DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE FOR SCIENCE

Sašo Džeroski
Head of Department of knowledge technologies
Jožef Stefan Institute, Ljubljana, Slovenia

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

WHY AI FOR SCIENCE IS IMPORTANT

Science has huge influence on our everyday lives

Science generates new knowledge and

- *Provides us with evidence-based understanding of the world*
- *It enables modern technologies*
  (incl. communication, transportation, energy)
- *It improves our quality of life*
  (through medicine, environmental protection, and food safety)

Artificial Intelligence provides us with ever more powerful tools

Fueled by

- *Ever faster hardware*
- *Ever more extensive data collection*
- *More and more powerful software*

AI can help science solve the global challenges better and faster

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

AI FOR SCIENCE: HISTORY

The field of AI is changing with lightning speed. *Today, AI (in the public eye) is all about Generative AI and Large Language Models (foundation models), underpinned by transformer architecture.*

But, AI has a long history. *And so has AI for science, starting with expert systems (DENDRAL, late 1960s) and computational scientific discovery (BACON, late 1970s).*

Lessons learned in computational scientific discovery *(early 2000s):*

- Output should be easy to communicate to domain scientists
- We should take advantage of background/ domain knowledge
- Should be able to infer knowledge from small data sets
- Produce models that provide explanations of data
- Support interaction with domain scientists

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE FOR SCIENCE:
How can AI help scientists

AI can be used to automate science

Or at least certain aspects of the scientific process

A prototypical example is finding scientific laws from data, e.g., equations

But other forms of scientific knowledge can also be generated, such as taxonomies

Classification of individual objects within the taxonomies, e.g., living organisms or celestial objects

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE FOR SCIENCE: How can AI help scientists

Experiment design

Experiment execution (via laboratory robotics)

Hypothesis formation

Improving scientific communication

- *Summarizing papers*
- *Literature reviews*
- *Drafting papers/ reports/ proposals*
- *Translation tools improve accessibility to scientific knowledge*

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE FOR SCIENCE
Course outline

A very quick introduction to artificial intelligence

- *What is AI: AI in a nutshell*
- *Knowledge-based AI*
- *Data-based AI*
- *Large language models*

From responsible use of AI to AI as infrastructure

- *Problems with LLMs*
- *Guidelines for responsible use of GenAI*
- *llm.ijs.si: A local infrastructure for science*

SLAIF: The Slovenian AI factory

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE FOR SCIENCE
Course outline

AI methods for use in science

- *Explainable ML for science,*
- *Foundation models for science,*
- *Automated scientific modelling, and*
- *Semantic technologies for open science.*

Example applications of AI in the sciences

- *Life sciences*
- *Environmental sciences*
- *Materials science*

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE: A VERY QUICK INTRO

Sašo Džeroski

Head of Department of knowledge technologies, Jožef Stefan Institute, Ljubljana, Slovenia

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

1. ARTIFICIAL INTELLIGENCE IN A NUTSHELL

2. KNOWLEDGE-BASED AI

3. DATA-BASED AI

4. LARGE LANGUAGE MODELS

WHAT IS ARTIFICIAL INTELLIGENCE

ARTIFICIAL INTELLIGENCE
is the science and engineering
of making intelligent machines,
especially intelligent computer programs

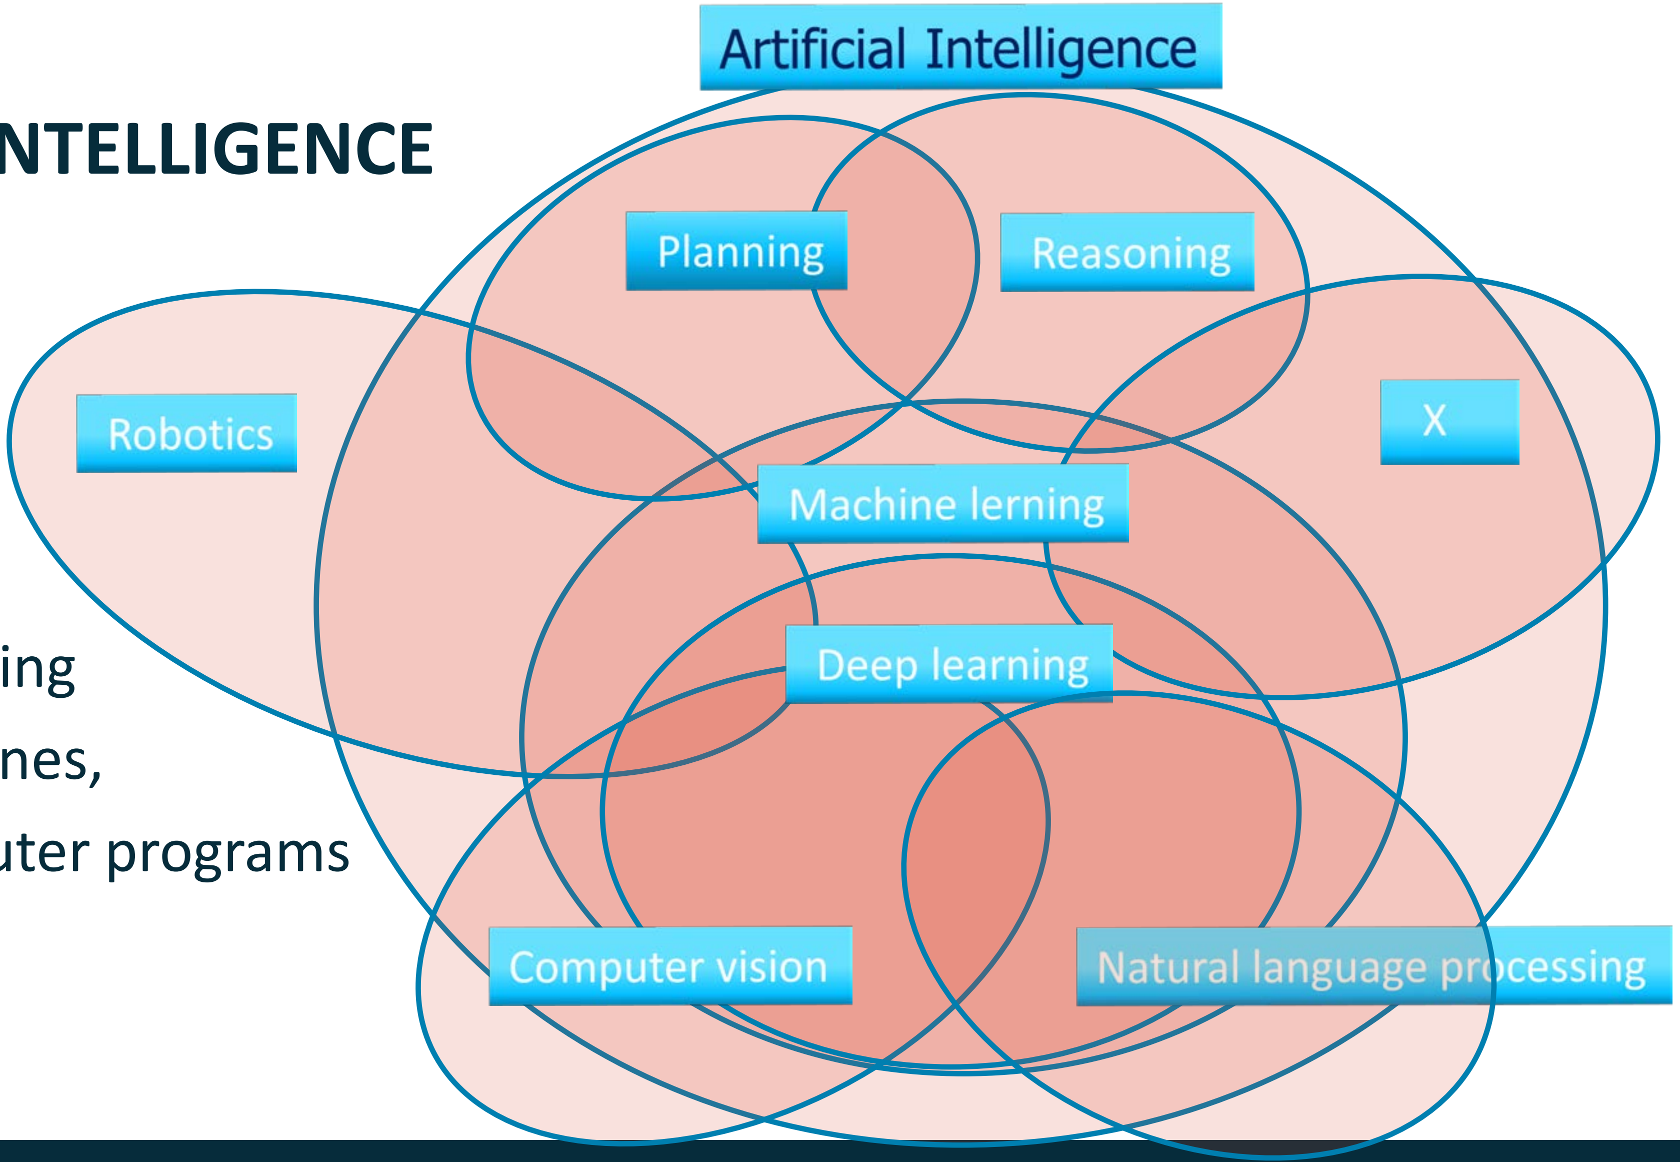

AI HISTORY AND EARLY OPTIMISM

Start in the 1950s, Dartmouth conference in 1956: Big expectations, underestimating the difficulties that lie ahead

- 1958, H. A. Simon and Allen Newell: "… within ten years a digital computer will discover and prove an important new mathematical theorem."
- 1965, H. A. Simon: "… machines will be capable, within twenty years, of doing any work a man can do."
- 1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."
- 1970, Marvin Minsky: "In from three to eight years we will have a machine with the general intelligence of an average human being."

