.app.goo.gl/sRaHue4TaNGuP5PQ9

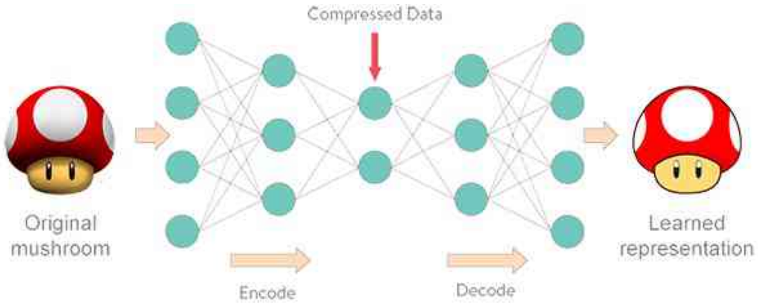

(Deep) reinforcement learning

https://images.app.goo.gl/A9PxDzq5FDcmT7Tn7

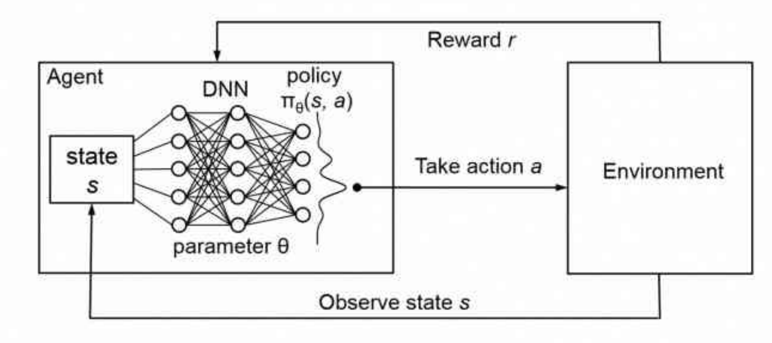

A playful demonstration example

https://playground.tensorflow.org/

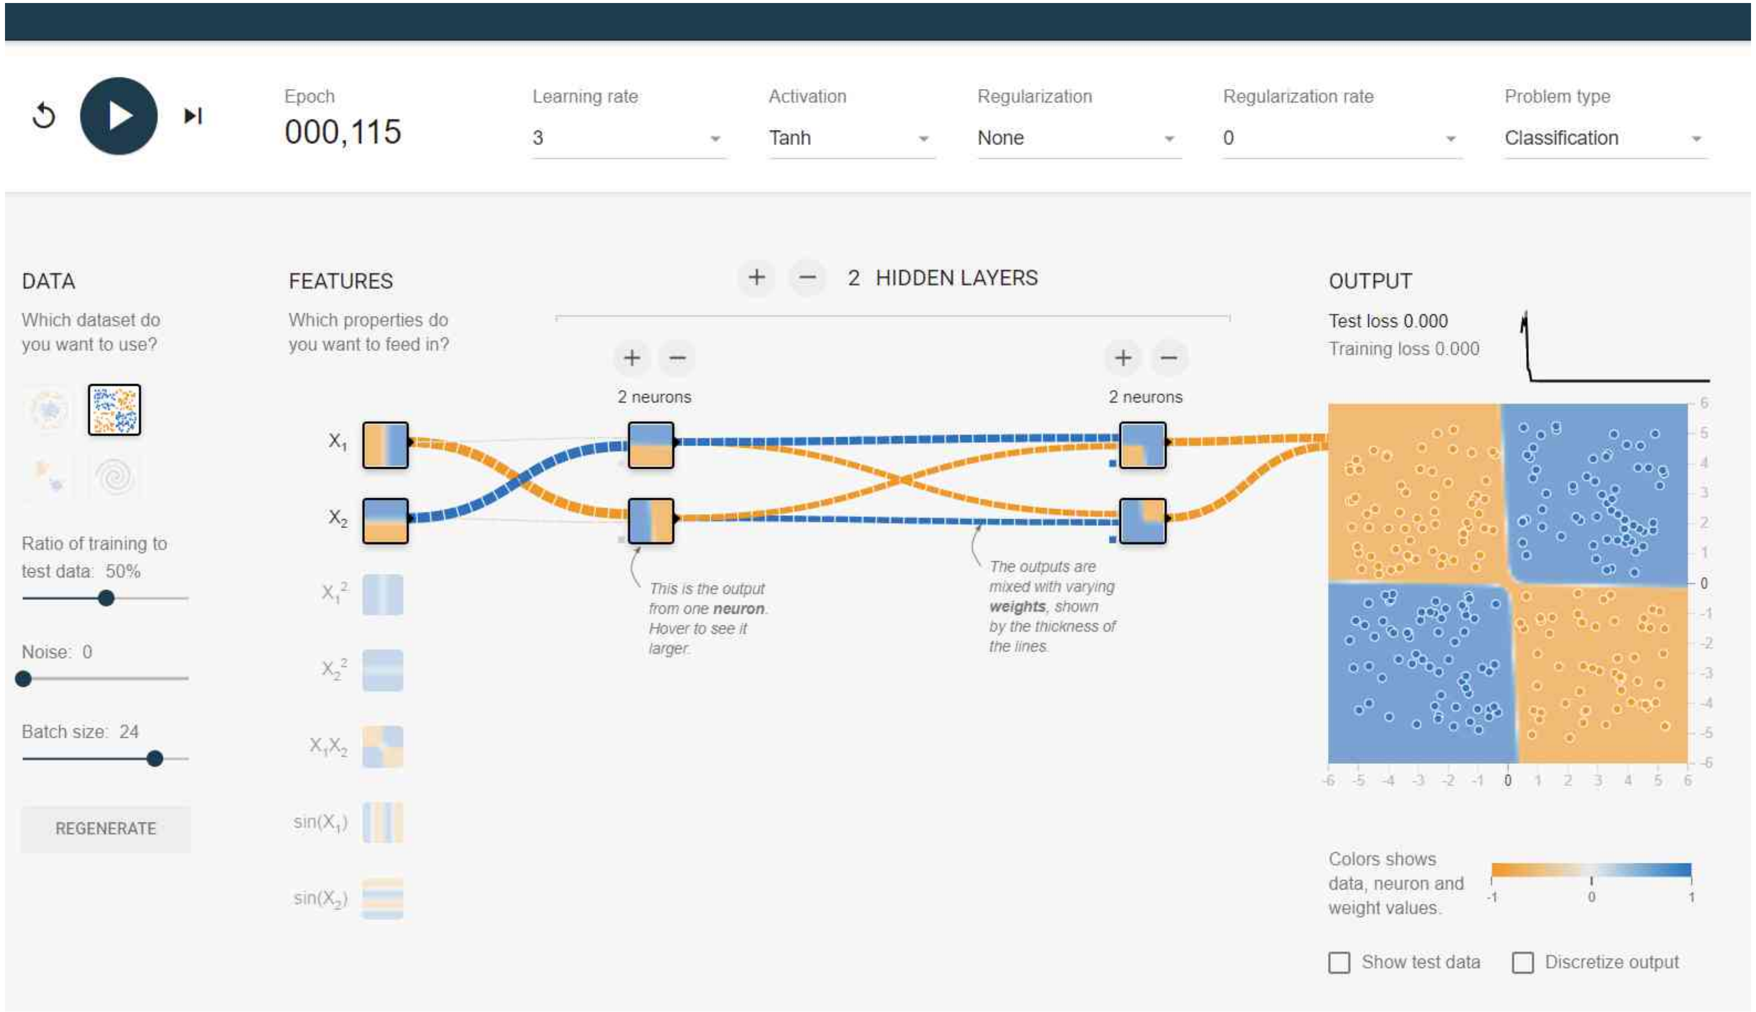

A playful demonstration example

https://playground.tensorflow.org/

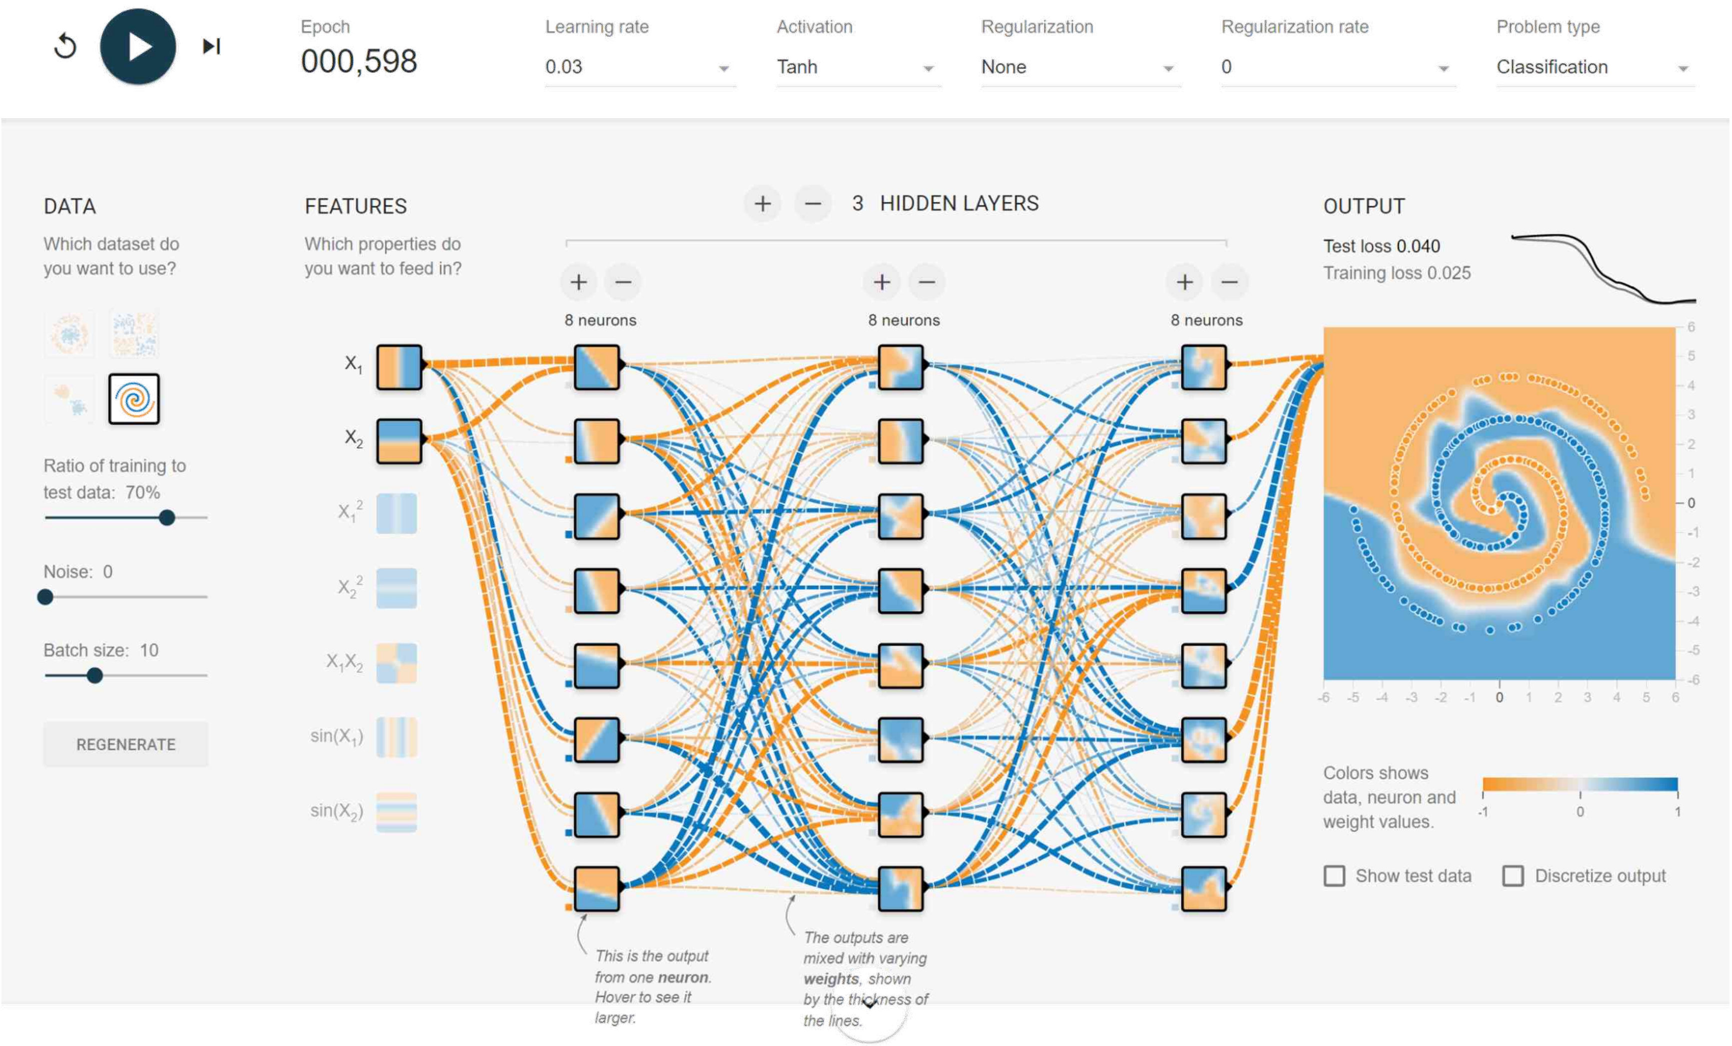

The superpower of deep neural networks

Deep learning = automated feature learning

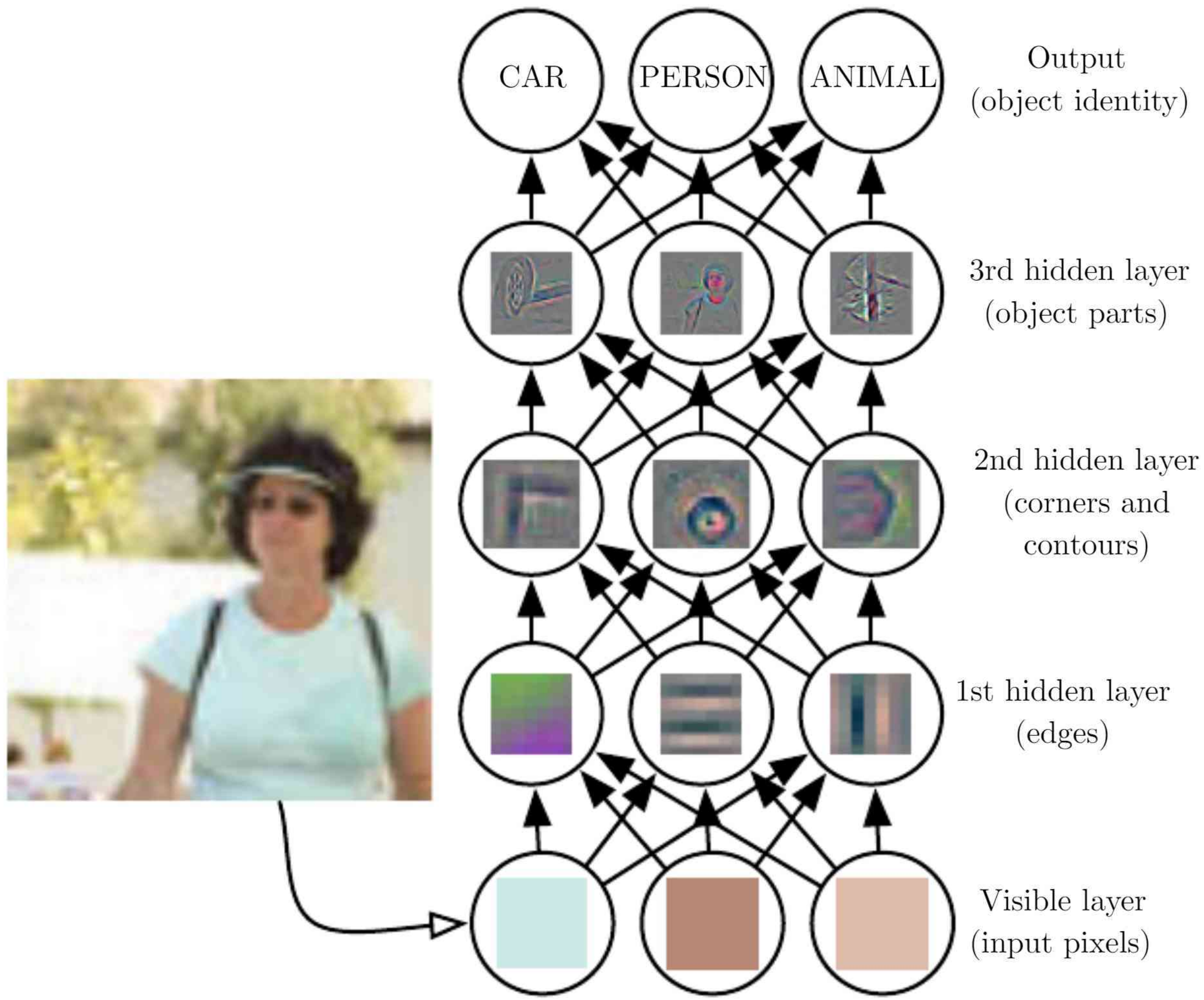

Goodfellow et al. (2016)