AI HISTORY: EVOLUTION OF TOPICS

- Neural networks (machine learning)
  - Built from artificial neurons (perceptrons)
- Expert systems (knowledge representation)
  - IF-THEN rules
  - Machine learning of symbolic rules
- Computer vision
- Natural language processing
- Machine learning of deep neural networks
  - Convolutional neural networks (CNNs) for images
  - Neural networks (transformers) for generating text: LLMs=Large language models)
  - Multi-modal neural networks

    (can process, e.g., text and images)

THE PERCEPTRON: AN ARTIFICIAL NEURON

- The perceptron can be viewed as an artificial neuron

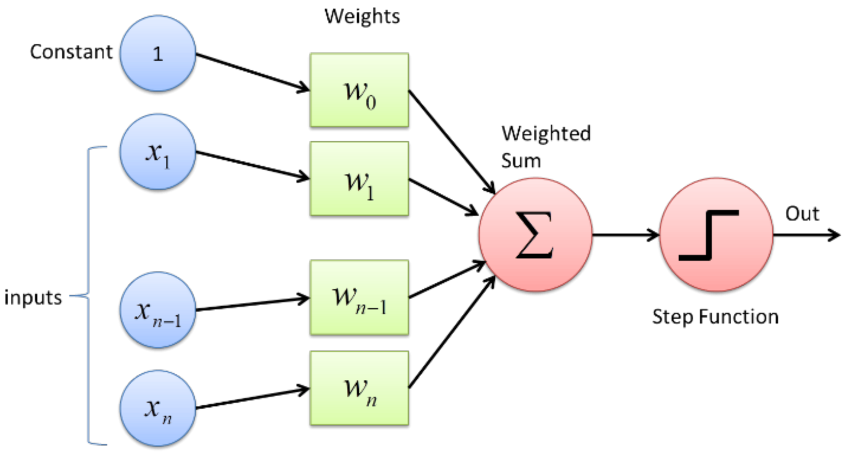

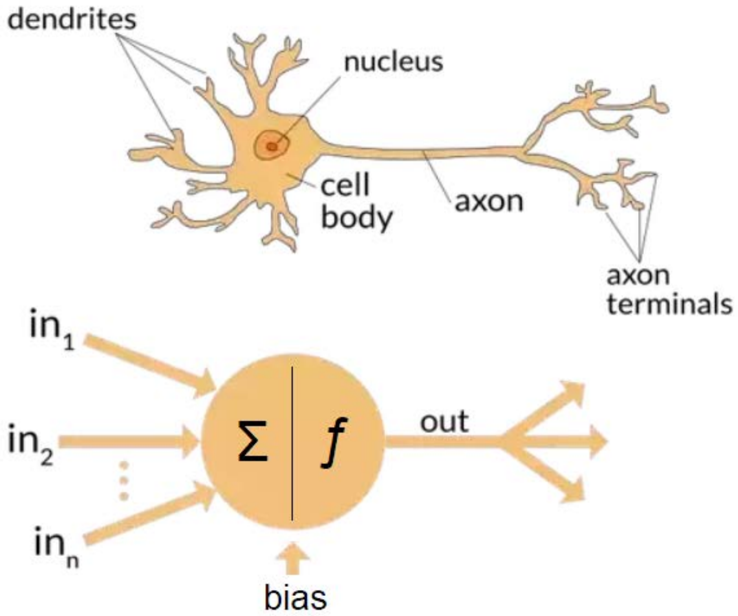

- The basic unit of Artificial Neural Networks (Rosenblatt, 1958)
- Performs binary classification via linear discrimination

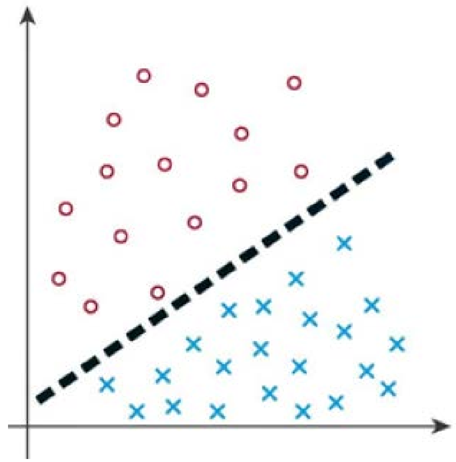

THE MULTI-LAYER PERCEPTRON AND NNs

While a single perceptron can learn a linear mapping,

MLPs can approximate an arbitrary non-linear mappings between inputs & outputs

Arbitrary depth or width may be needed

Note the hidden layers: If only a few, shallow NNs

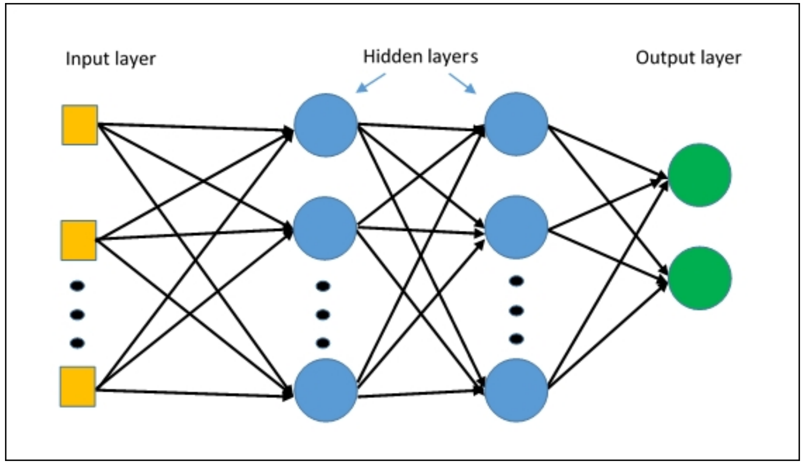

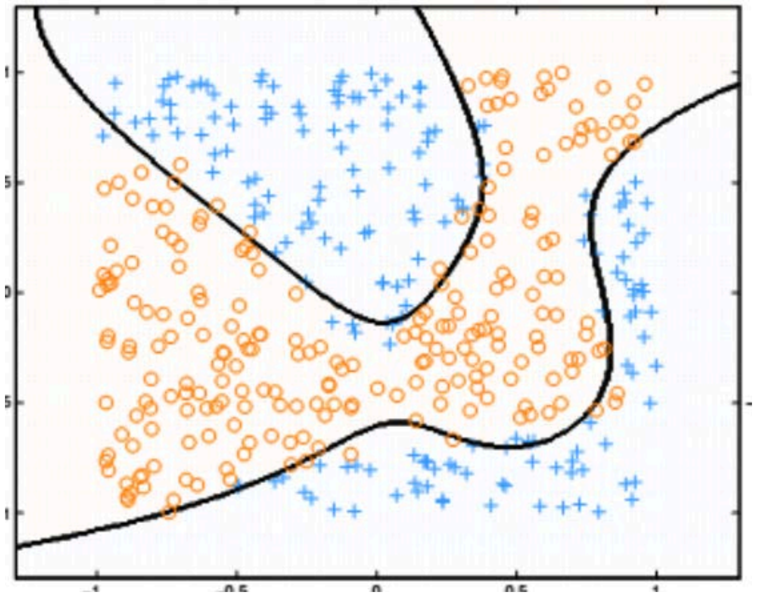

HOW TO LEARN A NEURAL NETWORK: FEED-FORWARD NNs and BACK-PROPAGATION

Making predictions with MLPs: Feed-forward reasoning

Learning the weights in a MLP: Back-propagation of error

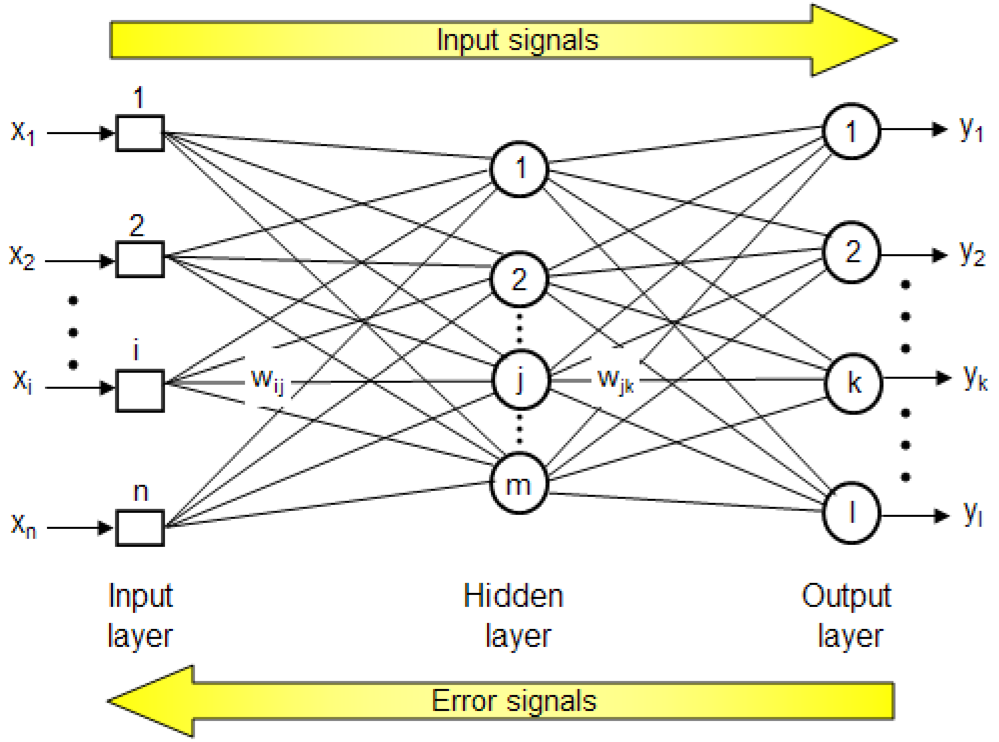

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

1. ARTIFICIAL INTELLIGENCE IN A NUTSHELL

2. KNOWLEDGE-BASED AI

3. DATA-BASED AI

4. LARGE LANGUAGE MODELS

EXPERT SYSTEMS

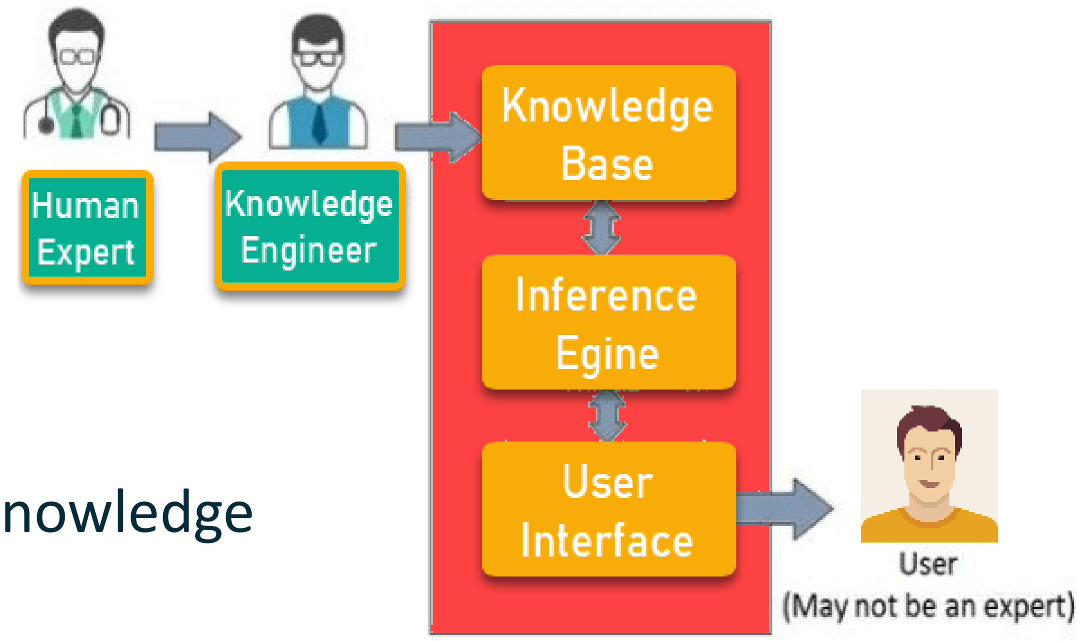

EXPERT SYSTEMS are computer systems emulating the decision-making ability of a human expert