Deep learning = learning of increasingly complex patterns

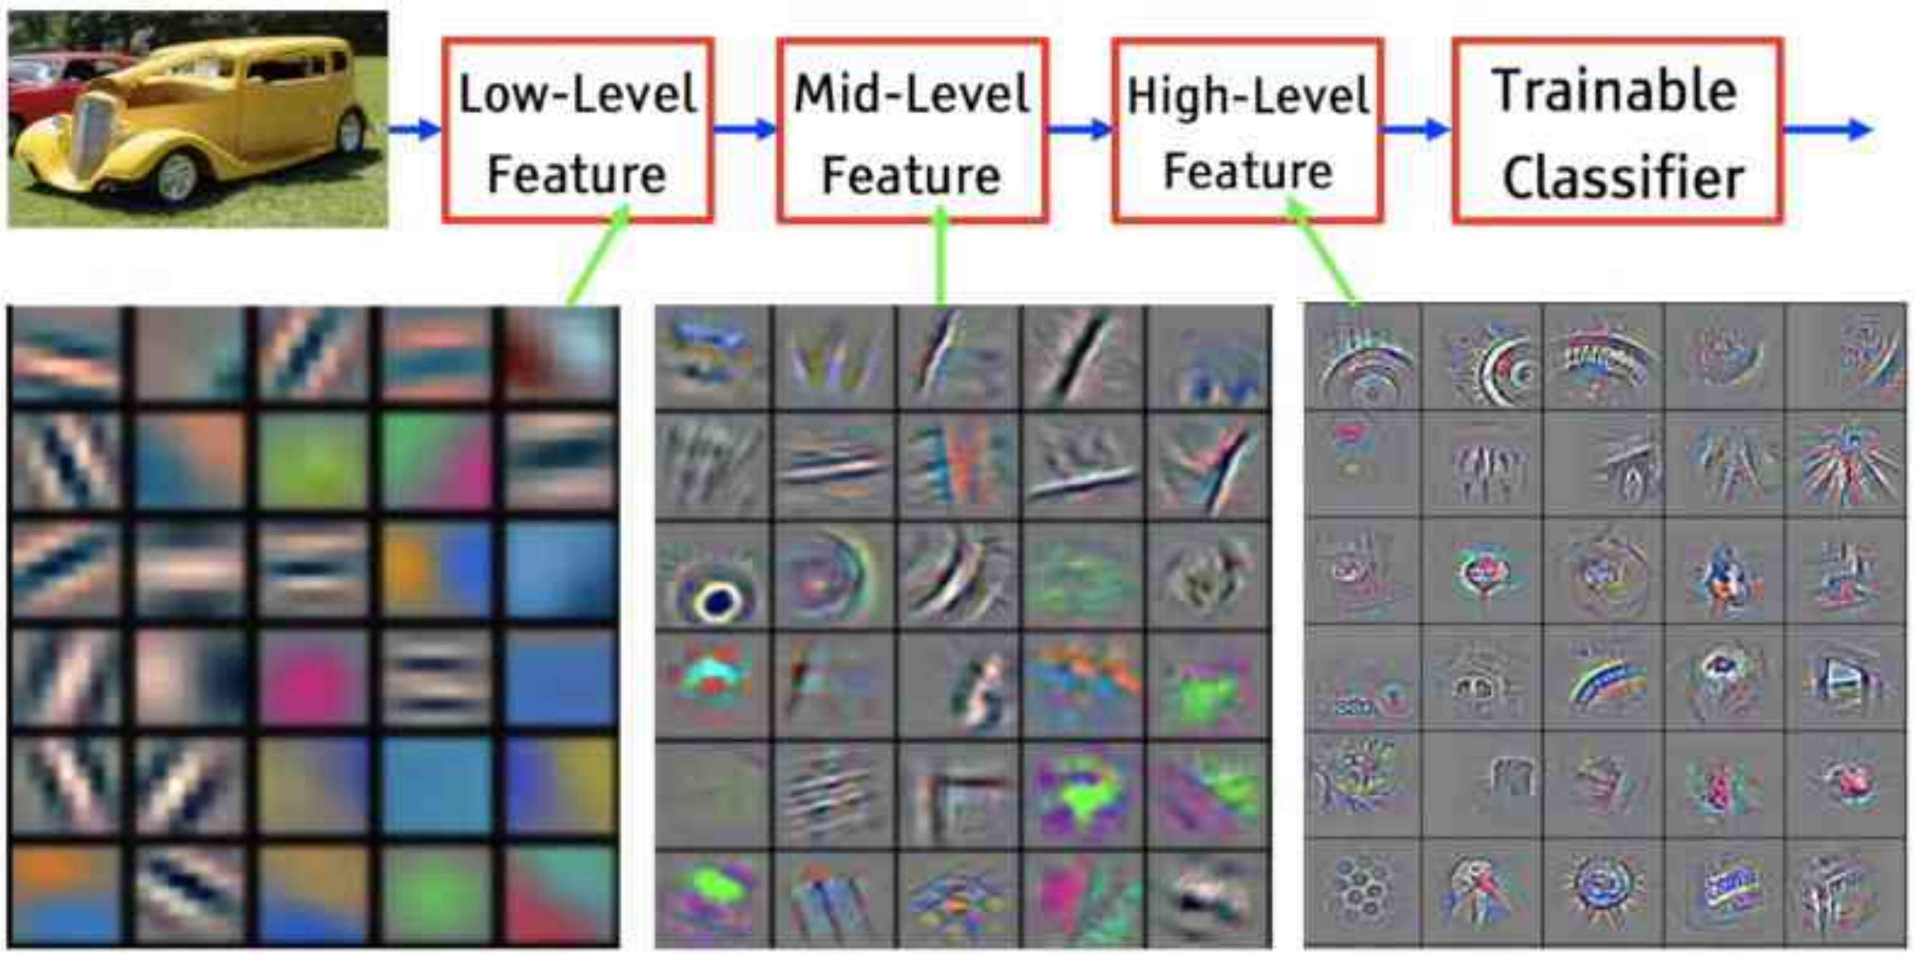

Zeiler & Fergus (2013)

Before the era of deep learning, developers defined patterns manually

Hand-crafted features / patterns

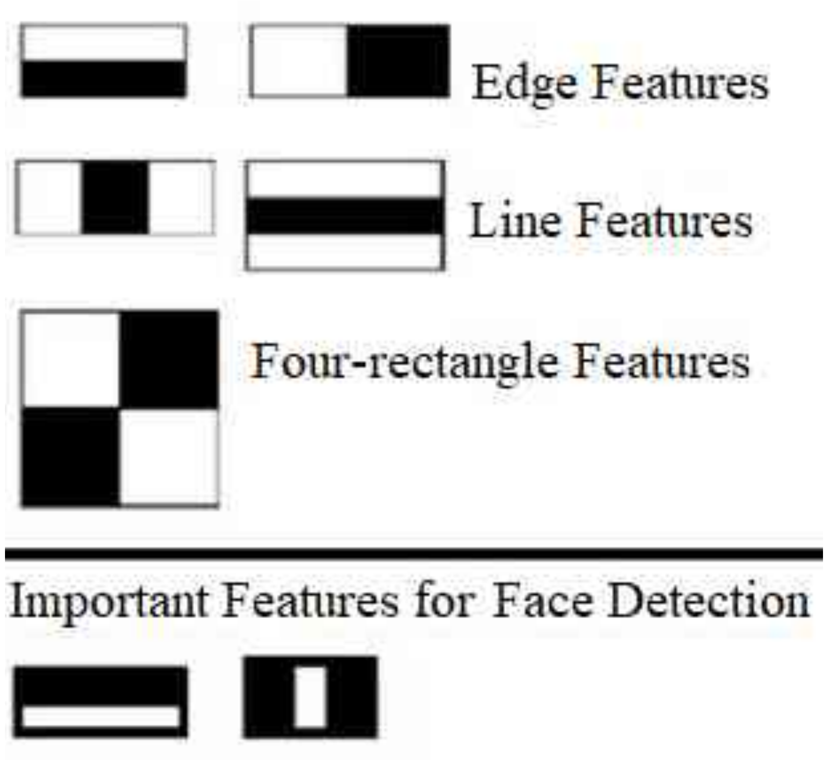

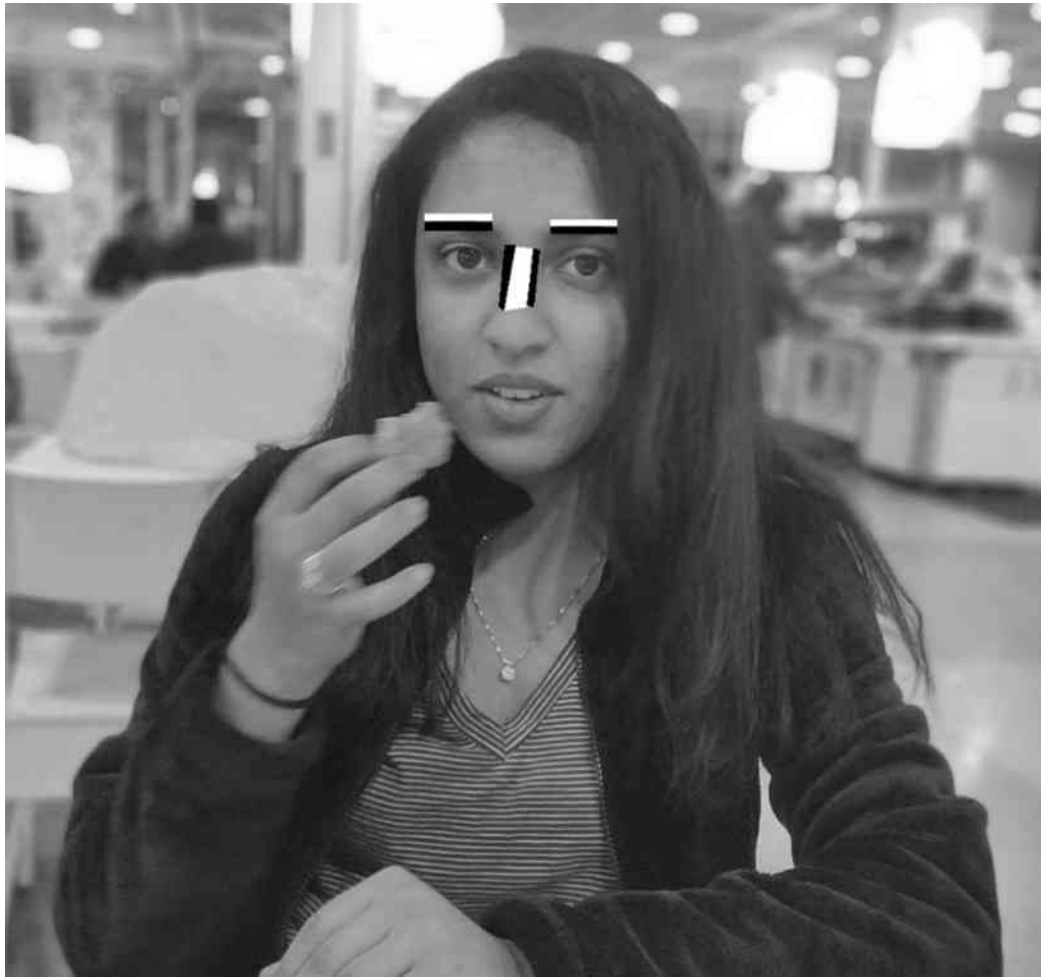

https://towardsdatascience.com/the-intuition-behind-facial-detection-the-viola-jones-algorithm-29d9106b6999

The power of increasingly large deep learning models

Puppy or muffin?

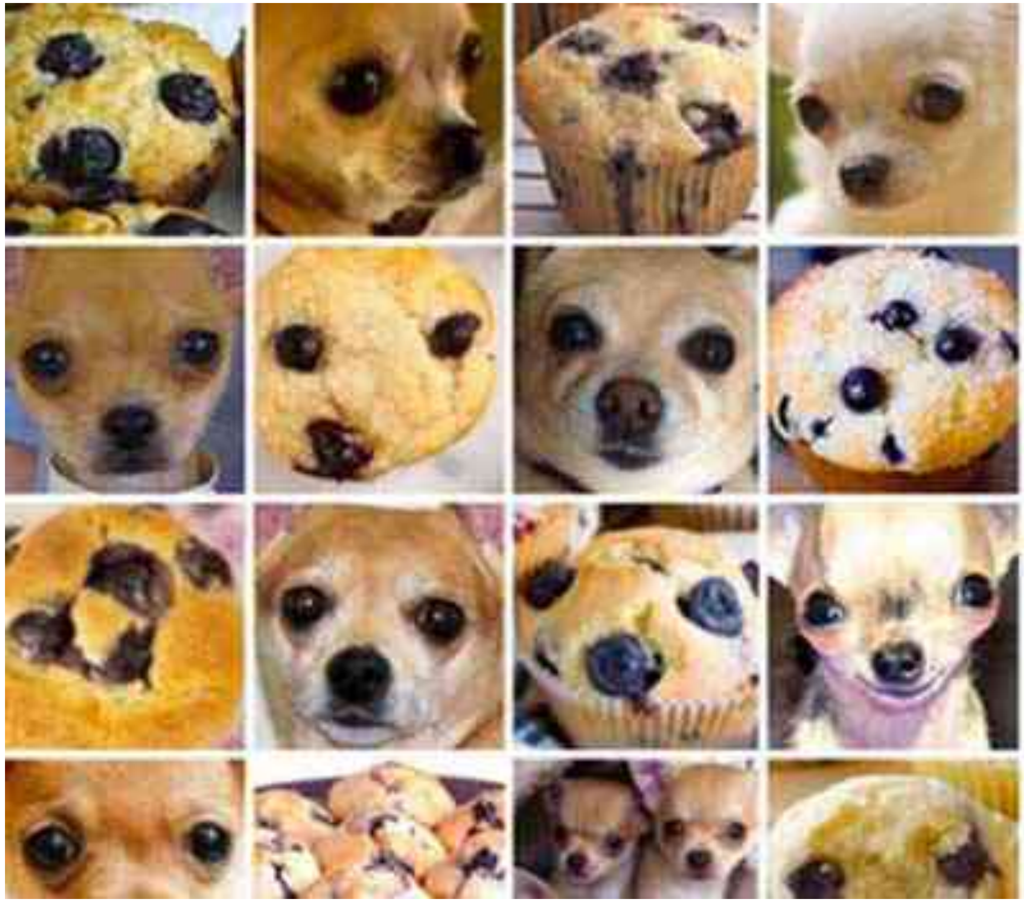

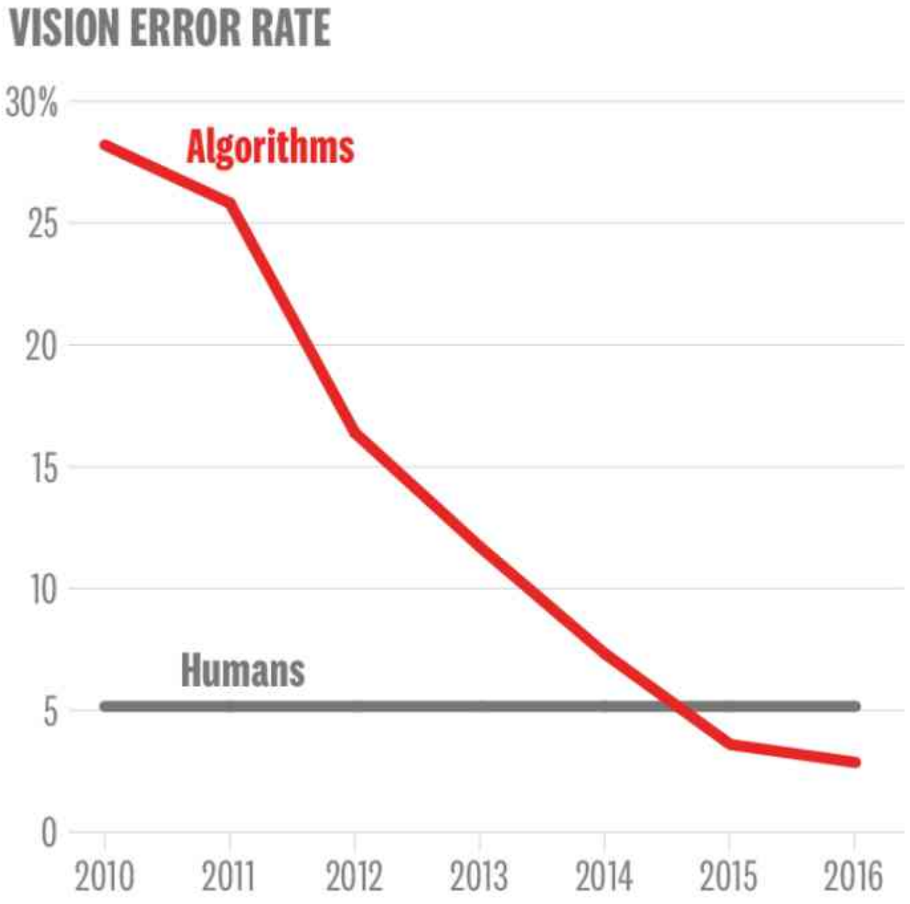

Brynjolfsson, E. and Mcafee A. "The business of artificial intelligence." *Harvard business review* 7.1 (2017): 1-2.