- Designed to solve complex problems by using knowledge
- Knowledge stored in knowledge bases, represented mainly as if–then rules
- Problems solved by reasoning over the knowledge base by inference engine
- Decisions can be explained

EXPERT SYSTEMS IN MEDICINE

MYCIN was an early expert system developed at Stanford in the 1970s (work on it started in 1972)

Its area of expertise was treating blood infections

An example rule from the MYCIN knowledge-base

IF the stain of the organism is gram-positive
AND the morphology of the organism is coccus
AND the growth conformation of the organism is clumps
THEN (0.7) the identity of the organism is staphyloccus

RULE037

IF: 1) *The identity of the organism is not known with* certainty, and
2) The stain of the organism is gramneg, and
3) The morphology of the organism is rod, and
4) The aerobicity of the organism is aerobic

THEN: There is strongly suggestive evidence (.8) that the class of the organism is enterobacteriaceae

RULE145

IF: 1) The therapy under consideration is one of: cephalothin clindamycin erythromycin lincomycin vancomycin, and
2) Meningitis is an infectious disease diagnosis *for the patient*

THEN: It is definite (1) the therapy under consideration is not a potential therapy for use against the organism

RULE060

IF: *The identity of the organism is bacteroides*

THEN: I recommend therapy chosen from among the following drugs:

1 - clindamycin	(.99)
2 - chloramphenicol	(.99)
3 - erythromycin	(.57)
4 - tetracycline	(.28)
5 - carbenicillin	(.27)

EXPERT SYSTEMS ARE STILL ALIVE
Qualitative, hierarchical multi-attribute decision models

B1) Decision modelling

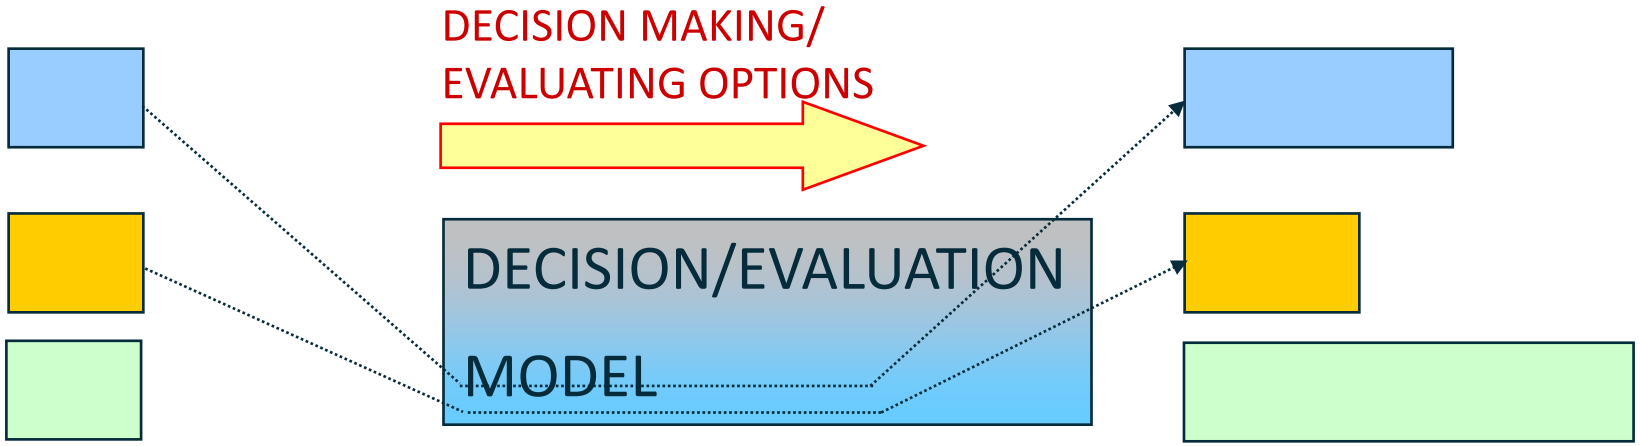

Hierarchical multi-criteria decision modelling method

- employs qualitative variables and decision rules
- is knowledge-based (rather than data-based), supports explanation

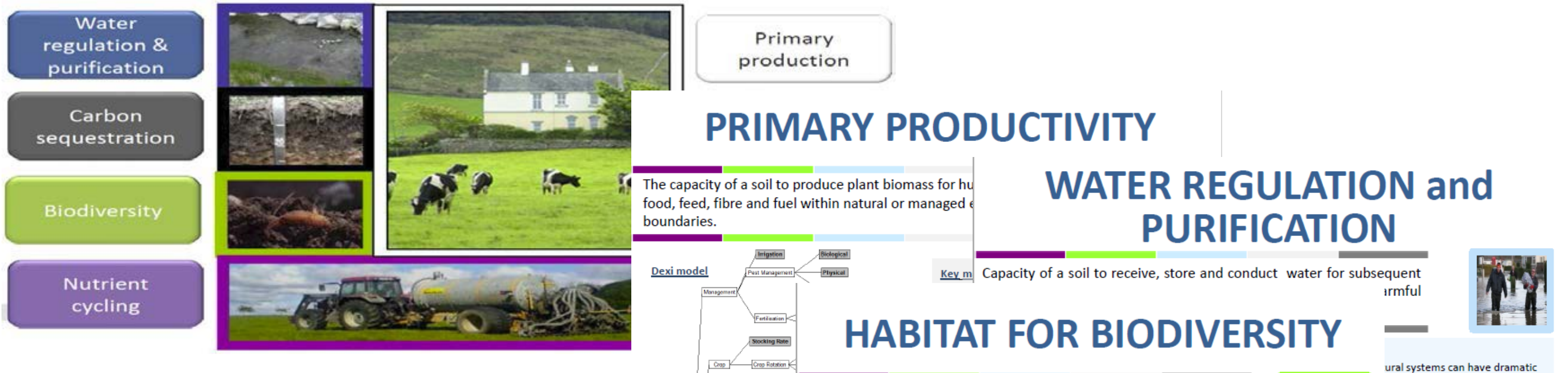

EXPERT SYSTEMS ARE STILL ALIVE

Qualitative, hierarchical multi-attribute models as knowledge bases

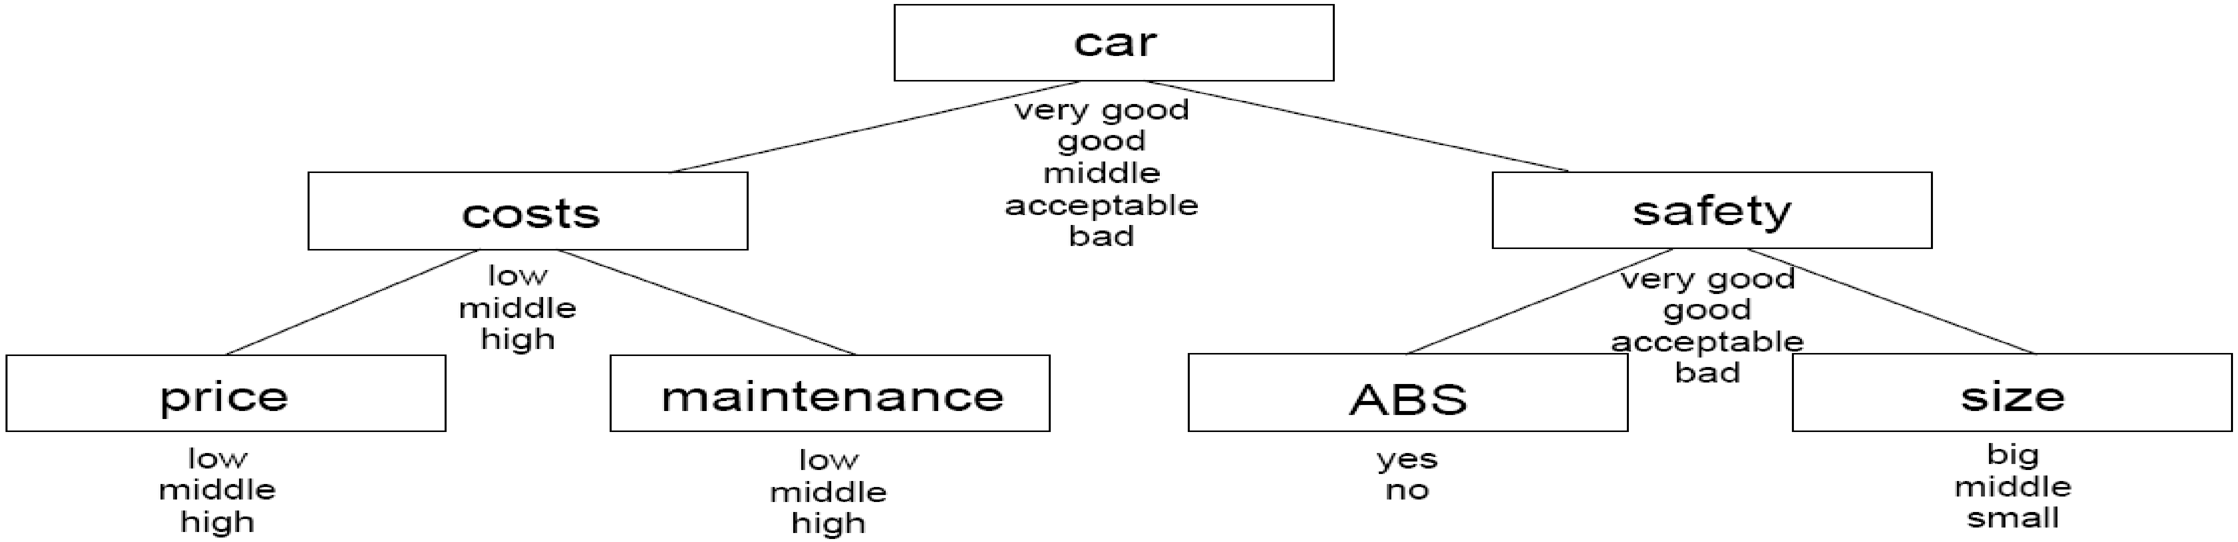

price	maint.	costs
low	low	low
low	mid.	low
low	high	mid.
mid.	low	low
mid.	mid.	mid.
mid.	high	high
high	low	mid.
high	mid.	high
high	high	high

costs	safety	car
low	v. good	v. good
low	good	good
low	acc.	mid.
low	bad	bad
mid.	v. good	good
mid.	good	good
mid.	acc.	acc.
mid.	bad	bad
high	v. good	good
high	good	mid.
high	acc.	bad
high	bad	bad

IF ABS=no AND size=small
THEN safety = bad

ABS	size	safety
no	small	bad
no	mid.	acc.
no	big	good
yes	small	bad
yes	mid.	good
yes	big	v. good

THE KNOWLEDGE ACQUISITION BOTTLENECK: STRONG MOTIVATION FOR MACHINE LEARNING

To create an expert system, you need a knowledge-base.