A computer vision model that reads players‘ faces to flag likely bluffs

https://www.instagram.com/p/Dapvu1TgMX8/?img_index=1&igsh=eWdqbmZzdnRlazh1

How do ChatGPT & Co. work?

**But how does this work
with text data and
natural language processing?**

For example, with ChatGPT?

How do ChatGPT & Co. work?

For a better understanding, we need to know two important concepts

1. Language modeling

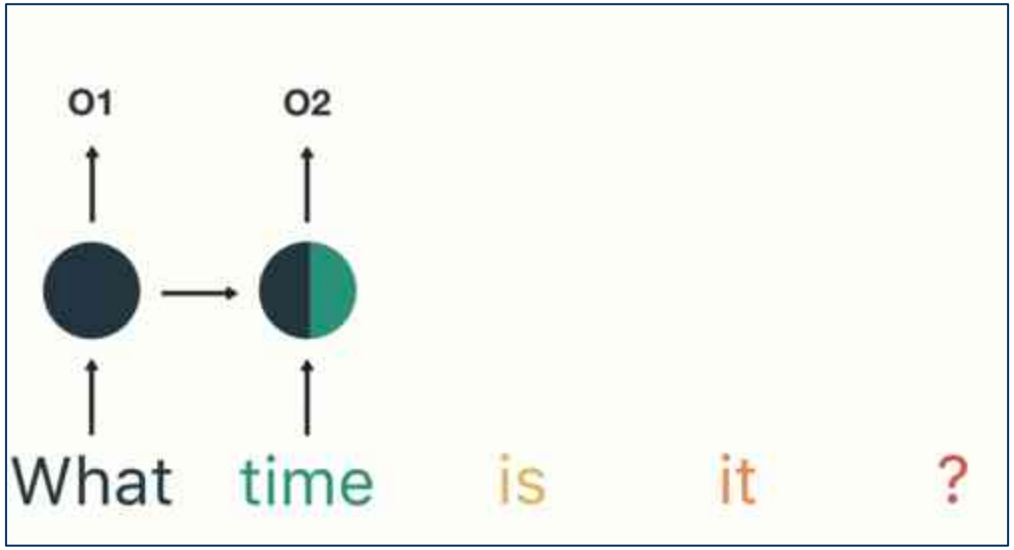

2. Embeddings

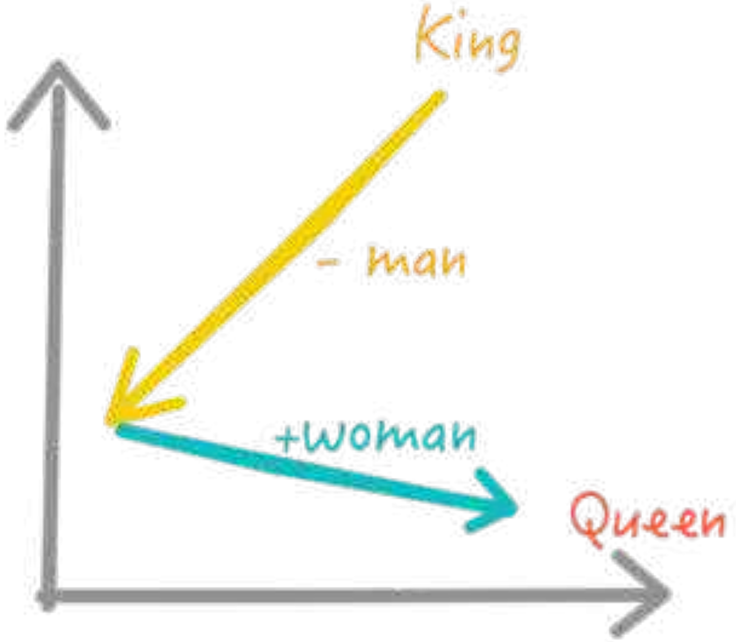

https://www.kdnuggets.com/2021/09/text-preprocessing-methods-deep-learning.html

ChatGPT & Co. are large language models (LLMs)

A **language model** estimates the probability of a word or a sequence of words given a preceding sequence of words.

Binge ... on | - | and | of | is
Binge drinking ... is | and | had | in | was
Binge drinking may ... be | also | have | not | increase
Binge drinking may not ... be | have | cause | always | help
Binge drinking may not necessarily ... be | lead | cause | results | have
Binge drinking may not necessarily kill ... you | the | a | people | your
Binge drinking may not necessarily kill or ... even | injure | kill | cause | prevent
Binge drinking may not necessarily kill or even ... kill | prevent | cause | reduce | injure
Binge drinking may not necessarily kill or even damage ... your | the | a | you | someone
Binge drinking may not necessarily kill or even damage brain ... cells | functions | tissue | neurons
Binge drinking may not necessarily kill or even damage brain cells, ... some | it | the | is | long

This also works for whole sentences and even paragraphs (e.g., questions + answers).

https://doi.org/10.1098/rsos.211837

How can ChatGPT & Co. understand and process natural language?

To a machine, all words actually look the same, because it doesn't really understand their meaning.

Original text

T1: "A woman arrived by car."

T2: "A man arrived by car."

T3: "A man married a woman."

Conversion into a numerical form
(for mathematical operations)

	a	woman	man	arrived	married	by	car
T1	1	1	0	1	0	1	1
T2	1	0	1	1	0	1	1
T3	2	1	1	0	1	0	0

(simple counting of words)

Representation of individual words in 3D vector space
(only for woman, man, car)

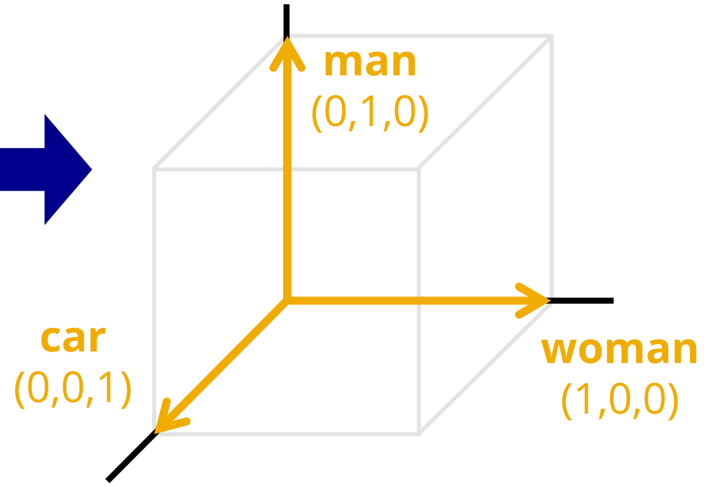

So, how can we extract the meaning of words and/or phrases in a machine-accessible way?

The solution is to create so-called “word embeddings”

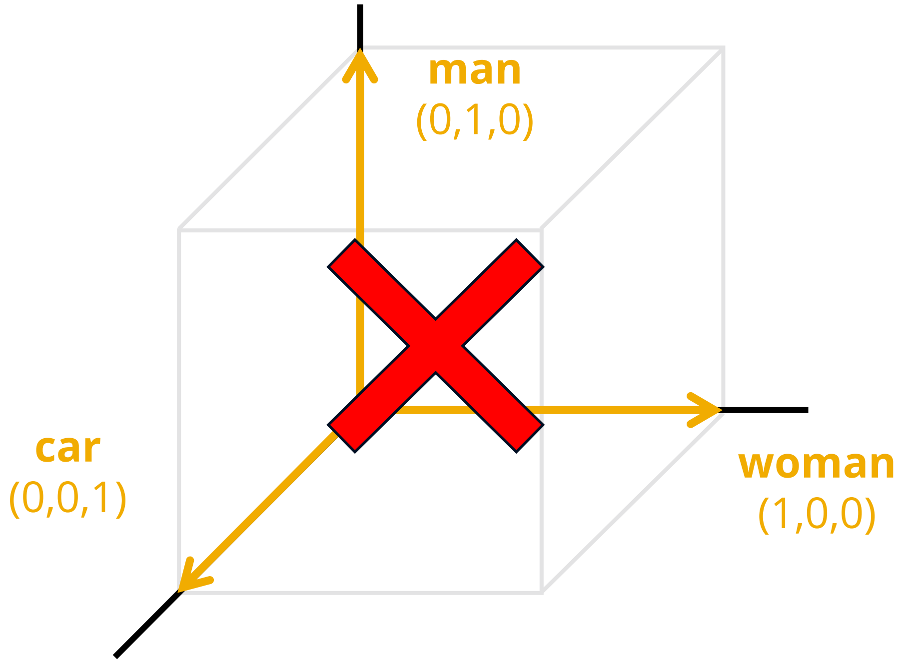

- Words are represented by **vectors (= “word embeddings”)**
- Words with **similar semantics** are **close** to each other (similarity is mathematically measurable)
- This makes it rather simple to **transfer learnings** from one word to a similar word (e.g., if algorithms know that a man can arrive by car, they can easily learn that a woman can also arrive by car)
- **Length** of the **vector representation** is usually much **smaller** than the length of the total vocabulary

The solution is to create so-called “word embeddings”

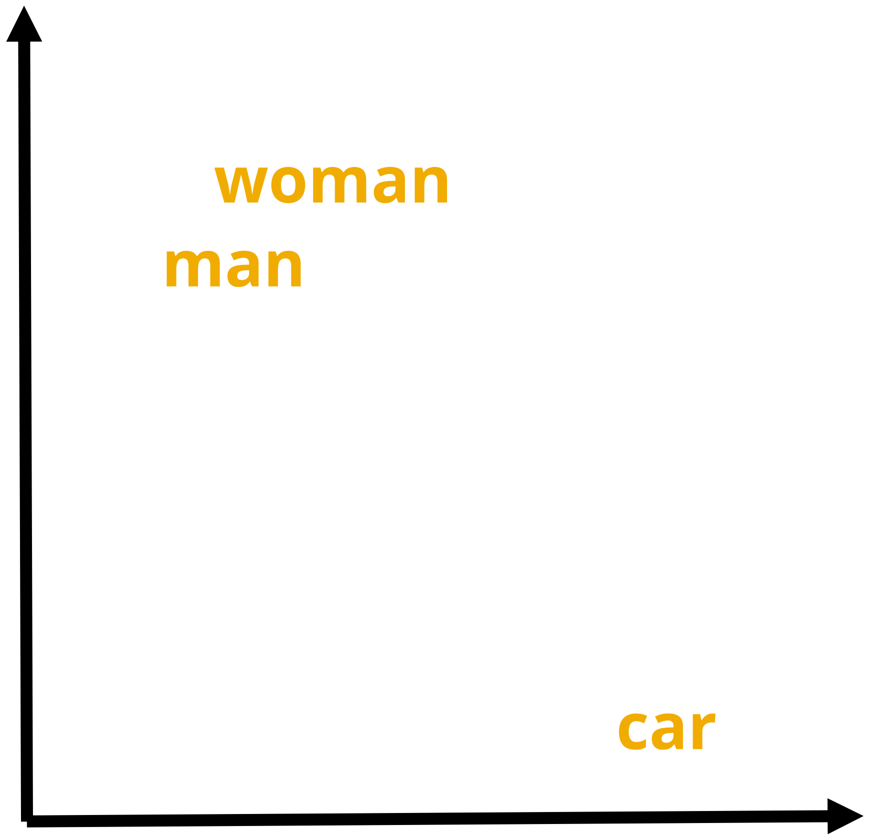

- Words are represented by **vectors** **(= “word embeddings”)**
- Words with **similar semantics** are **close** to each other (similarity is mathematically measurable)
- This makes it rather simple to **transfer learnings** from one word to a similar word (e.g., if algorithms know that a man can arrive by car, they can easily learn that a woman can also arrive by car)
- **Length** of the **vector representation** is usually much **smaller** than the length of the total vocabulary

Some examples of word embeddings

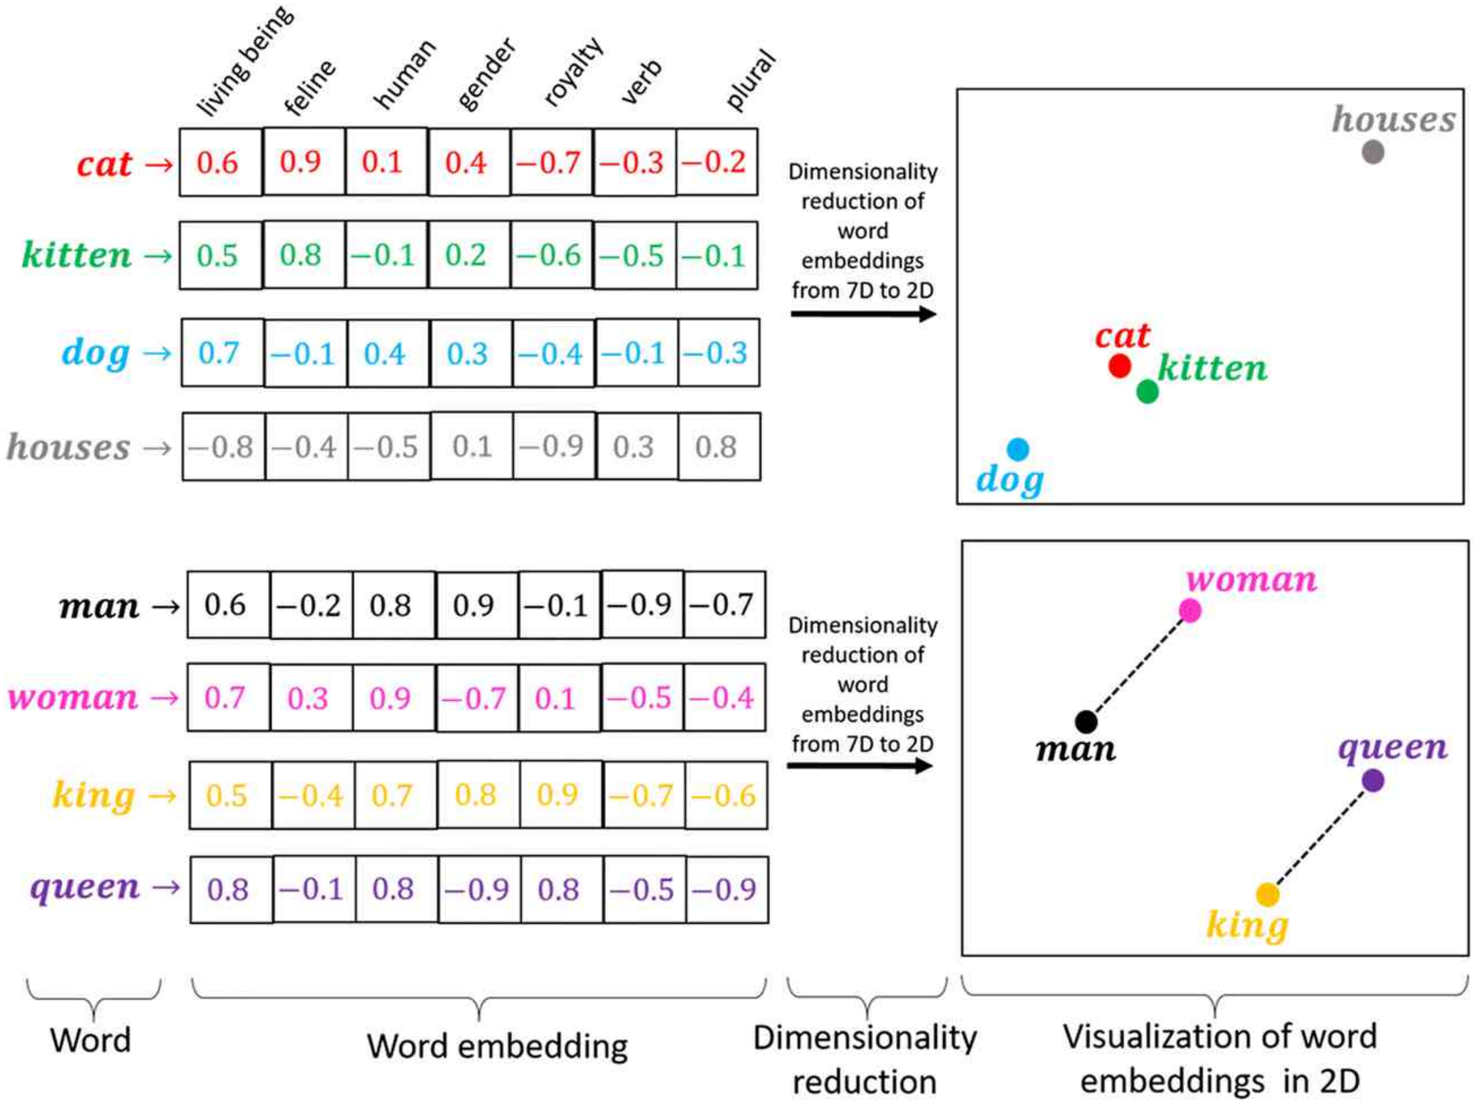

Benefits of word embeddings

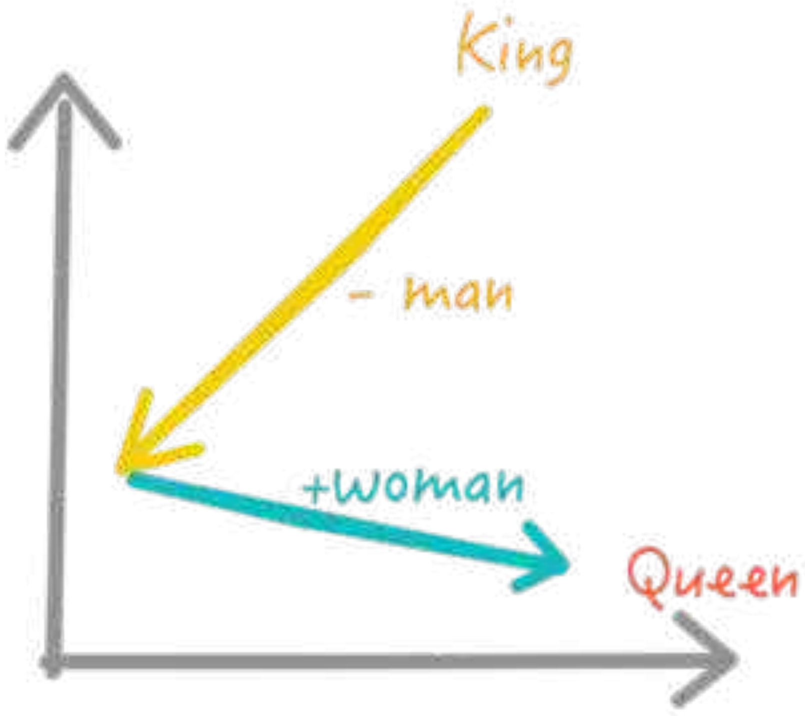

a) Learns Analogy

b) Similar Words have same angles

Benefits of word embeddings

More advanced examples of embeddings

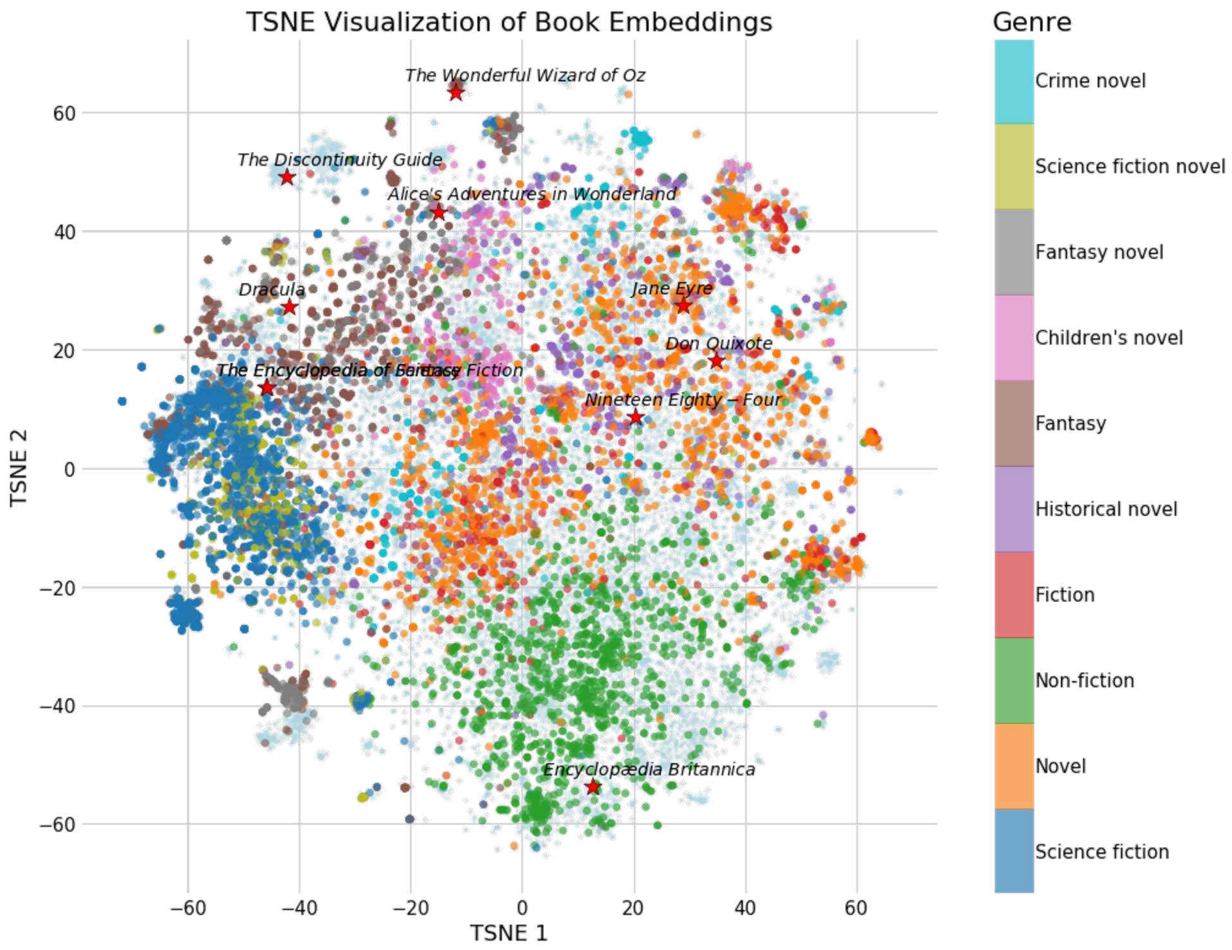

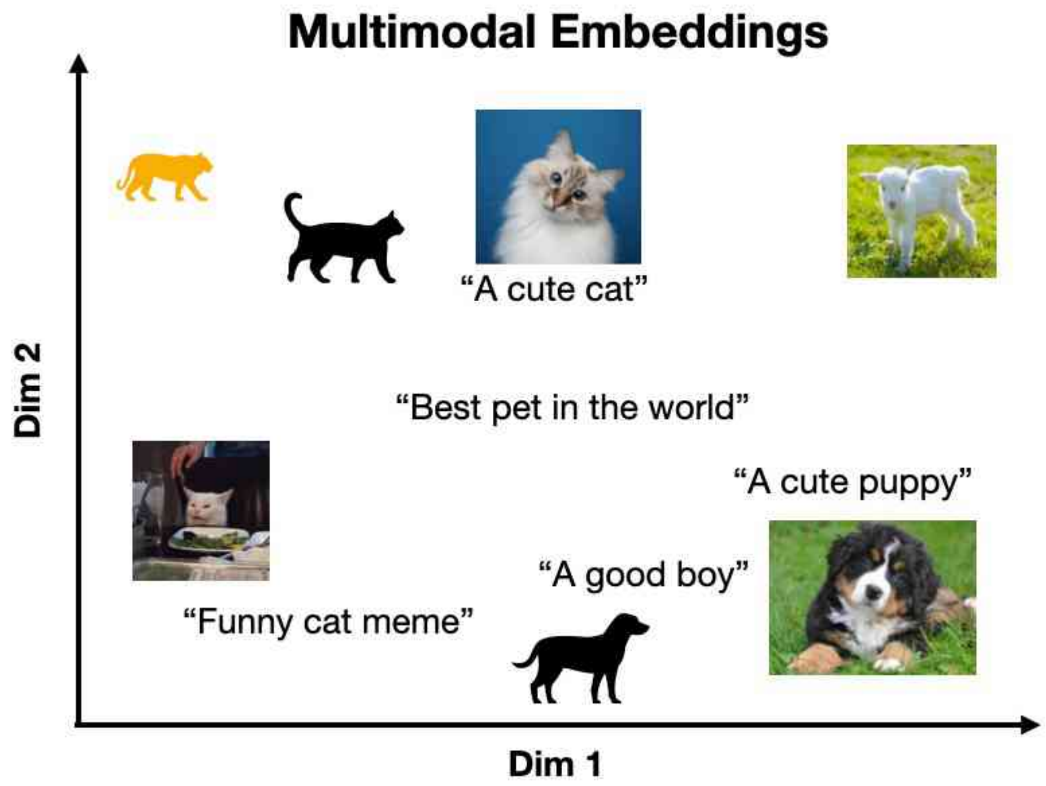

Recap: Examples of fascinating applications

Some examples of supervised learning tasks

Features / input (= x)	Target / output (= y)	Application
Voice recording	Transcript	Speech recognition
Historical market data	Future market data	Trading bots
Photograph	Caption	Image tagging
Drug chemical properties	Treatment efficacy	Pharma R&D
Store transaction details	Is the transaction fraudulent?	Fraud detection
Recipe ingredients	Customer reviews	Food recommendations
Purchase histories	Future purchase behavior	Customer retention
Car locations and speed	Traffic flow	Traffic lights
Faces	Names	Face recognition

1) Symbolic AI and rule-based automation
2) (Shallow) machine learning
3) Neural networks and deep learning
4) **<u>Foundation models and generative AI</u>**
5) Tool learning and agentic AI

From traditional ML to (multi-modal) foundation models

Traditional ML

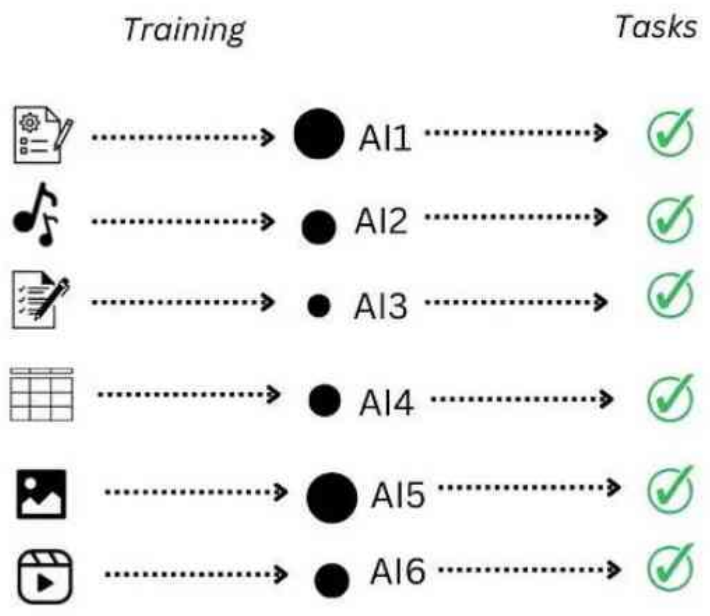

Foundation models

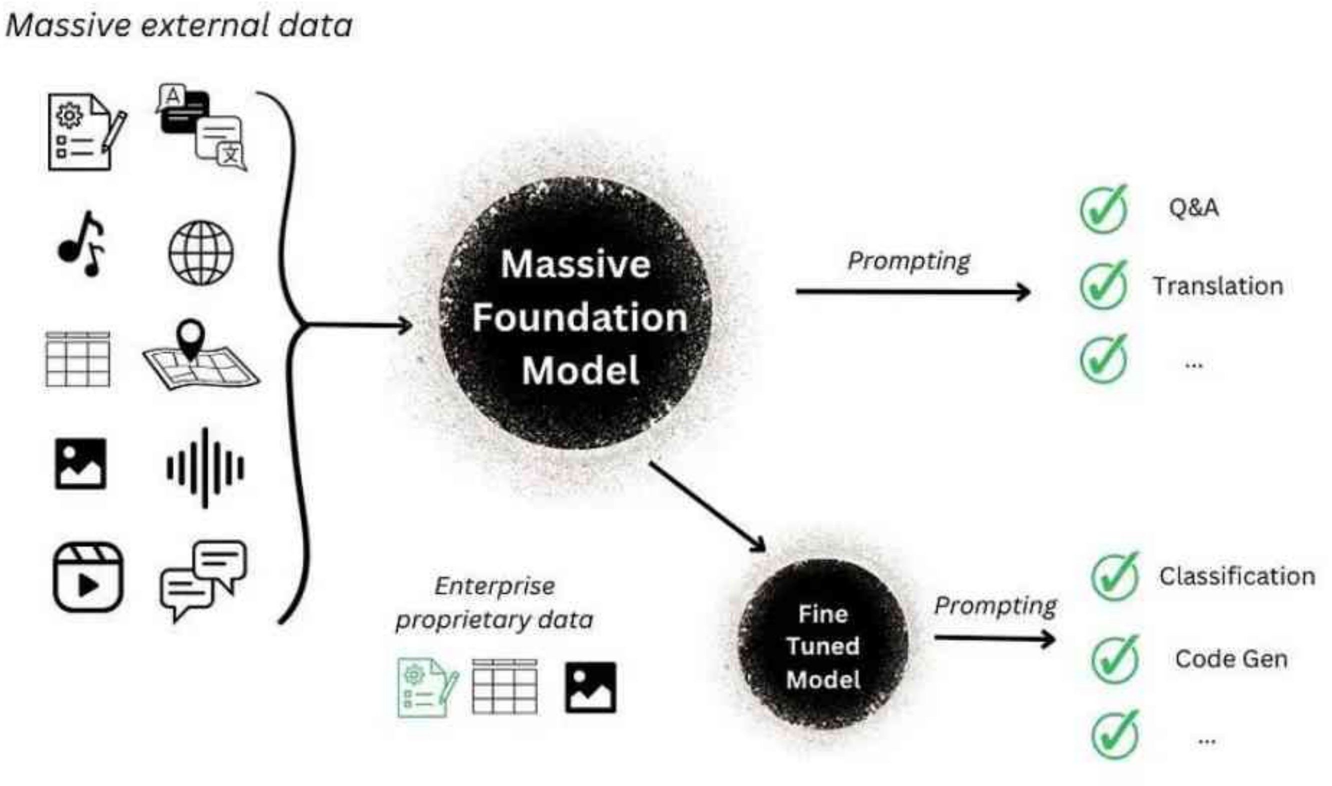

https://medium.com/data-science-collective/what-are-foundation-models-and-why-should-data-scientists-care-c4a3aa7e9589