The knowledge needs to go from the experts to the computer.

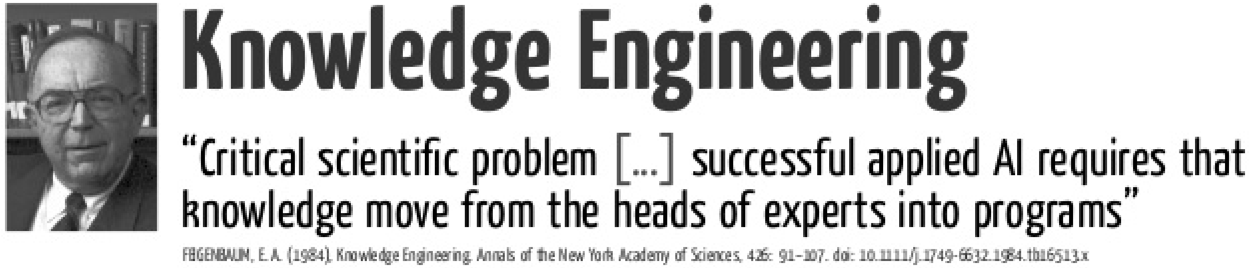

Experts notoriously bad at explaining how they solve problems, but can provide many examples of solved problems

Solution: Learn rules from examples solved by the domain experts (i.e., from data)

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

COMPUTER VISION

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

How computers can be made to gain high-level understanding from digital images or videos

- Typical tasks in computer vision: Recognition, Motion analysis, Scene reconstruction, Image restoration
- Recognition tasks: Object recognition/classification, Identification, Detection, Content-based image retrieval

Even in traditional computer vision, machine learning plays an important role (e.g., learning to classify images)

- Feature construction: knowledge-driven, by human expert (manual, laborious, time consuming)
- Hand-crafted features fed into classifier
- Typically a shallow neural network

NATURAL LANGUAGE PROCESSING

Three levels of analysis

- **Syntax** — the structure of language: how words combine into grammatically valid sentences (tokenisation, POS tagging, parsing)
- **Semantics** — the meaning of language: what words and sentences mean (word sense disambiguation, relationship extraction, semantic role labelling)
- **Pragmatics** — the use of language in context: what the speaker intends, how context shapes interpretation

Downstream tasks: machine translation, sentiment analysis, information extraction, chatbots

Classical/ knowledge-driven NLP: hand-written grammars, rules

Statistical NLP: learned from corpora — currently predominant, but started quite early

1. ARTIFICIAL INTELLIGENCE IN A NUTSHELL

2. KNOWLEDGE-BASED AI

3. DATA-BASED AI: MACHINE LEARNING

4. LARGE LANGUAGE MODELS

MACHINE LEARNING TO THE RESCUE

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

Addressing bottlenecks in three major knowledge-based approaches to artificial intelligence

- Expert systems
- Computer vision
- Natural language processing

Recently shift of focus from knowledge and reasoning towards perception and action

From tasks that can be only done by experts to tasks addressed by everyday people

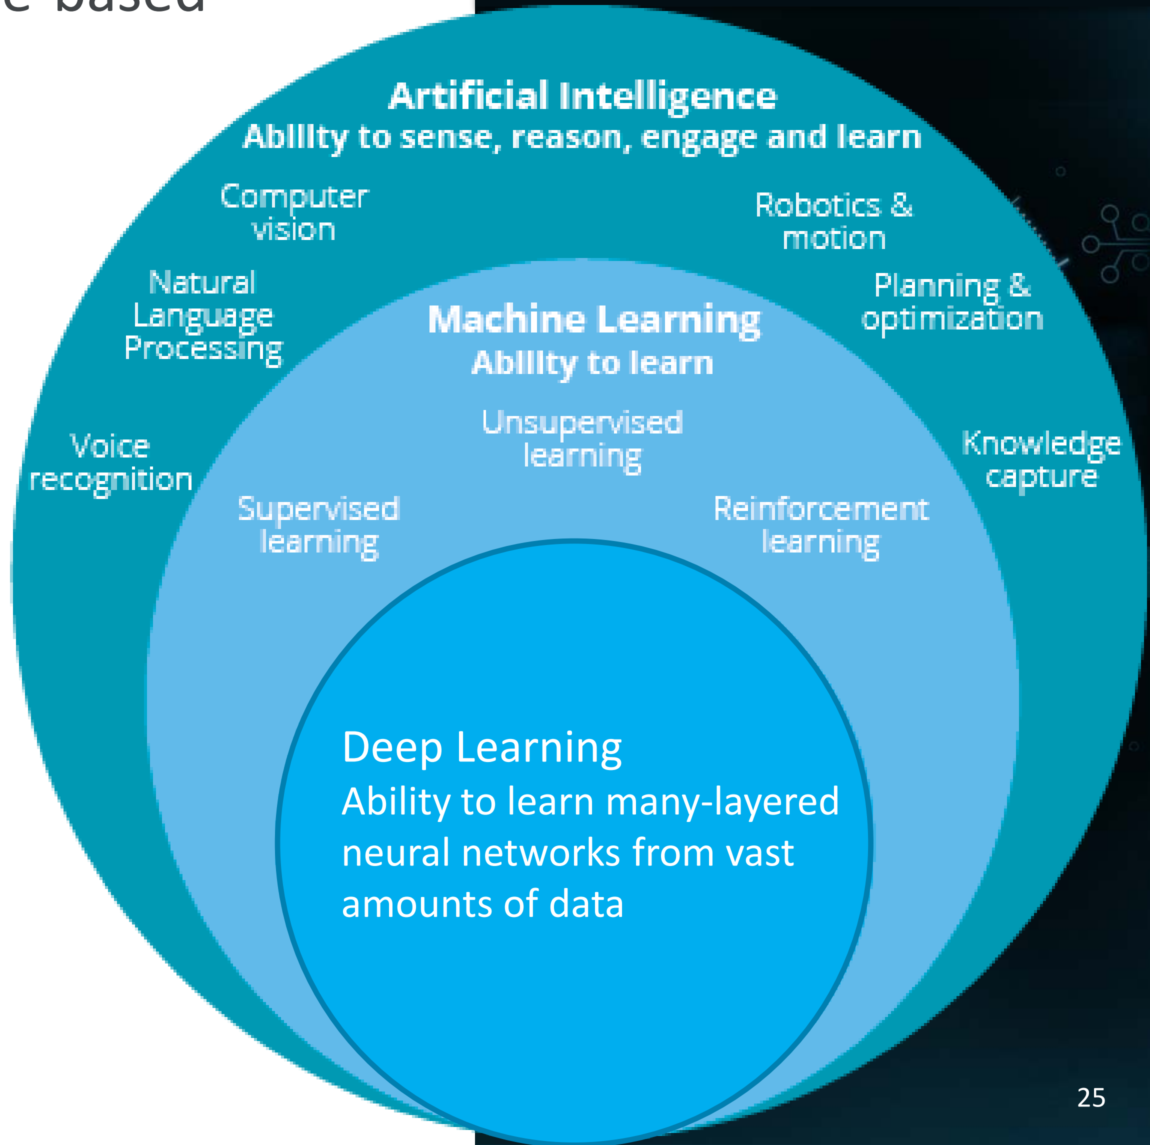

MACHINE LEARNING OF EXPLAINABLE MODELS

Input: Table of data on brown bears

Location	PLC	FOREST-ABUNDANCE	PTS	OtherEnvVariables	BBH
l1	Forest	80	21.4	...	Yes
l2	Forest	66	13.9	...	Yes
l3	Forest	55	50.0	...	No
l4	Forest	72	1.2	...	No
l5	Grassland	6	19.1	...	No
l6	Grassland	0	11.4	...	No
l7	Wetland	3	5.8	...	No
l8	Water	0	3.9	...	No

Output: Models in the form of IF-THEN rules

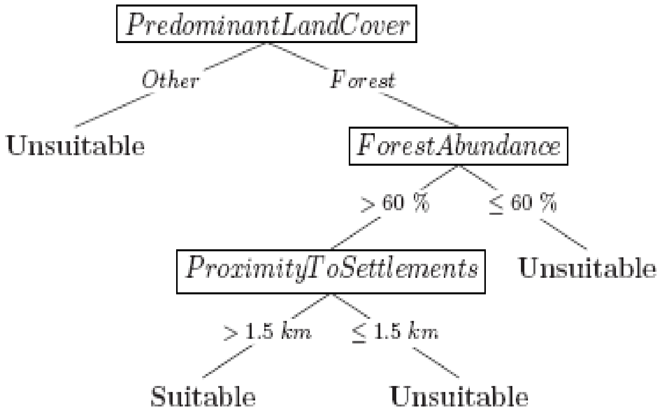


```
IF   PREDOMINANT-LAND-COVER = Forest
AND  FOREST-ABUNDANCE > 60%
AND  PROXIMITY-TO-SETTLEMENTS > 1.5 km
THEN BrownBearHabitat = Suitable
```

or decision trees, predicting habitat suitability

HABITAT MODELS FOR PESTS AND BENEFICIAL SPECIES

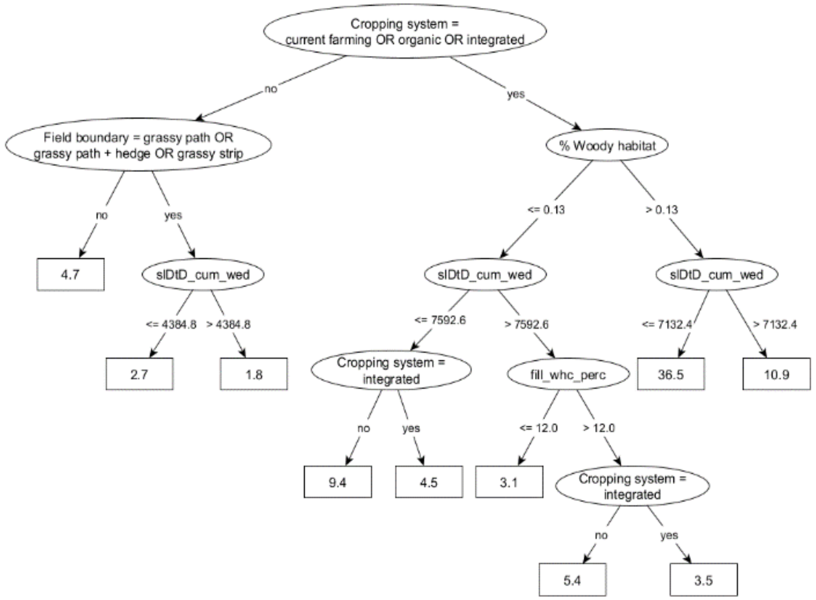

ACCURACY VS EXPLAINABILITY

Two complementary goals:

- Accuracy
- Explainability

Trees and rules are interpretable by design, but are not very accurate

In some domains, such as medicine (or science), explanations are of key importance

So we either

- Learn explainable models (like trees or rules)
- Or learn more complicated models (ensembles, NNs) and explain post-facto the models (var. importance) or their predictions (SHAP, LIME)

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

TREE ENSEMBLE MODELS

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

Ensembles are collections of models (e.g., forests of trees)

Whose predictions are combined to obtain a final prediction

Can have much higher predictive performance vs. single models

A single decision tree

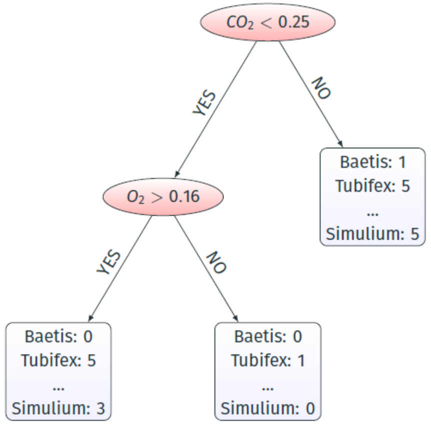

An ensemble of decision trees

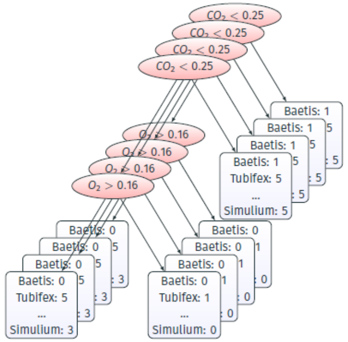

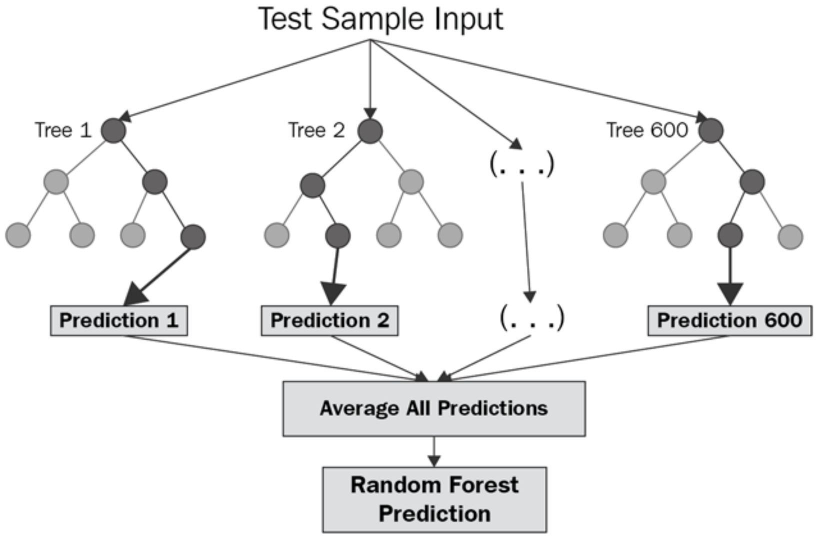

Can provide feature importance estimates
and thus some insight/ interpretation

E.g., random forests and associated variable importance estimates

Very efficient and accurate

DECISION BOUNDARIES IN MACHINE LEARNING

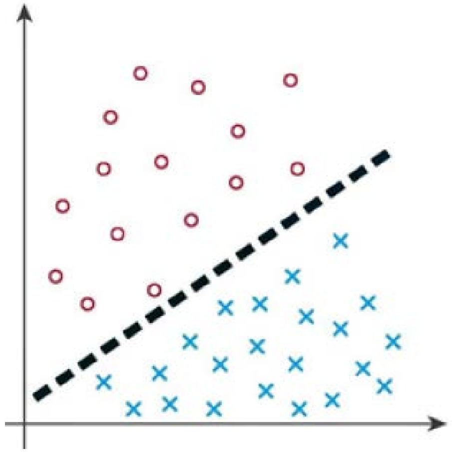

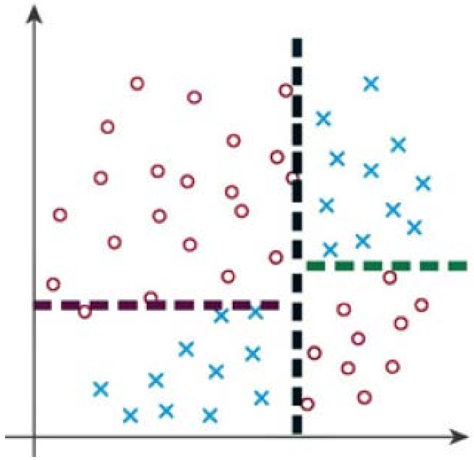

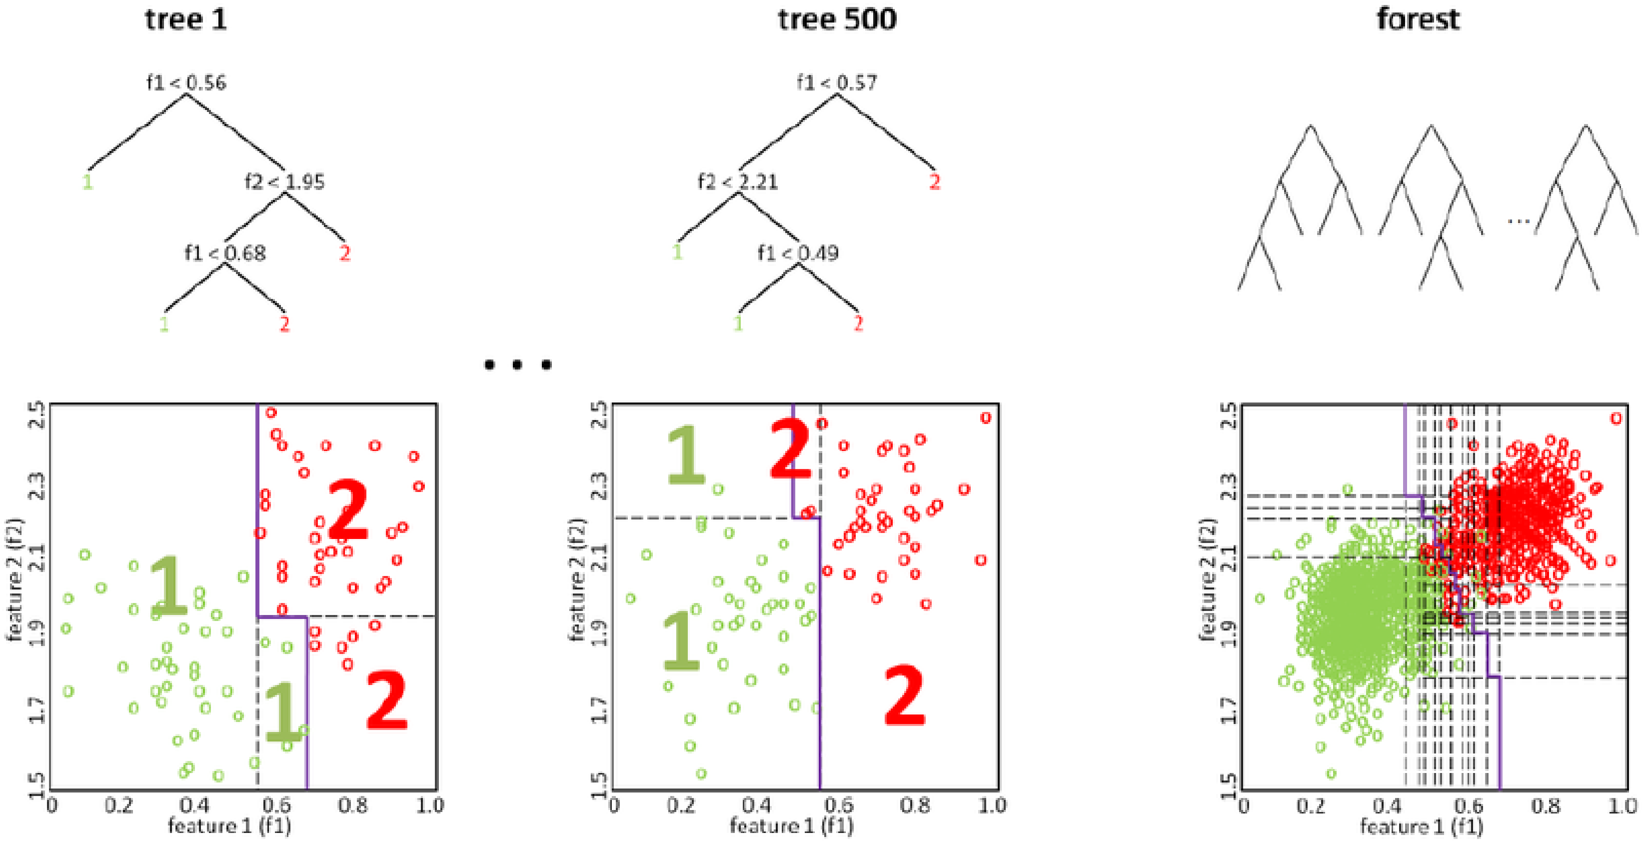