Foundation models for generalist medical AI

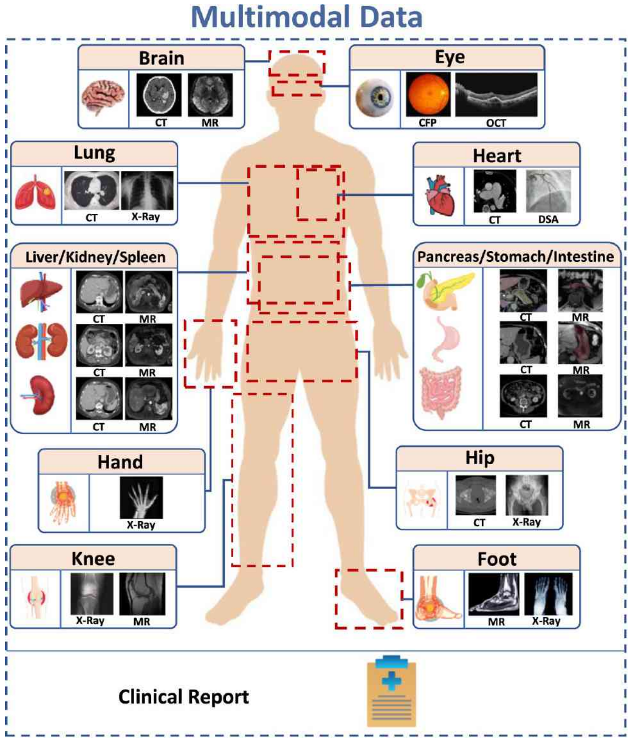

https://www.sciencedirect.com/science/article/pii/S0933365725002003

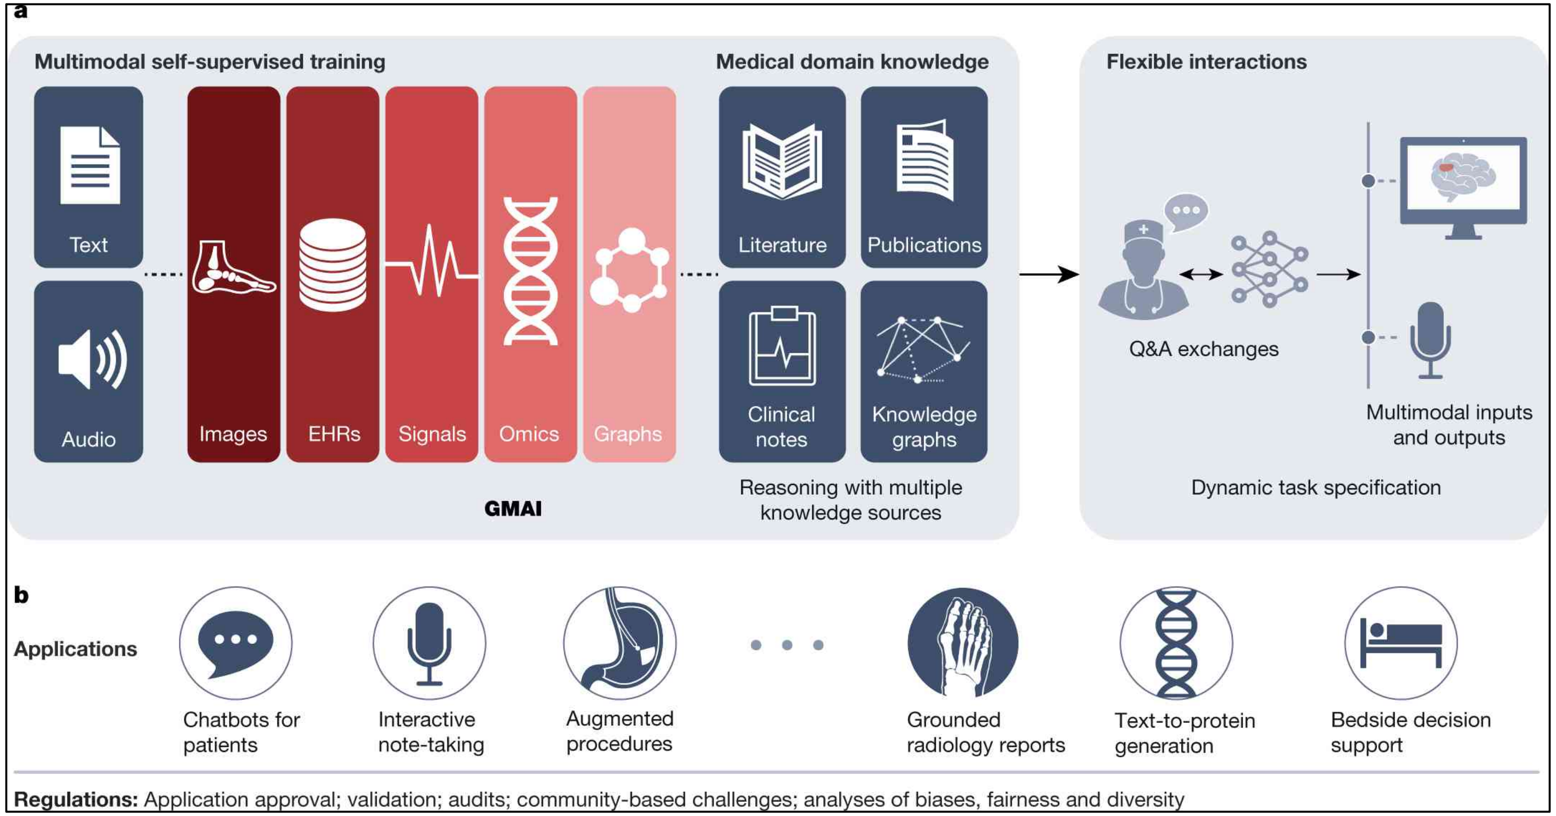

Moor, M., Banerjee, O., Abad, Z.S.H. *et al.* Foundation models for generalist medical artificial intelligence. *Nature* 616, 259–265 (2023). https://doi.org/10.1038/s41586-023-05881-4

Generative AI

Use of foundation models to translate one data modality / representation into another

KHARI JOHNSON BUSINESS JUL 5, 2023 7:00 AM

AI Could Change How Blind People See the World

Assistive technology services are integrating OpenAI's GPT-4, using artificial intelligence to help describe objects and people.

ILLUSTRATION: JACQUI VANLIEW; GETTY IMAGES

https://www.wired.com/story/ai-gpt4-could-change-how-blind-people-see-the-world/

Generative AI
Challenges with biased data

Hungarian has no gendered pronouns, so Google Translate had to make some assumptions

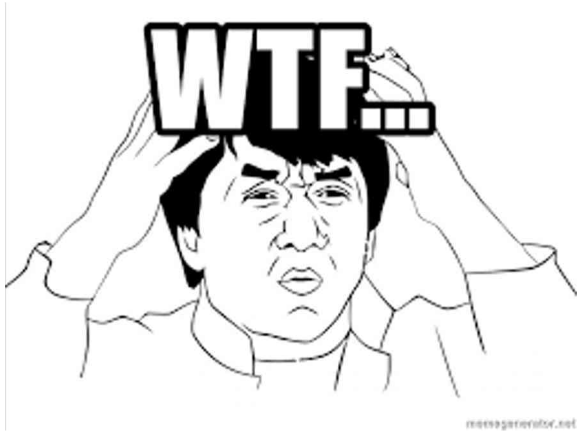

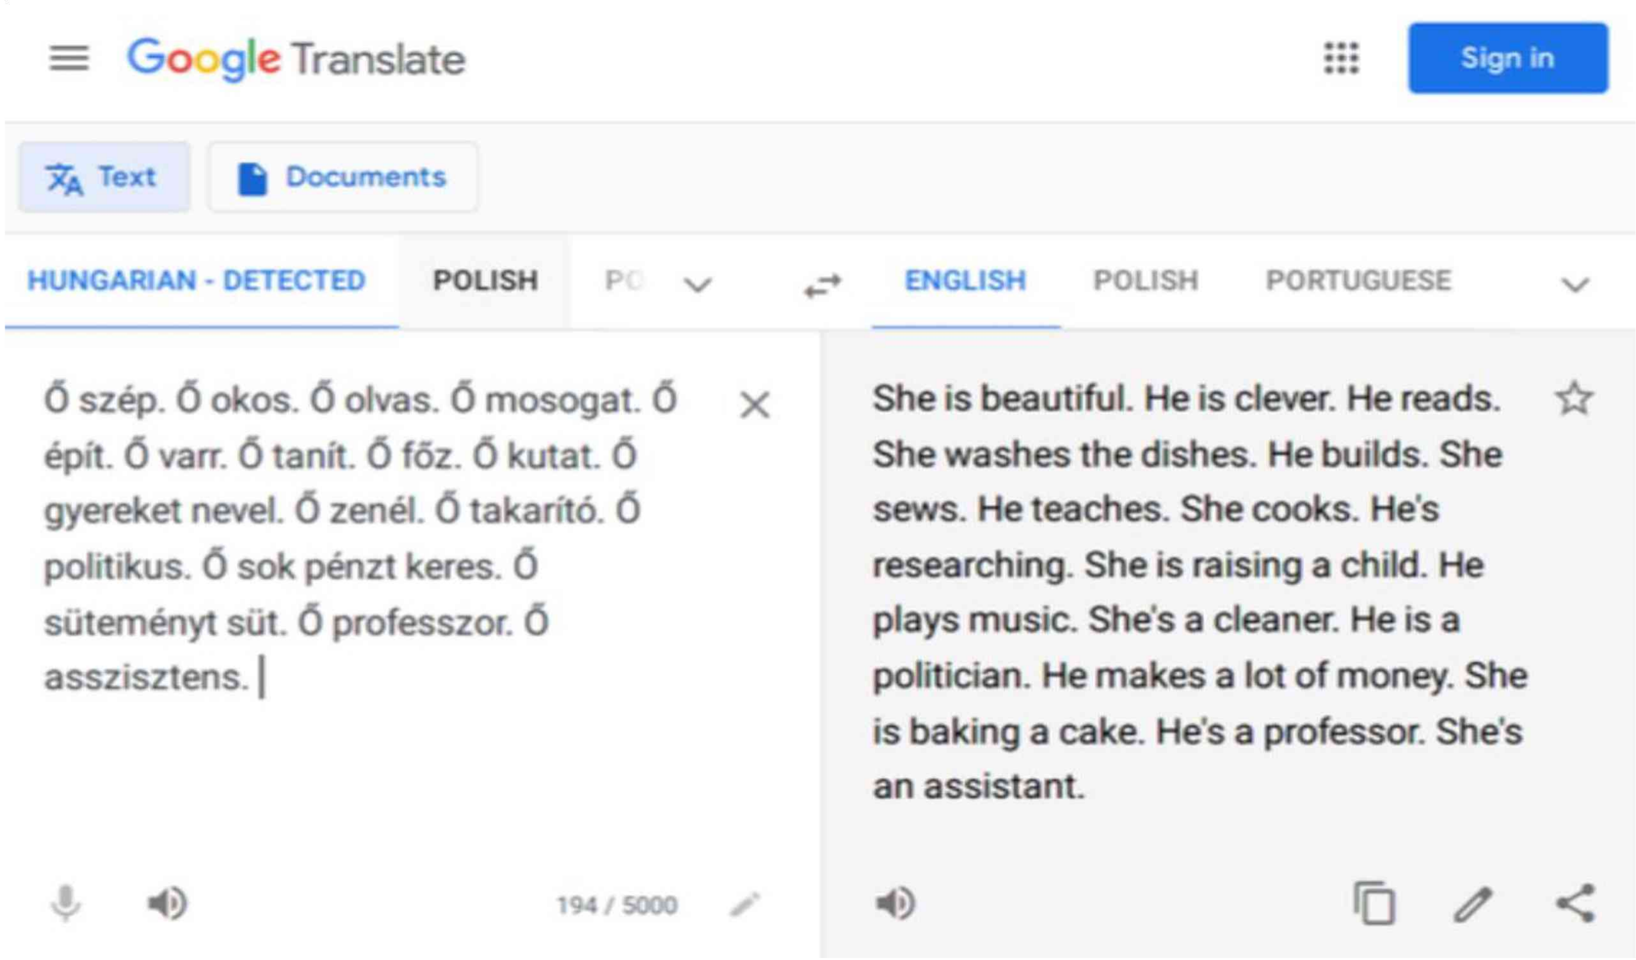

Generative AI
Generating realistic images

Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6(1), 60. https://doi.org/10.1186/s40537-019-0197-0

Generative AI
Generating realistic images

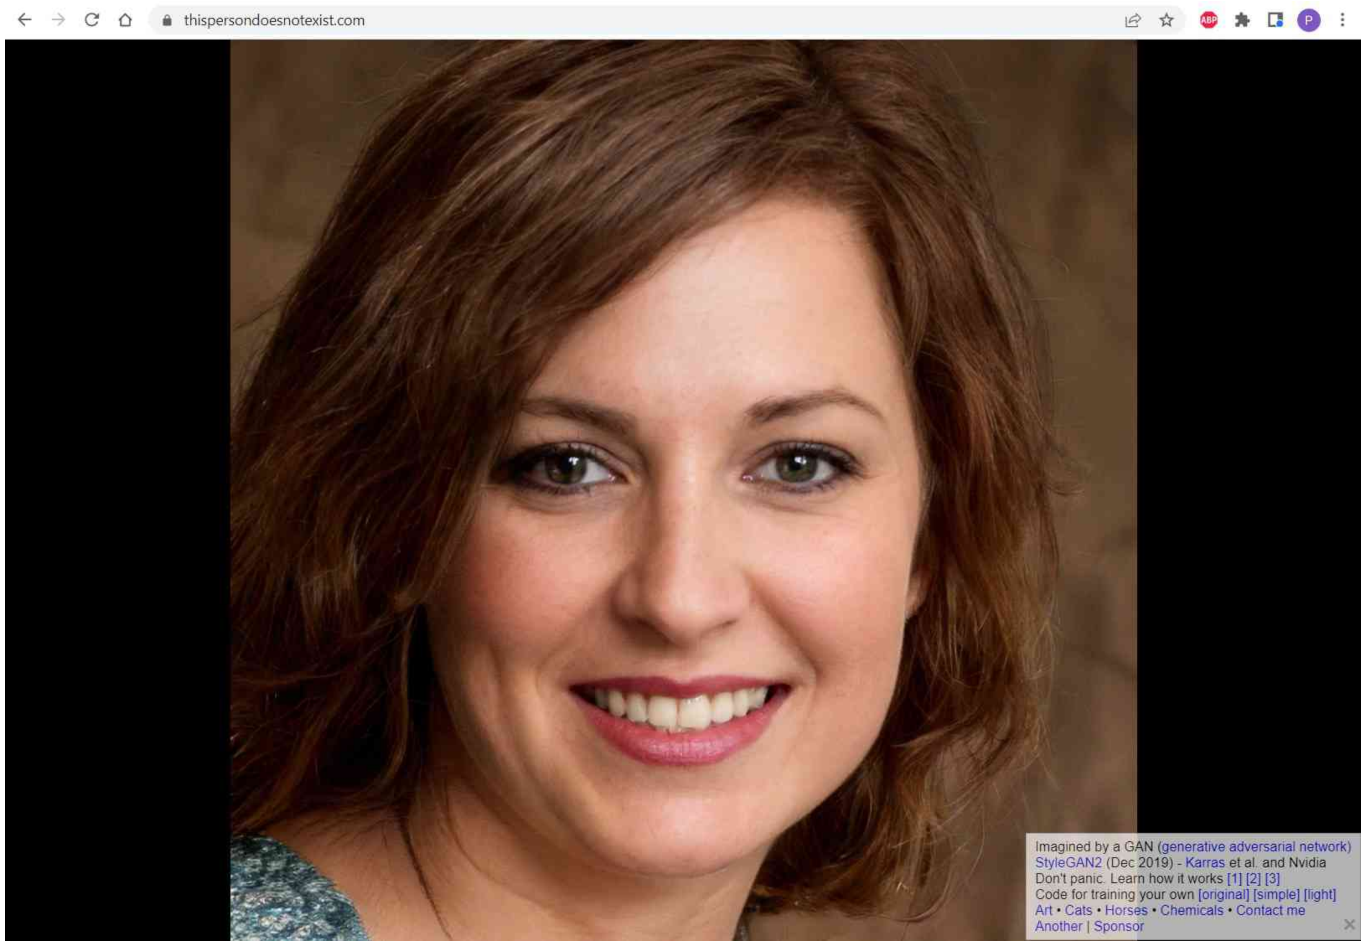

https://thispersondoesnotexist.com/

Generative AI
Generating realistic videos

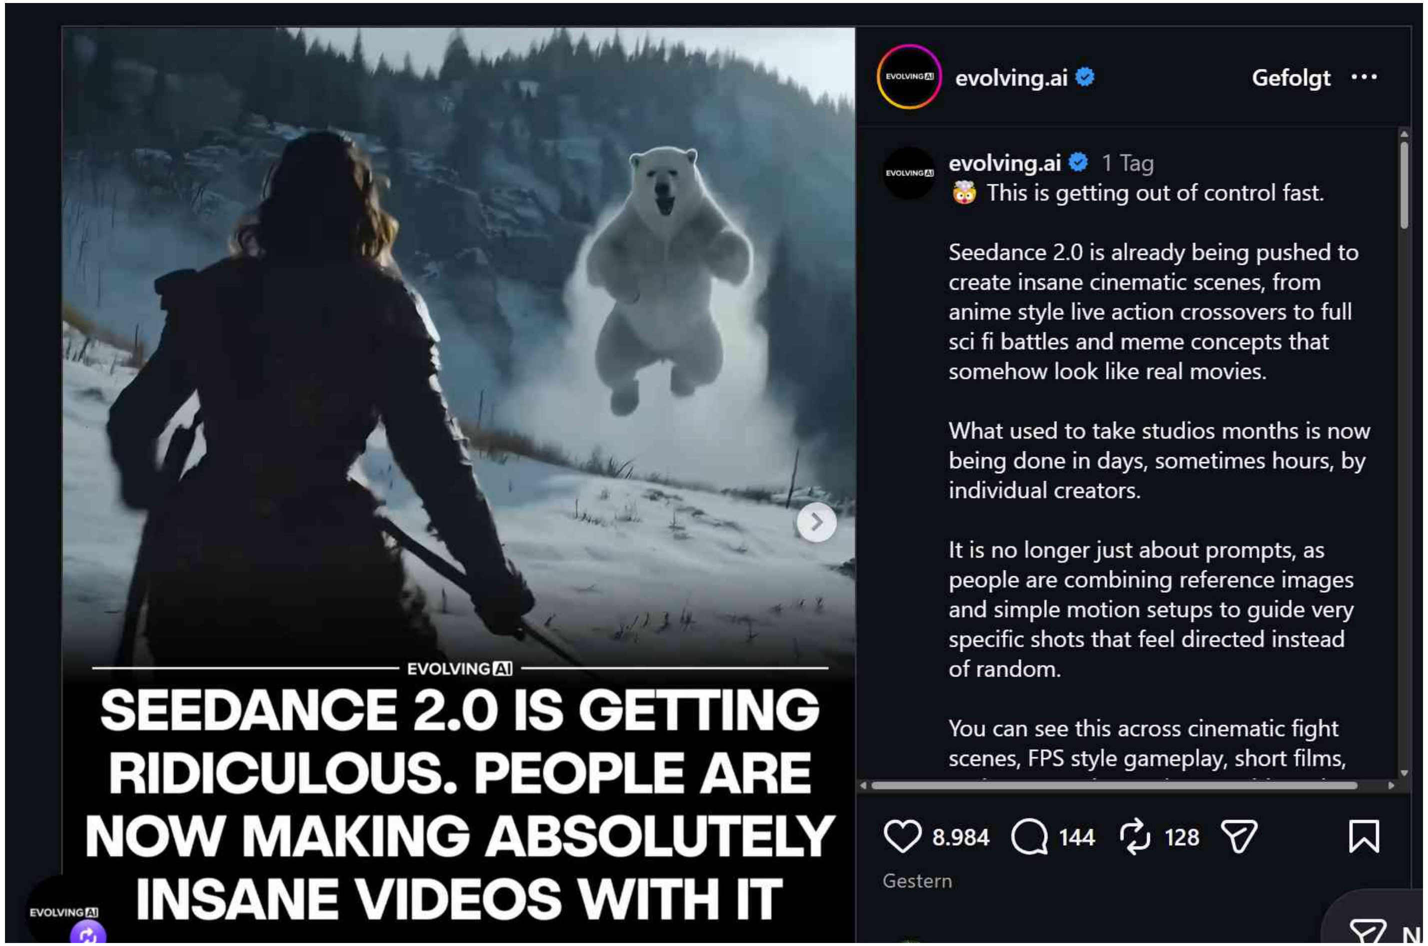

https://www.instagram.com/p/DXE-Nw4jWVw/

1) Symbolic AI and rule-based automation
2) (Shallow) machine learning
3) Neural networks and deep learning
4) Foundation models and generative AI
5) **Tool learning and agentic AI**

Solving math tasks with ChatGPT v1 vs. v5

ChatGPT v1 (back then 2022)

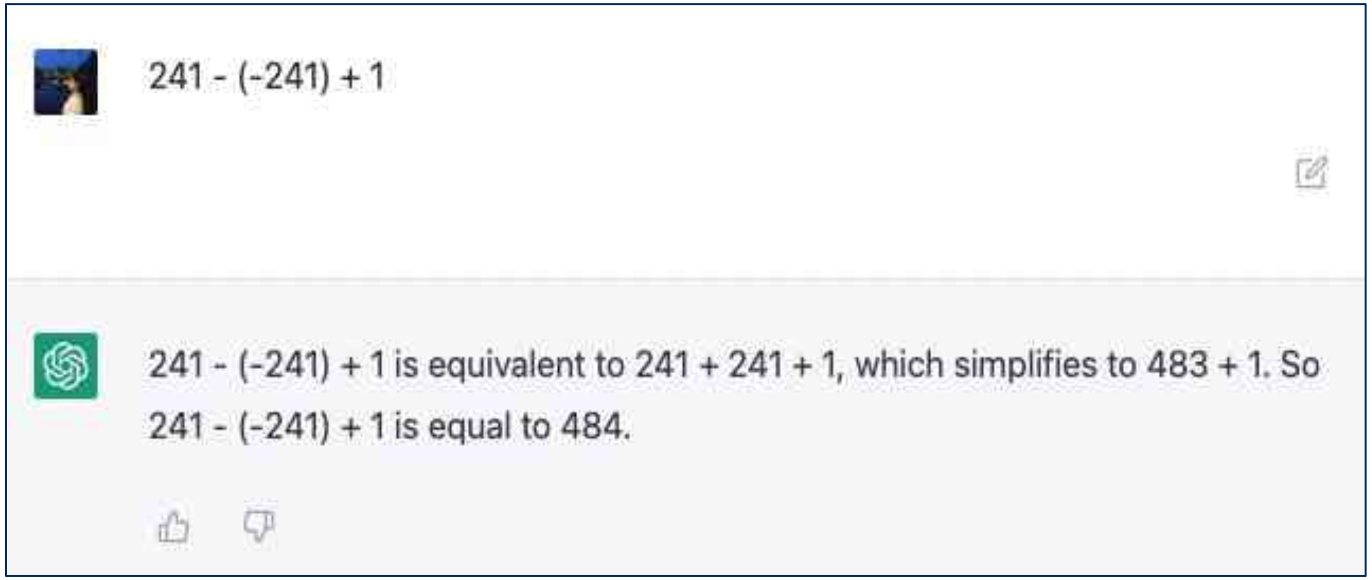

ChatGPT v5 era (today 2026)

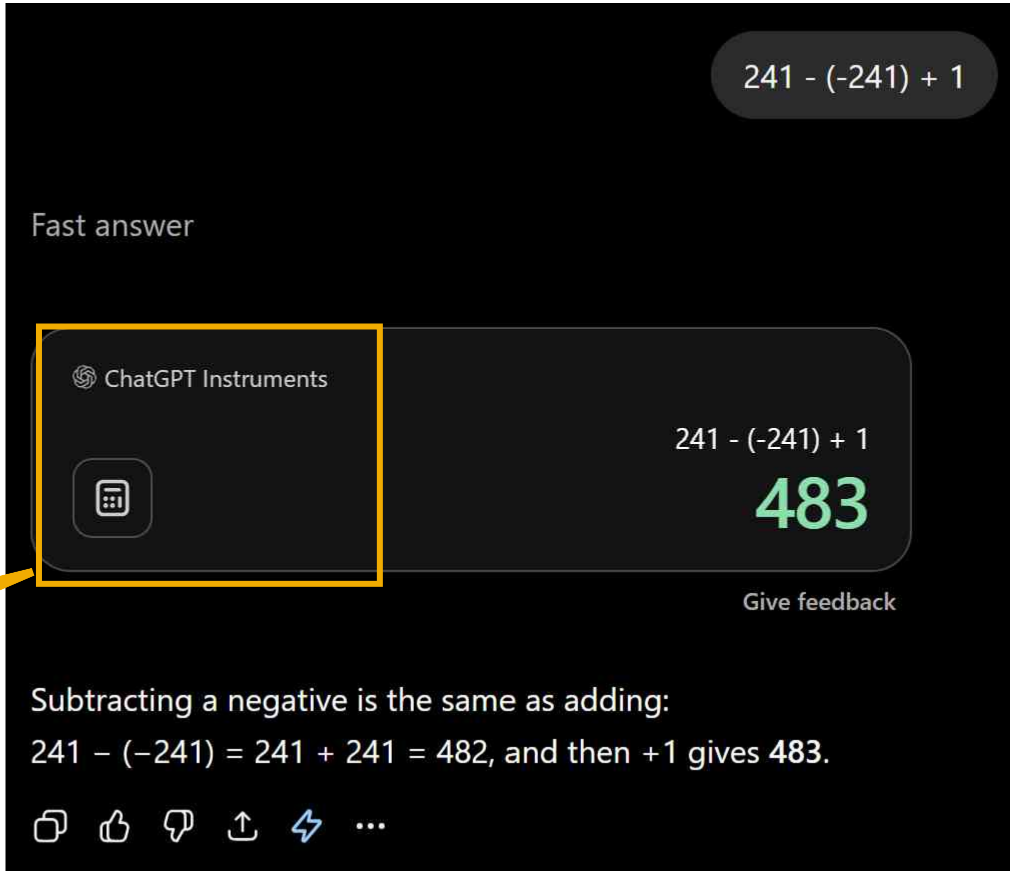

Integrated tool use (!)

Tool learning with foundation models

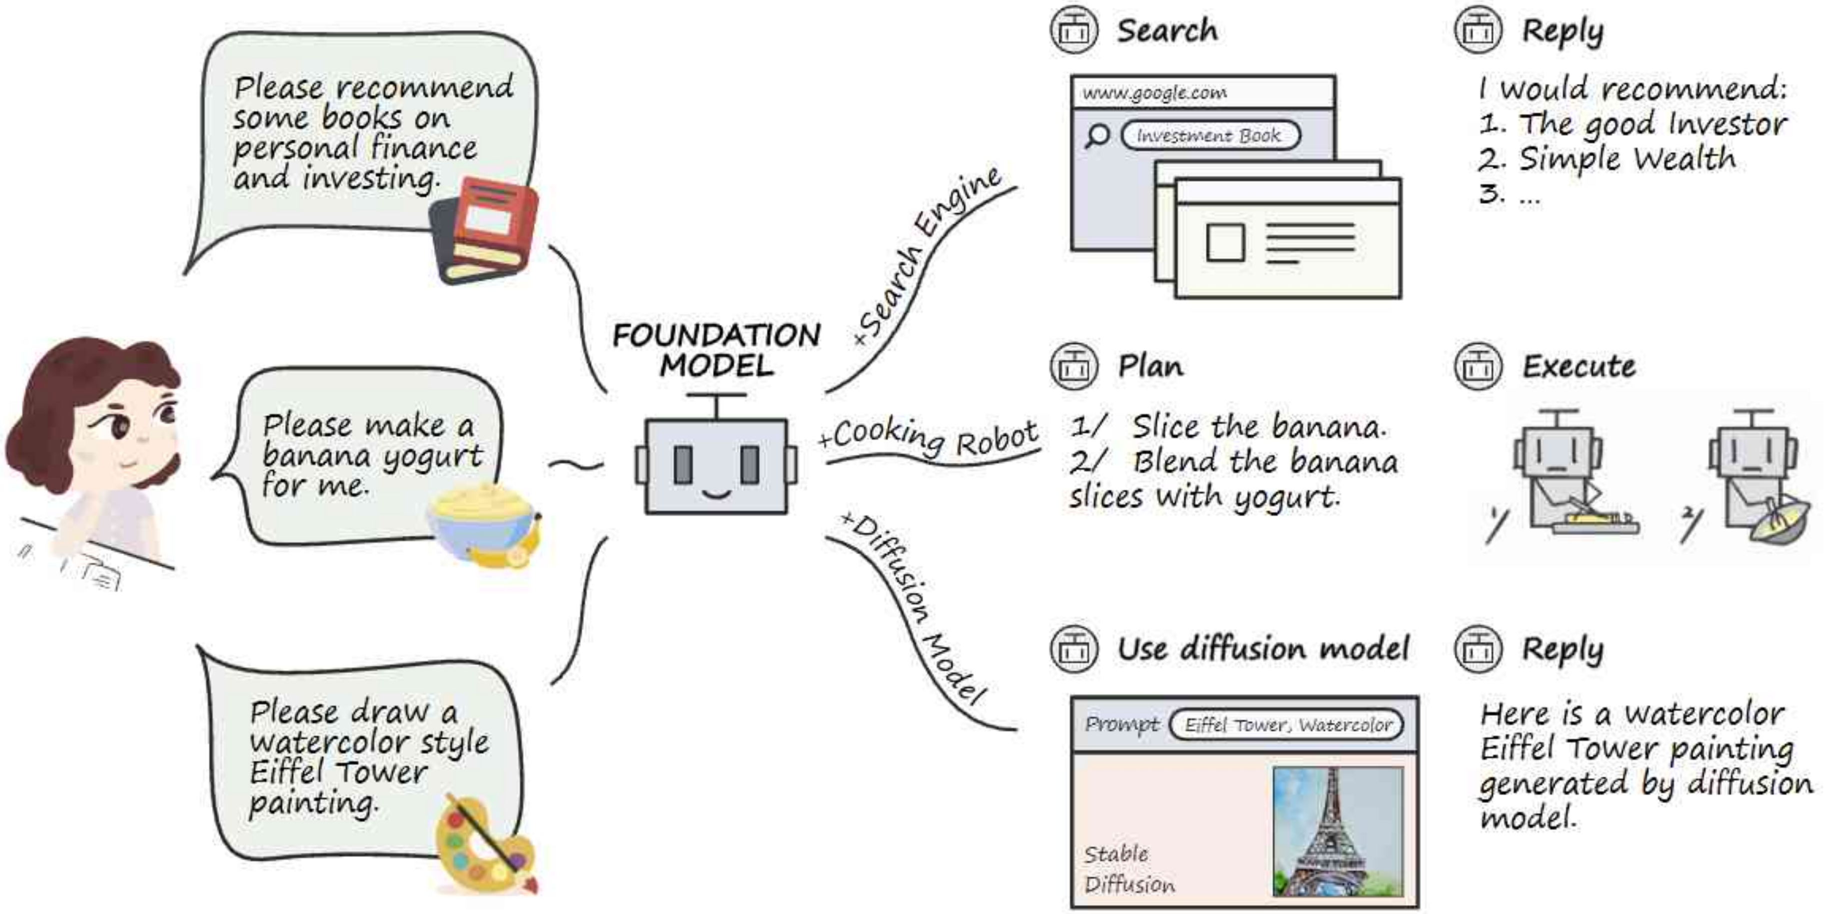

Qin Y, Hu S, Lin Y, et al (2024) Tool Learning with Foundation Models. *ACM Computing Surveys* 57:101:1-101:40. https://doi.org/10.1145/3704435

Tool learning and exemplary tools

Tool learning is the process by which an autonomous AI agent learns to **select**, **parameterize**, and **execute external tools**.

- Tool discovery and selection
- Parameter formulation
- Execution & feedback loop

Exemplary tools:

What makes an AI agent an agent?

More than an LLM with tools

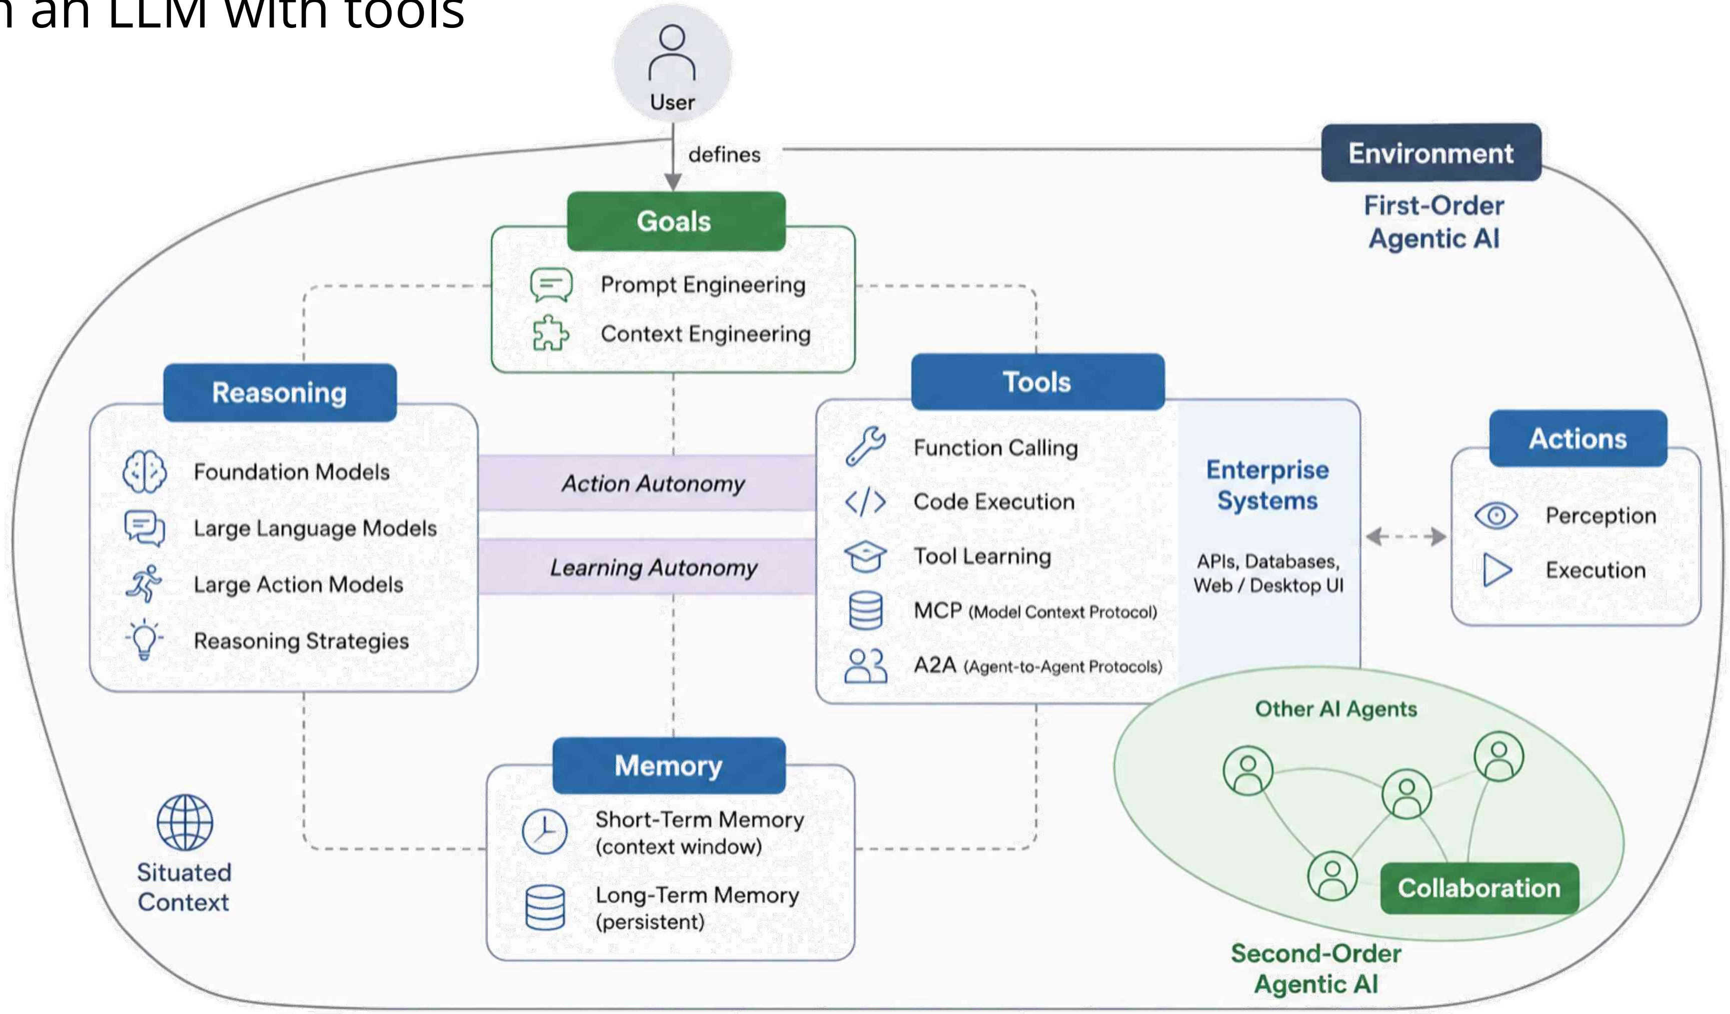

Adapted from Li MM, Reinhard P, Leimeister JM (2026) Agentic AI. *Business & Information Systems Engineering* 1–14. https://doi.org/10.1007/s12599-026-01009-w