SHALLOW AND DEEP NEURAL NETWORKS

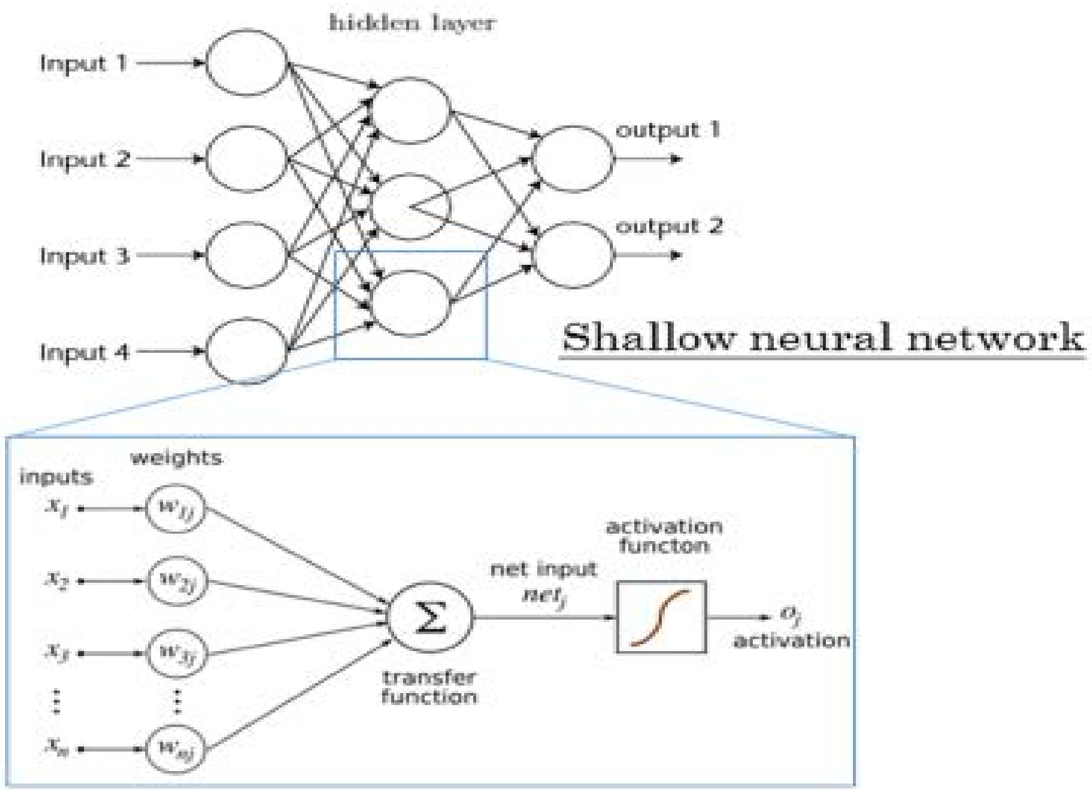

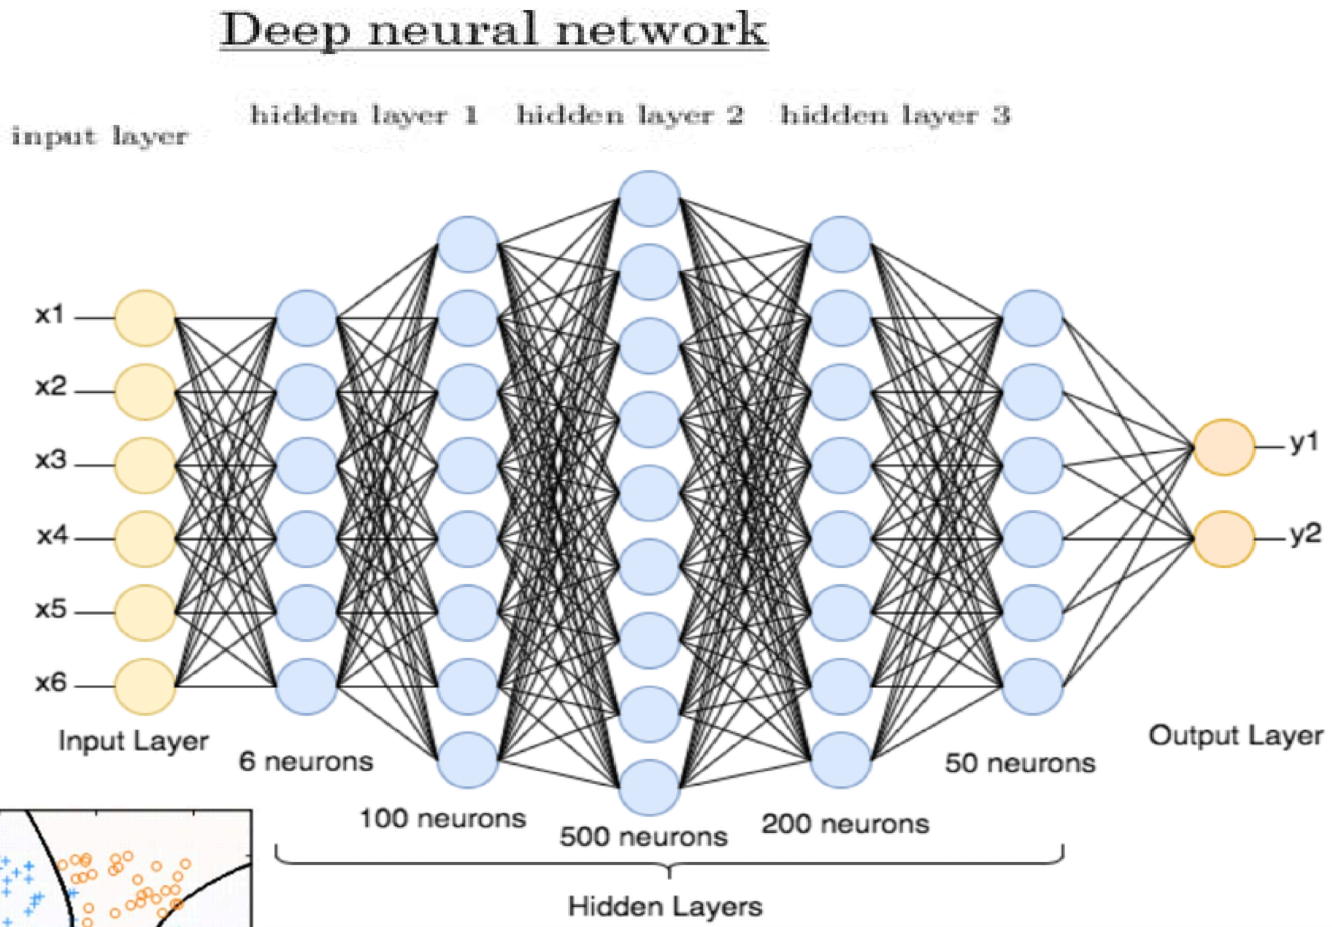

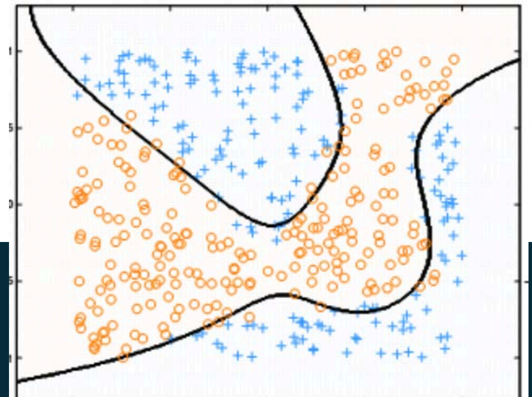

CONVOLUTIONAL DEEP NEURAL NETWORKS

CNNs are DNNs that include computer vision ideas (convolutional filters/ smoothing operators) and can learn features from images

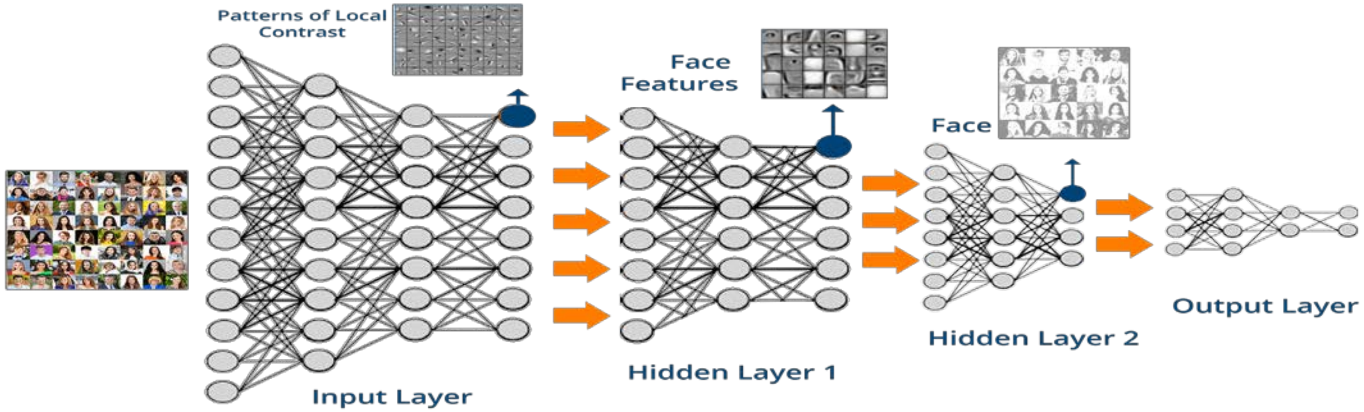

This is the key to the success of NNs: End-to-end learning

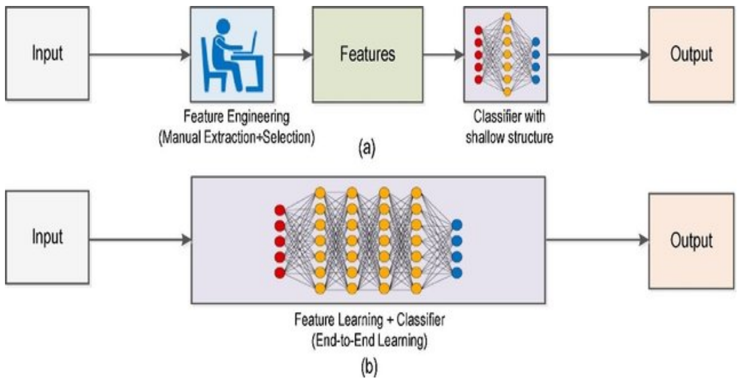

Example: Convolution with a filter

$$K = \begin{matrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{matrix}$$

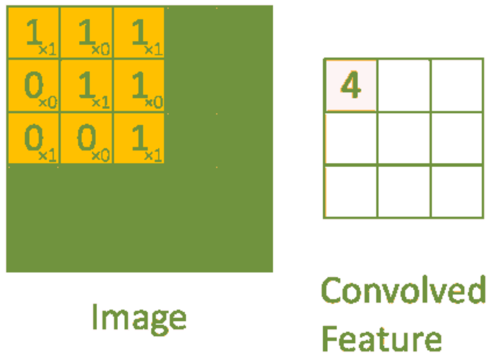

LIMITATIONS OF DNNs + OTHER ML PARADIGMS

Deep neural networks

- Cannot explain what they have learned
- Are data hungry
- Are computationally expensive to learn

But NNs and DNNs are not the only ML approaches

Variety of ML approaches

- trees, rules, kernel methods, (D)NNs

Variety of ML tasks

- Supervised, semi-supervised, unsupervised learning

Computational scientific discovery

Reinforcement learning

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

REINFORCEMENT LEARNING

- Agent learning from interaction with the environment
  - Performs actions
  - Receives feedback/ reinforcement
- Google’s AlphaGo