From answering once to pursuing a goal

Agents use feedback to repeatedly decide what to do next

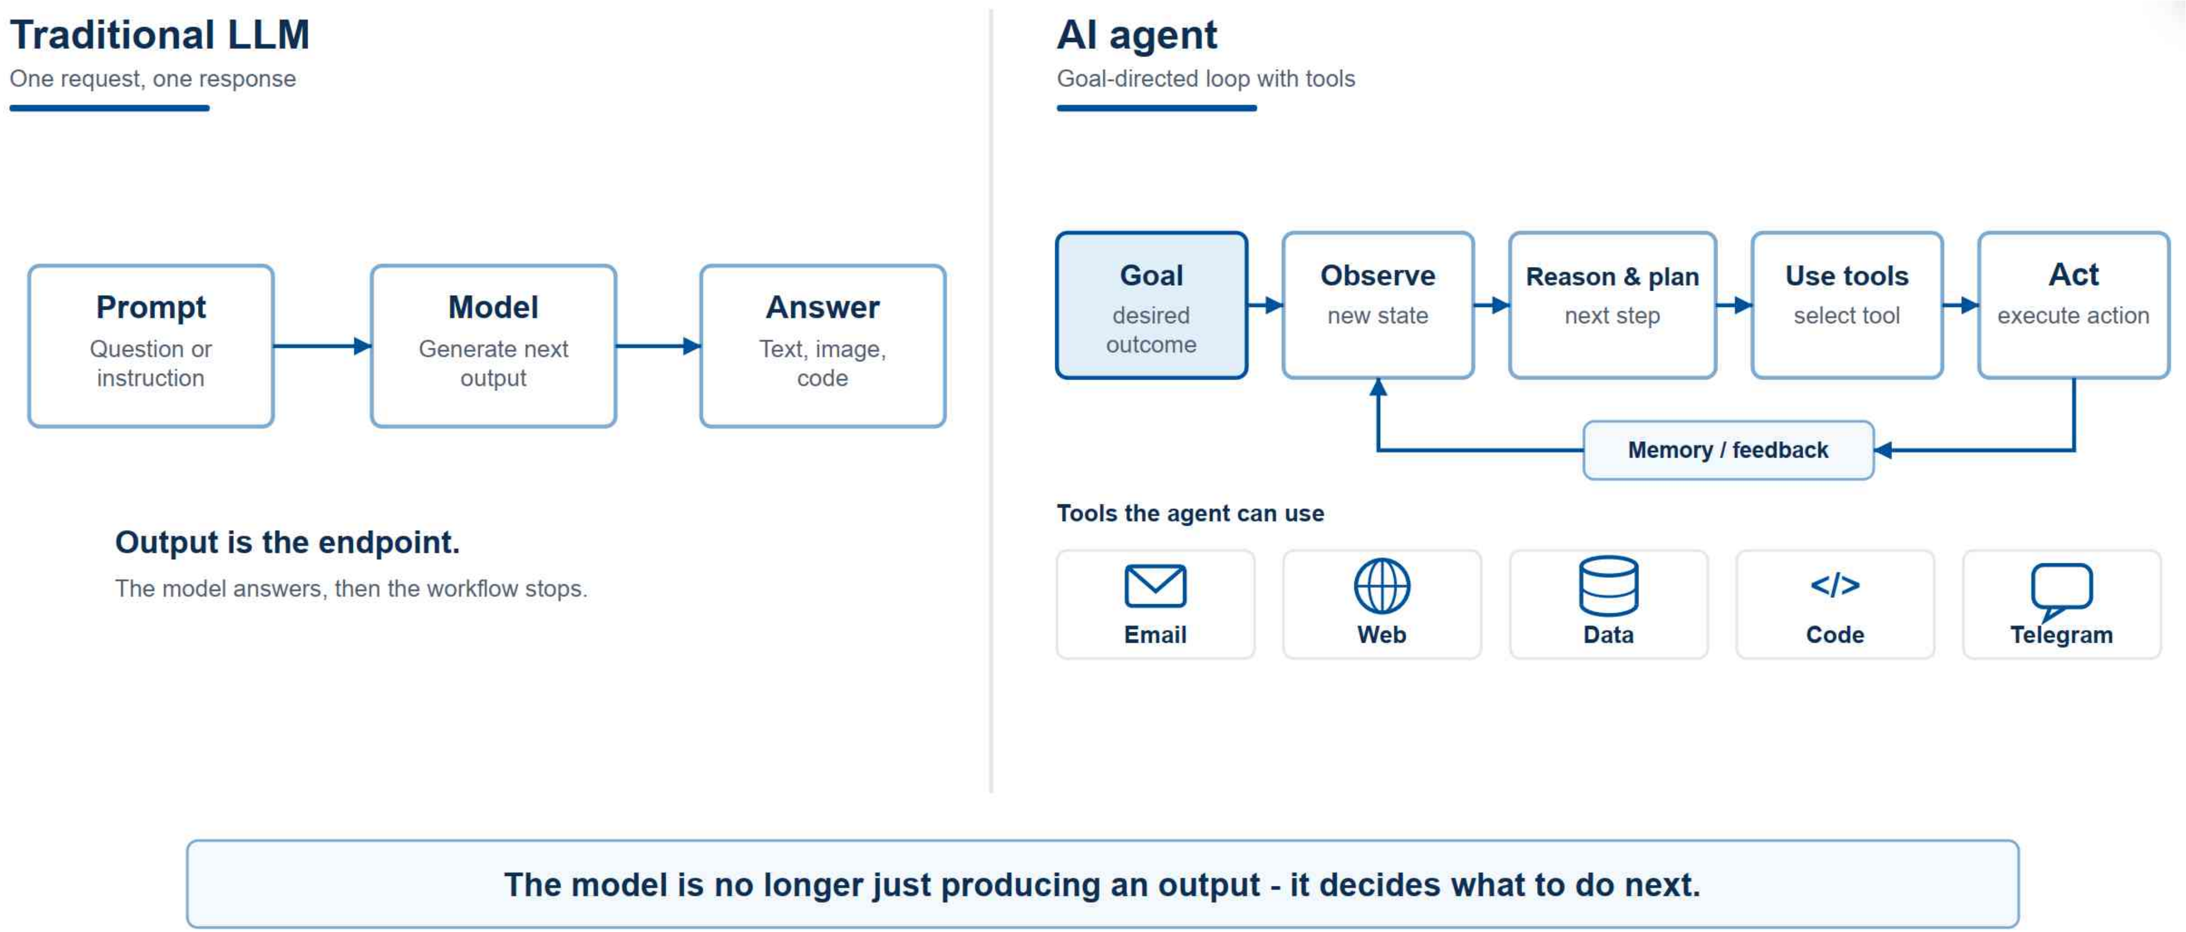

Let’s have a look at some examples

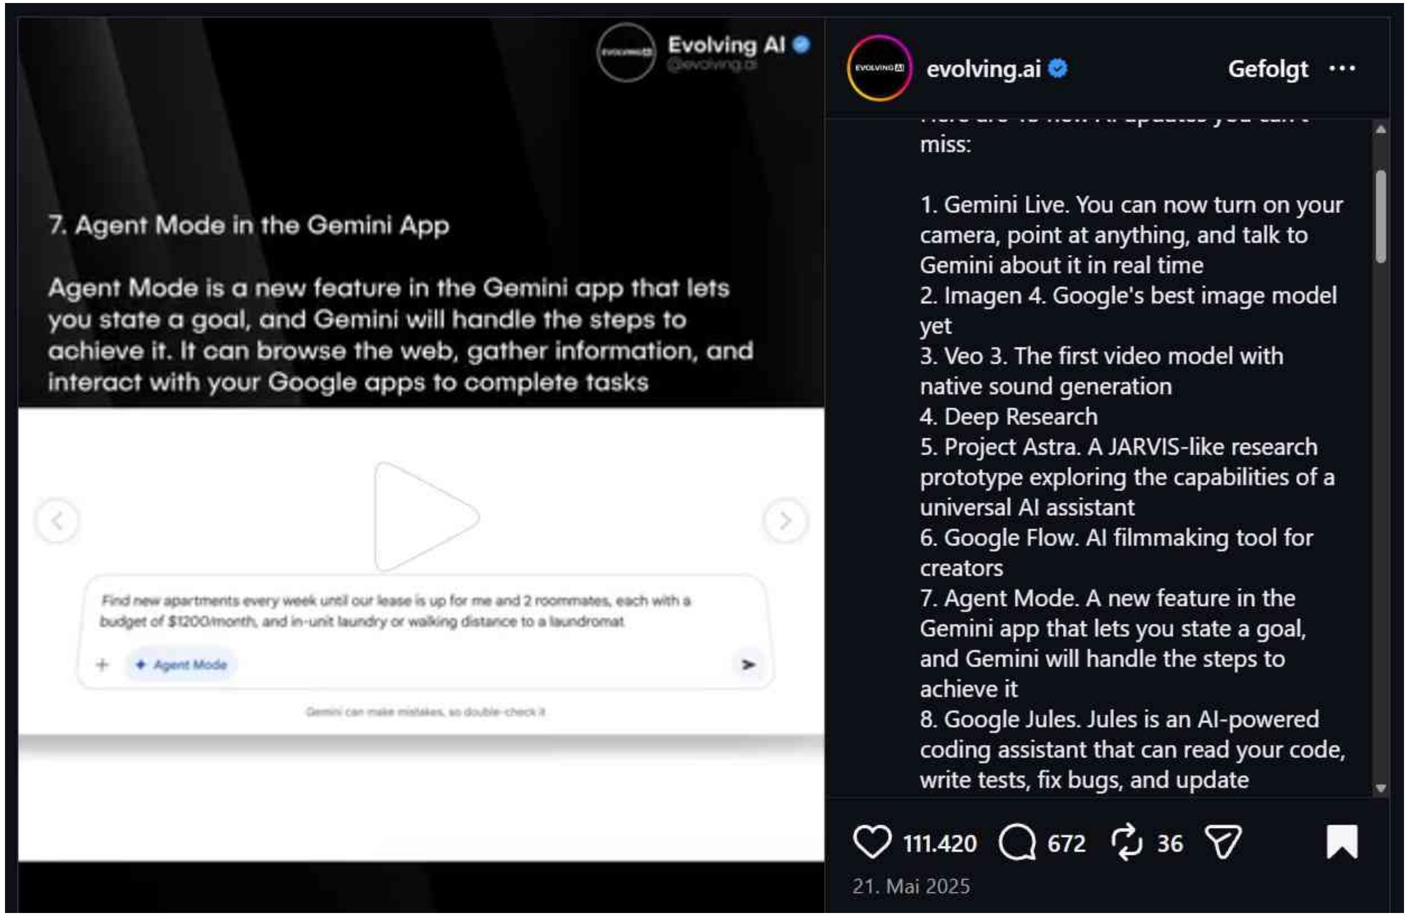

https://www.instagram.com/p/DJ6-swRCyiC?img_index=8

Traditional automation or agentic AI?

Apartment search in high-demand cities

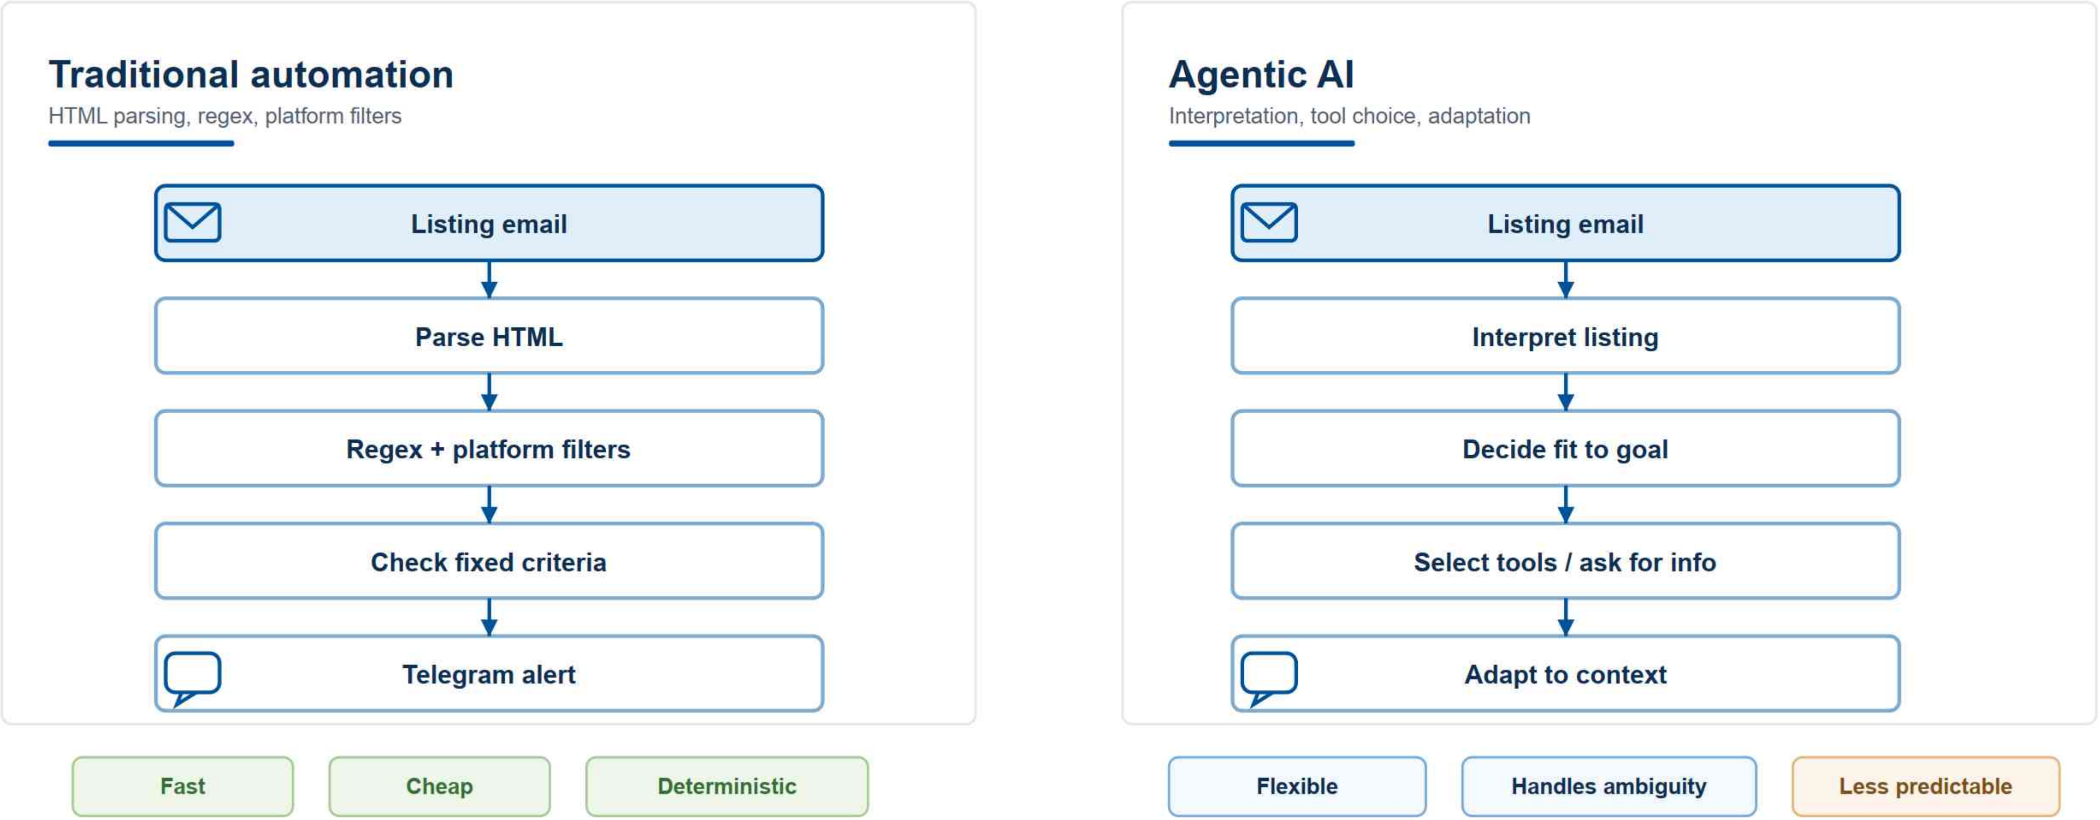

Use agents where contextual reasoning, tool choice, and adaptation add value – not where rules already solve the problem.

The agentic enterprise: AI agents as a team

Define the outcome – let specialized agents coordinate the work

1) Set the goal

“Organize a one-day AI workshop for 50 students within a €2,000 budget.”

2) Assign multiple agents

- Program agent,
- venue agent,
- communications agent,
- catering agent,
- etc. …

3) Approve & monitor

Review plans, budgets, risks, and progress.

Example

Paperclip turns a goal into coordinated agent work

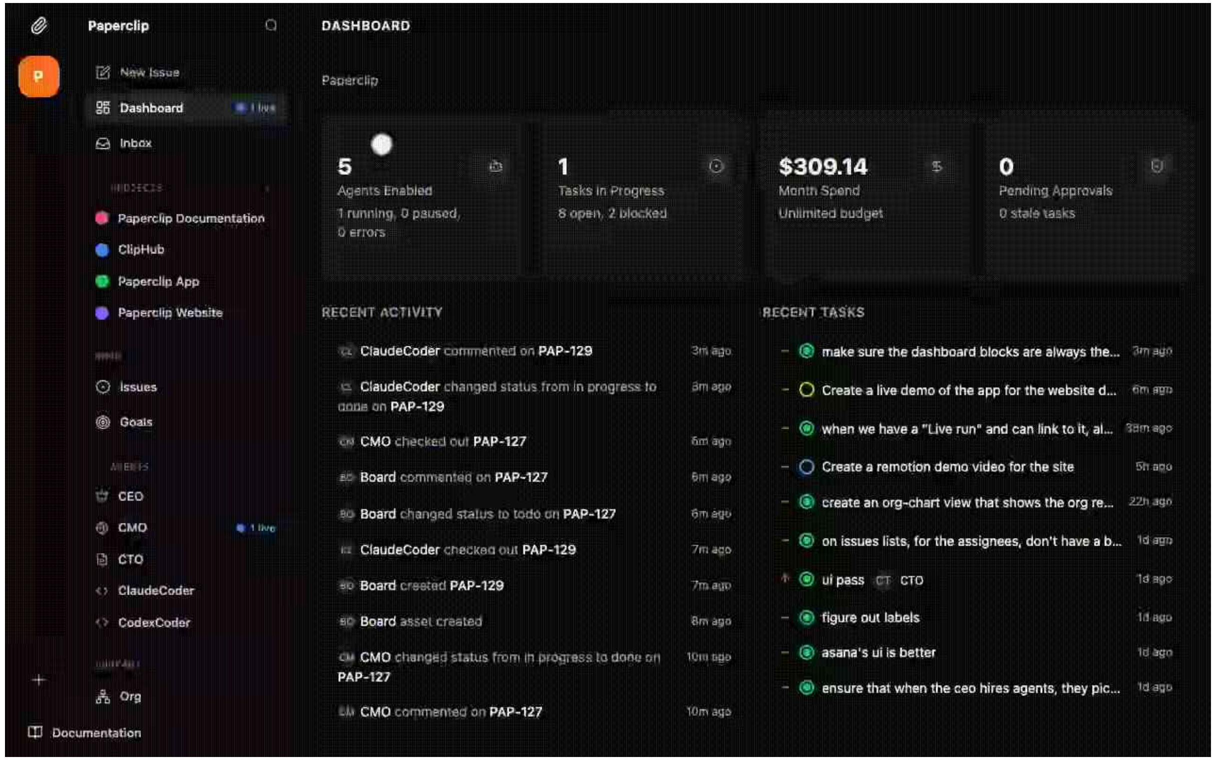

https://paperclip.ing/

The agentic enterprise
SAP Business AI

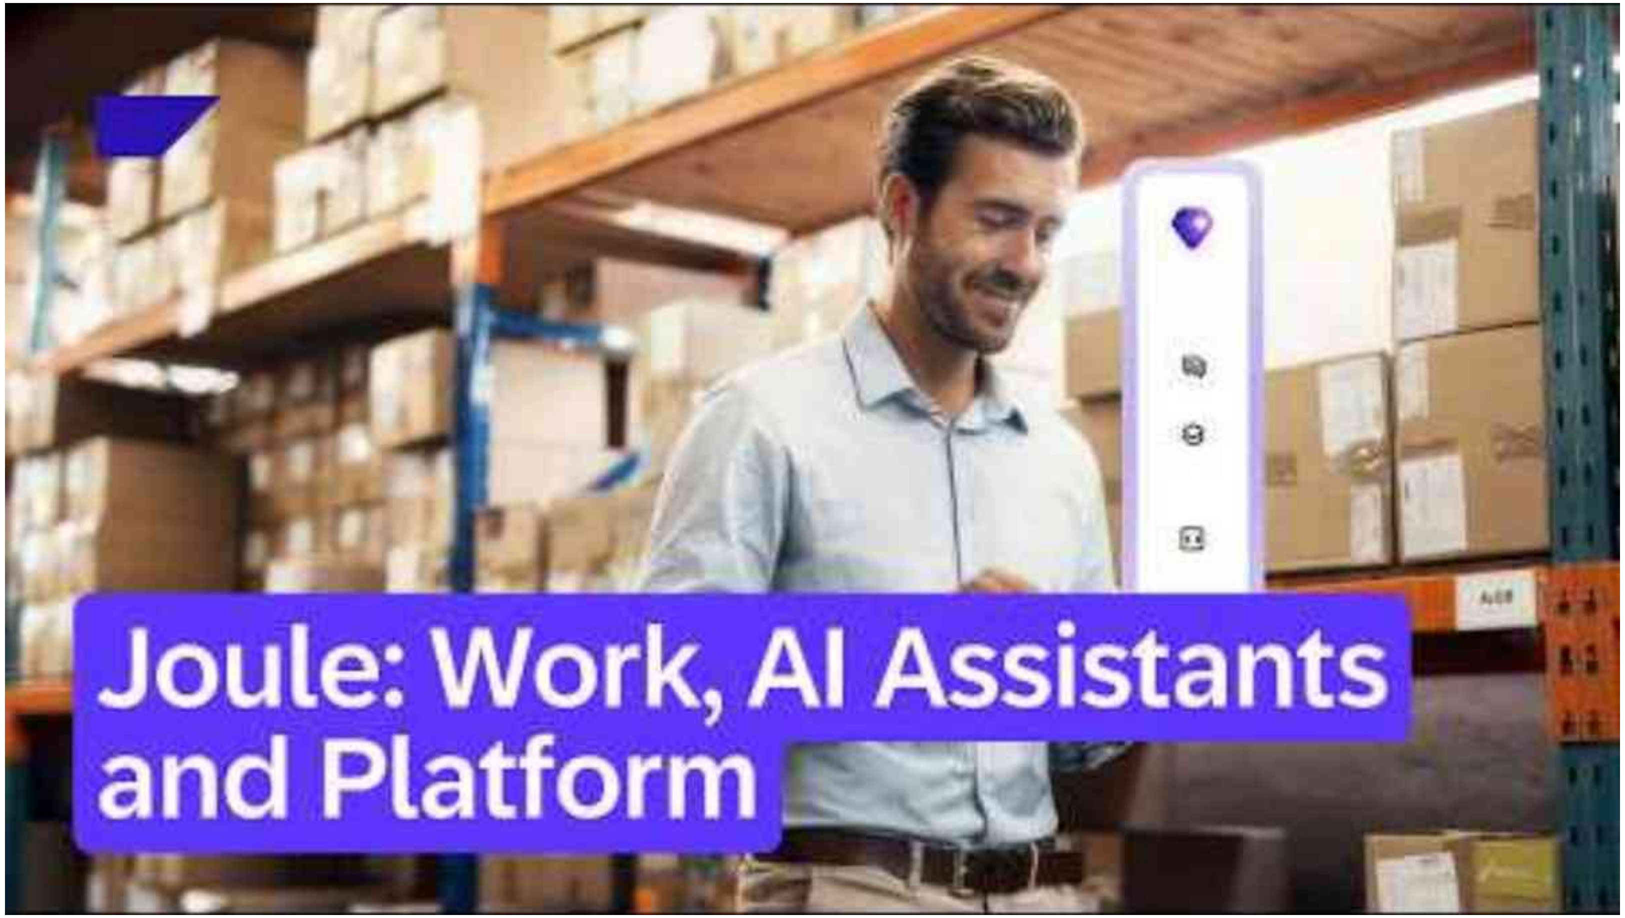

Overreliance on AI

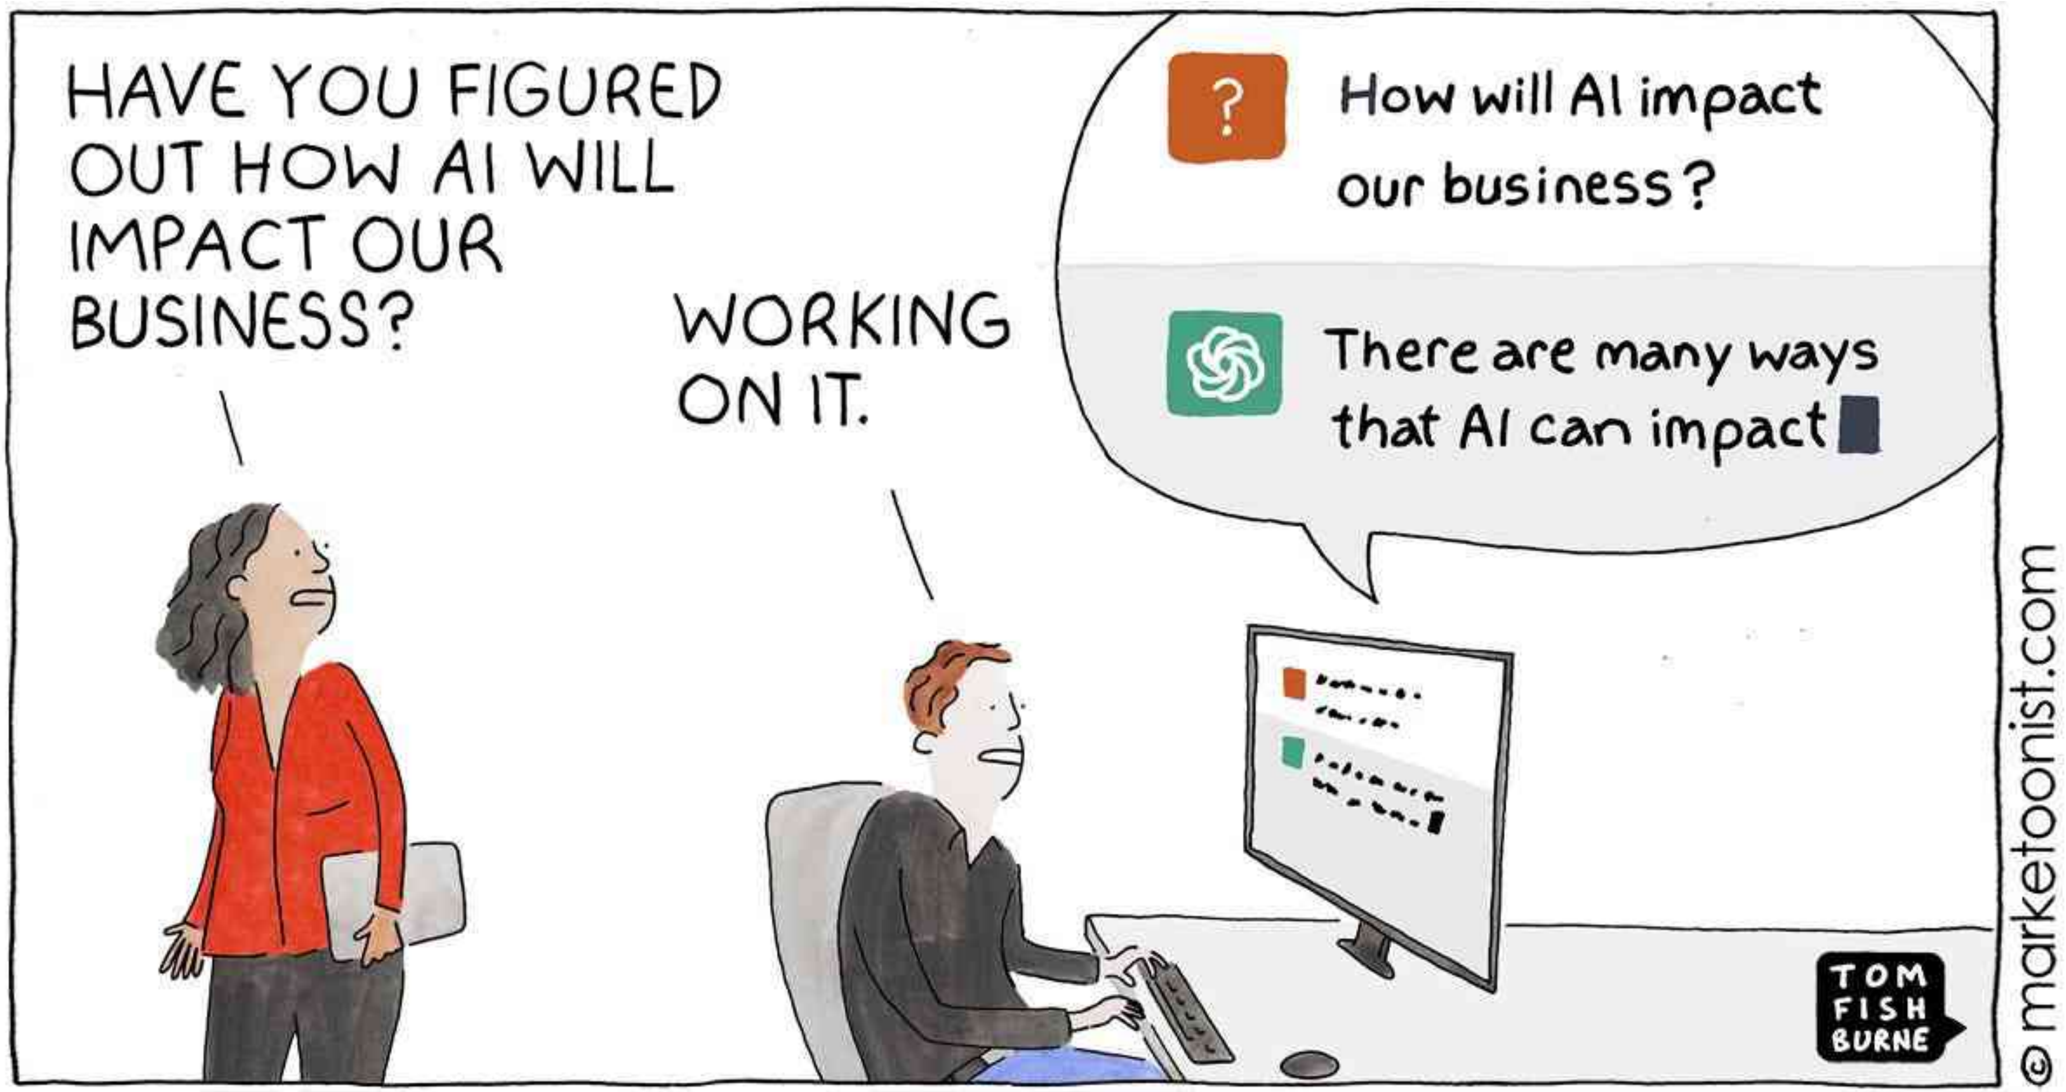

Thank you for your attention! ☺

Any questions?