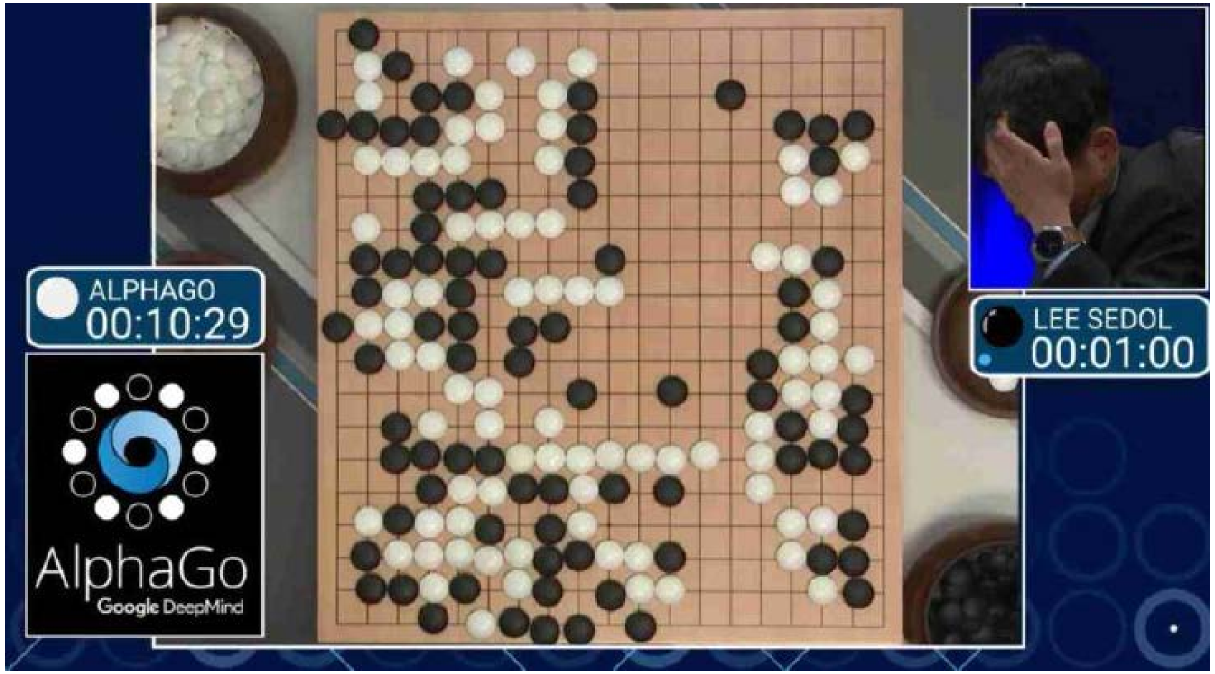

COMPUTATIONAL SCIENTIFIC DISCOVERY: FINDING LAWS FROM DATA

- Automate scientific activities, finding laws from data
- Earliest AI systems in the late 1970s (BACON)
- Input: Observations of distances between moons and Jupiter (d) and periods of their orbits (p)
- Output: Kepler’s third law $d^3/p^2 = k$
- Process of discovery
  - BACON (Langley 1978; Langley et al. 1987)
  - Carries out heuristic search through the space of numeric terms, looking for constant values and linear relations
- Latest work in the direction of robot scientists/ self-driving laboratories that automate the entire cycle

Moon	d	p
A	5.67	1.77
B	8.67	3.57
C	14.00	7.16
D	24.67	16.69

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

1. ARTIFICIAL INTELLIGENCE IN A NUTSHELL

2. KNOWLEDGE-BASED AI

3. DATA-BASED AI: MACHINE LEARNING

4. LARGE LANGUAGE MODELS

TRANSFER LEARNING

- To be able to learn from small data

- Exploit features/ representation learned on a large dataset
- Adapt the learned neural network (classification head)
  by using data from the small dataset

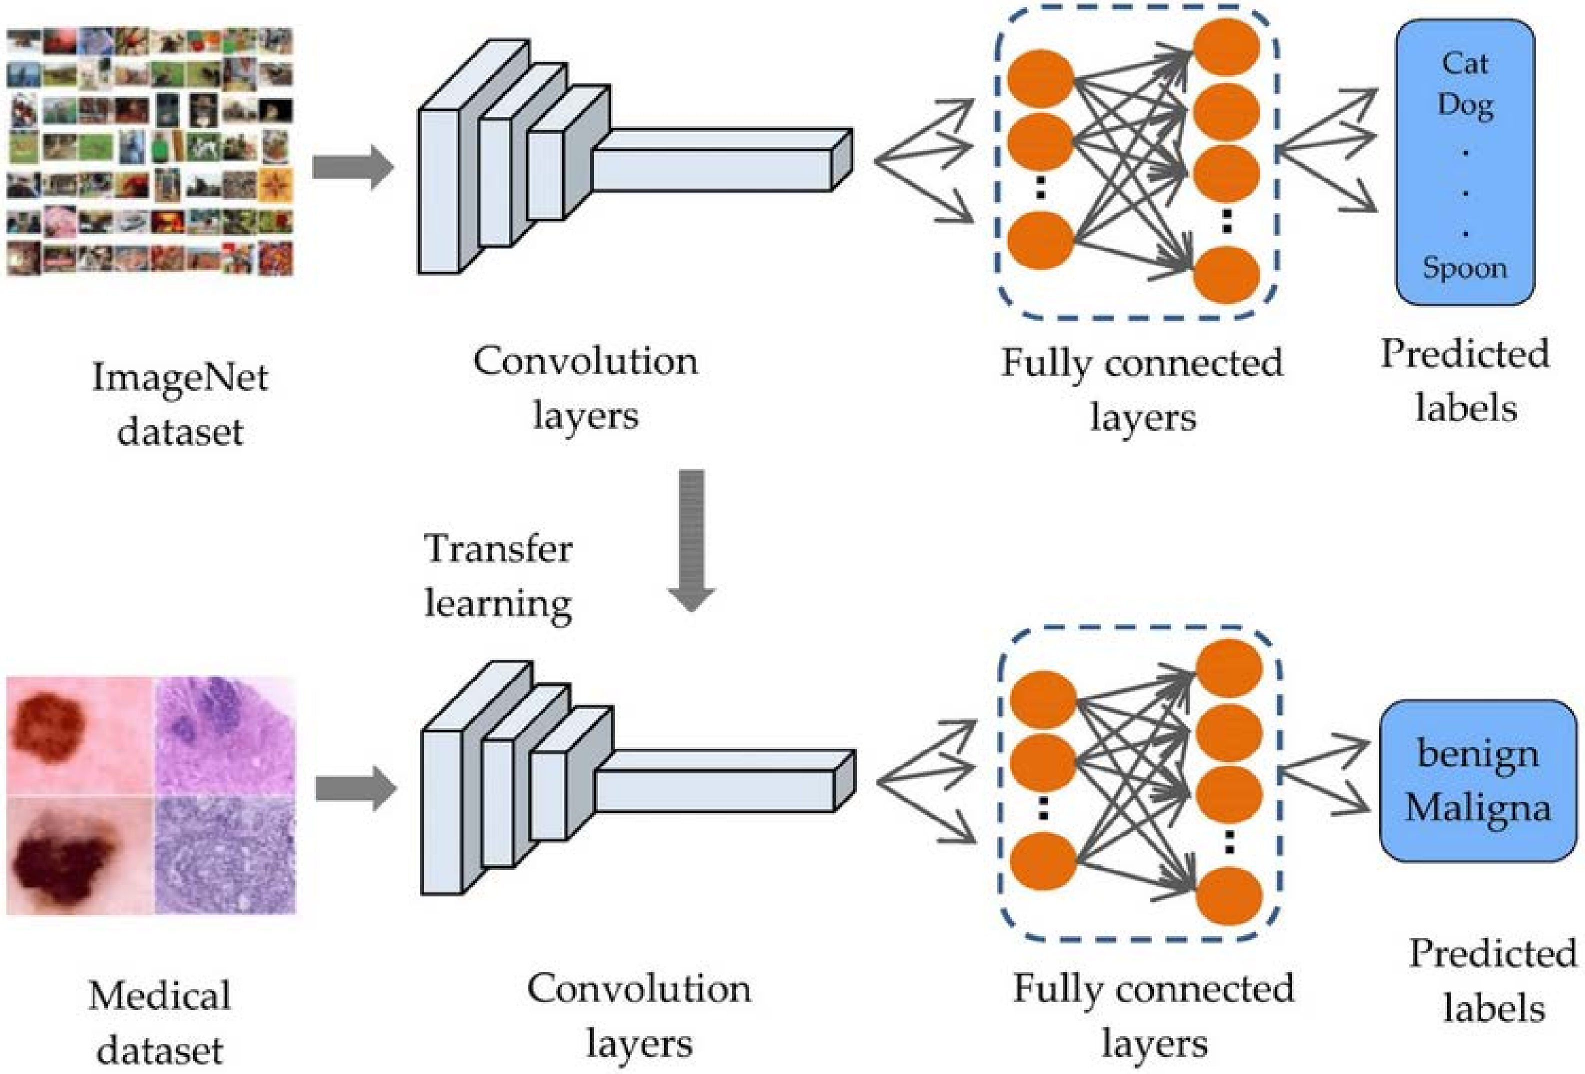

FOUNDATION MODELS

Foundation models (FMs) *are large models*
generated by applying ML (deep learning)
to a broad collection of data at scale
that can be adapted for use in a wide range of downstream tasks

Foundation models *are trained on pretext tasks,*
like reconstructing images or text from masked versions

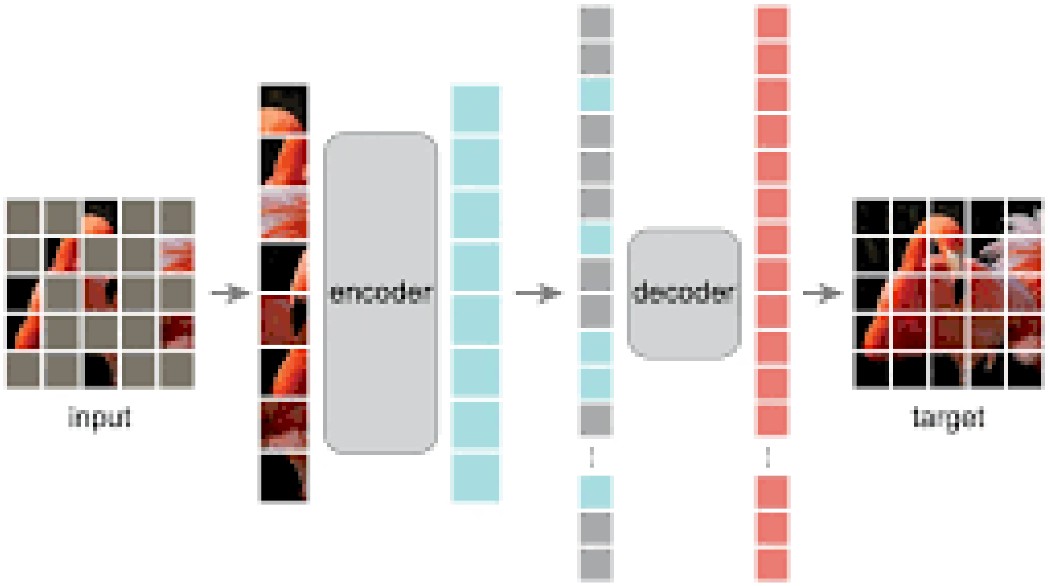

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

LARGE LANGUAGE MODELS

Large language models (LLMs) *are foundation models learned from large collections of text*

They are also trained on pretext tasks, like reconstructing text from masked versions or predicting the next word in a sequence

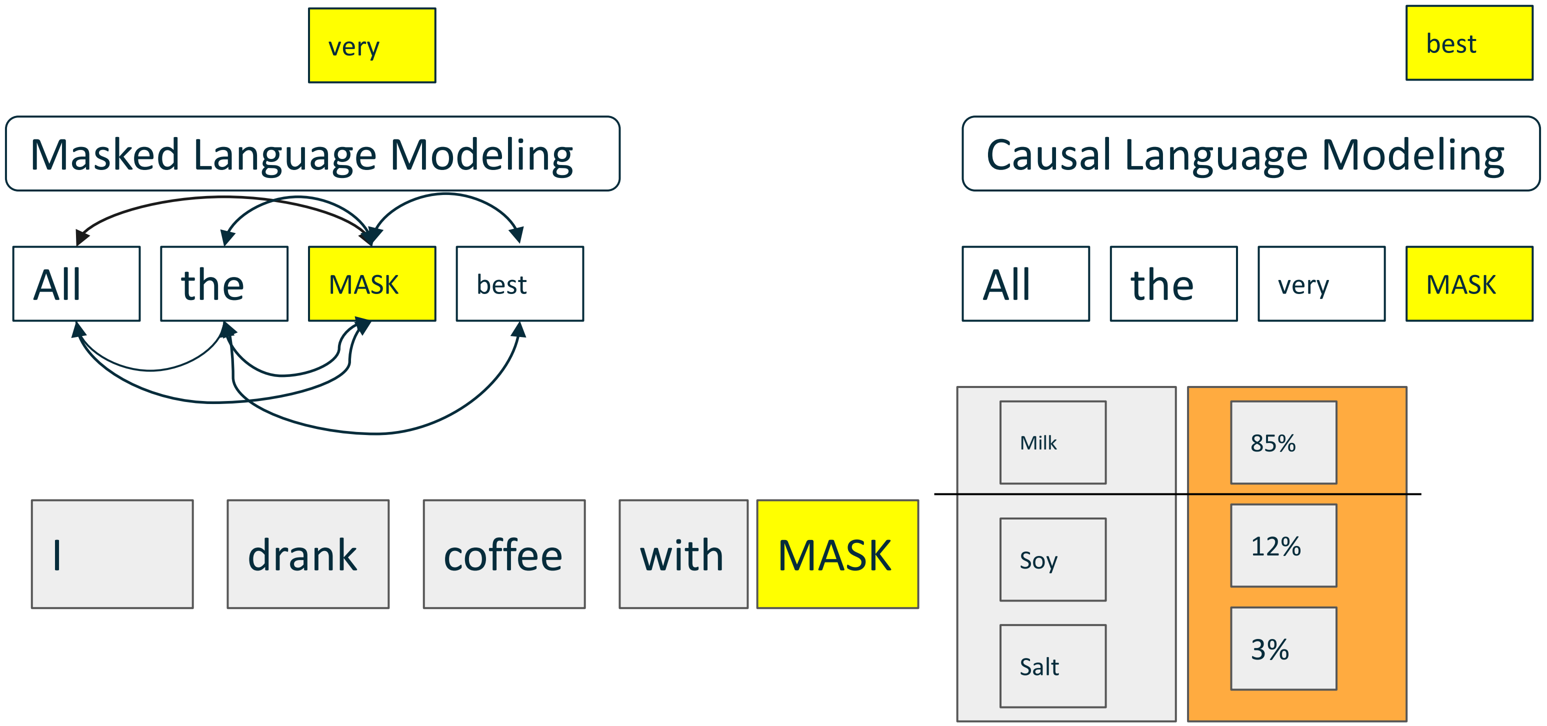

LARGE LANGUAGE MODELS AS GENERATIVE AI

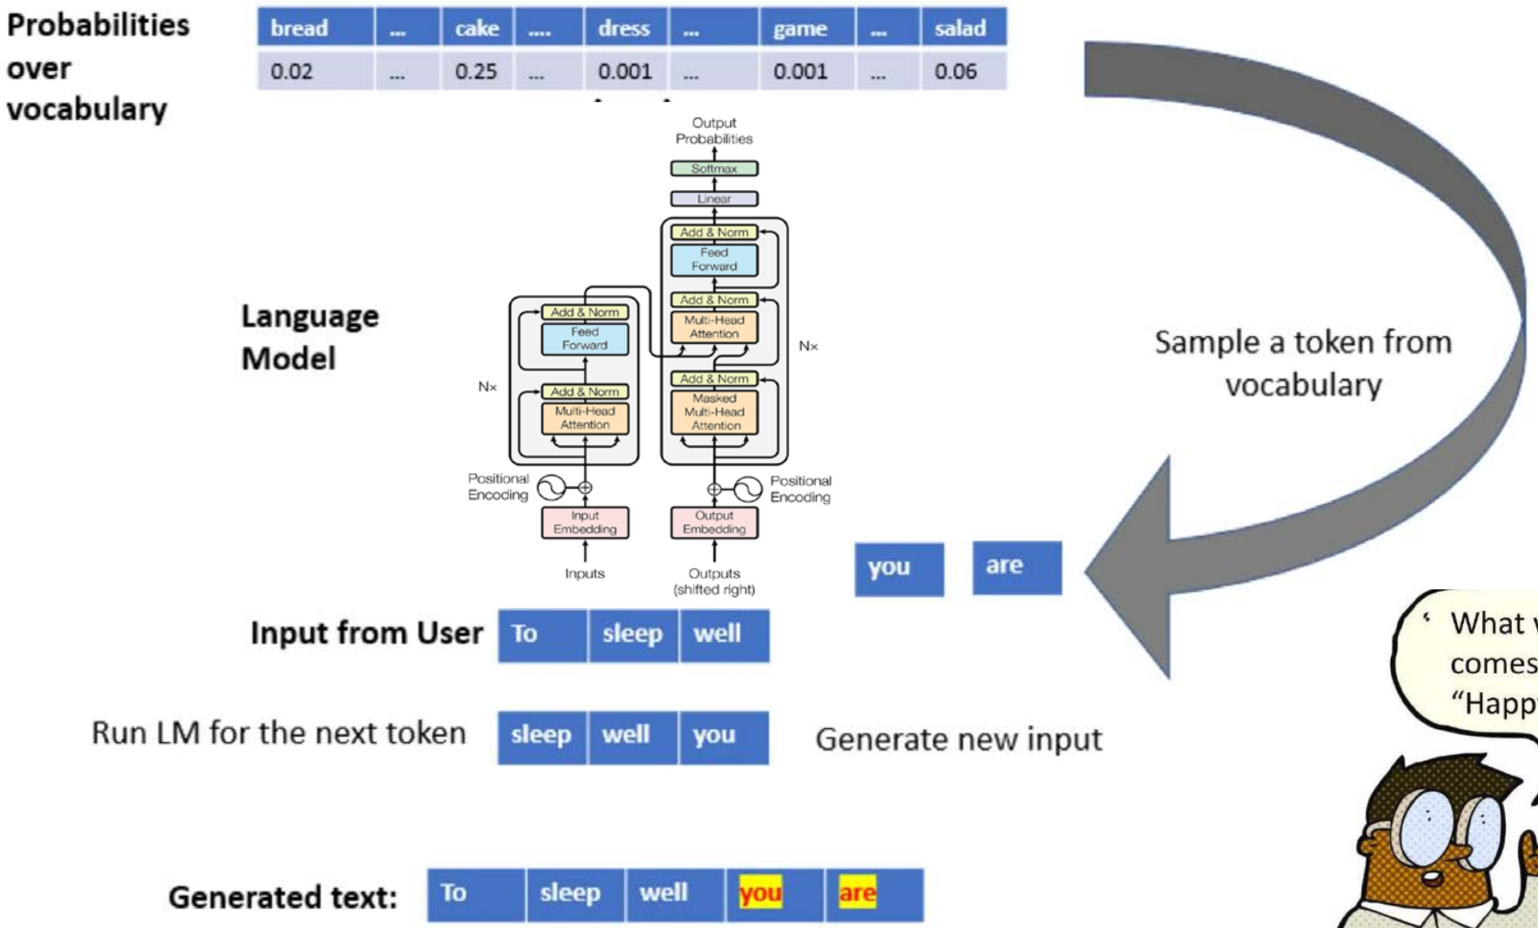

LARGE LANGUAGE MODELS (LLMs) are huge deep NNs with billions of parameters

They are trained on huge amounts of text

They predict the next word after given partial text and can be used to **generate** new text

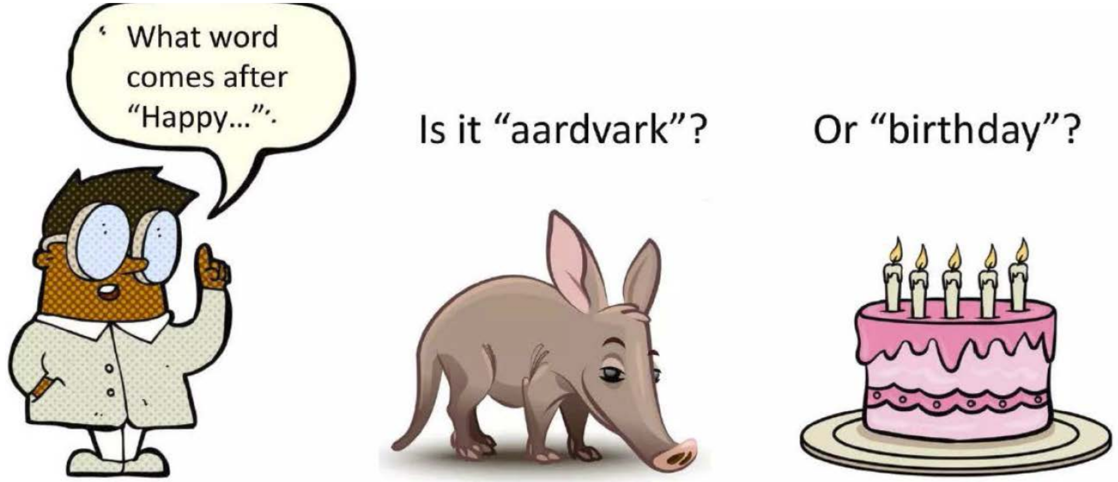

DOWNSTREAM TASKS FOR LARGE LANGUAGE MODELS

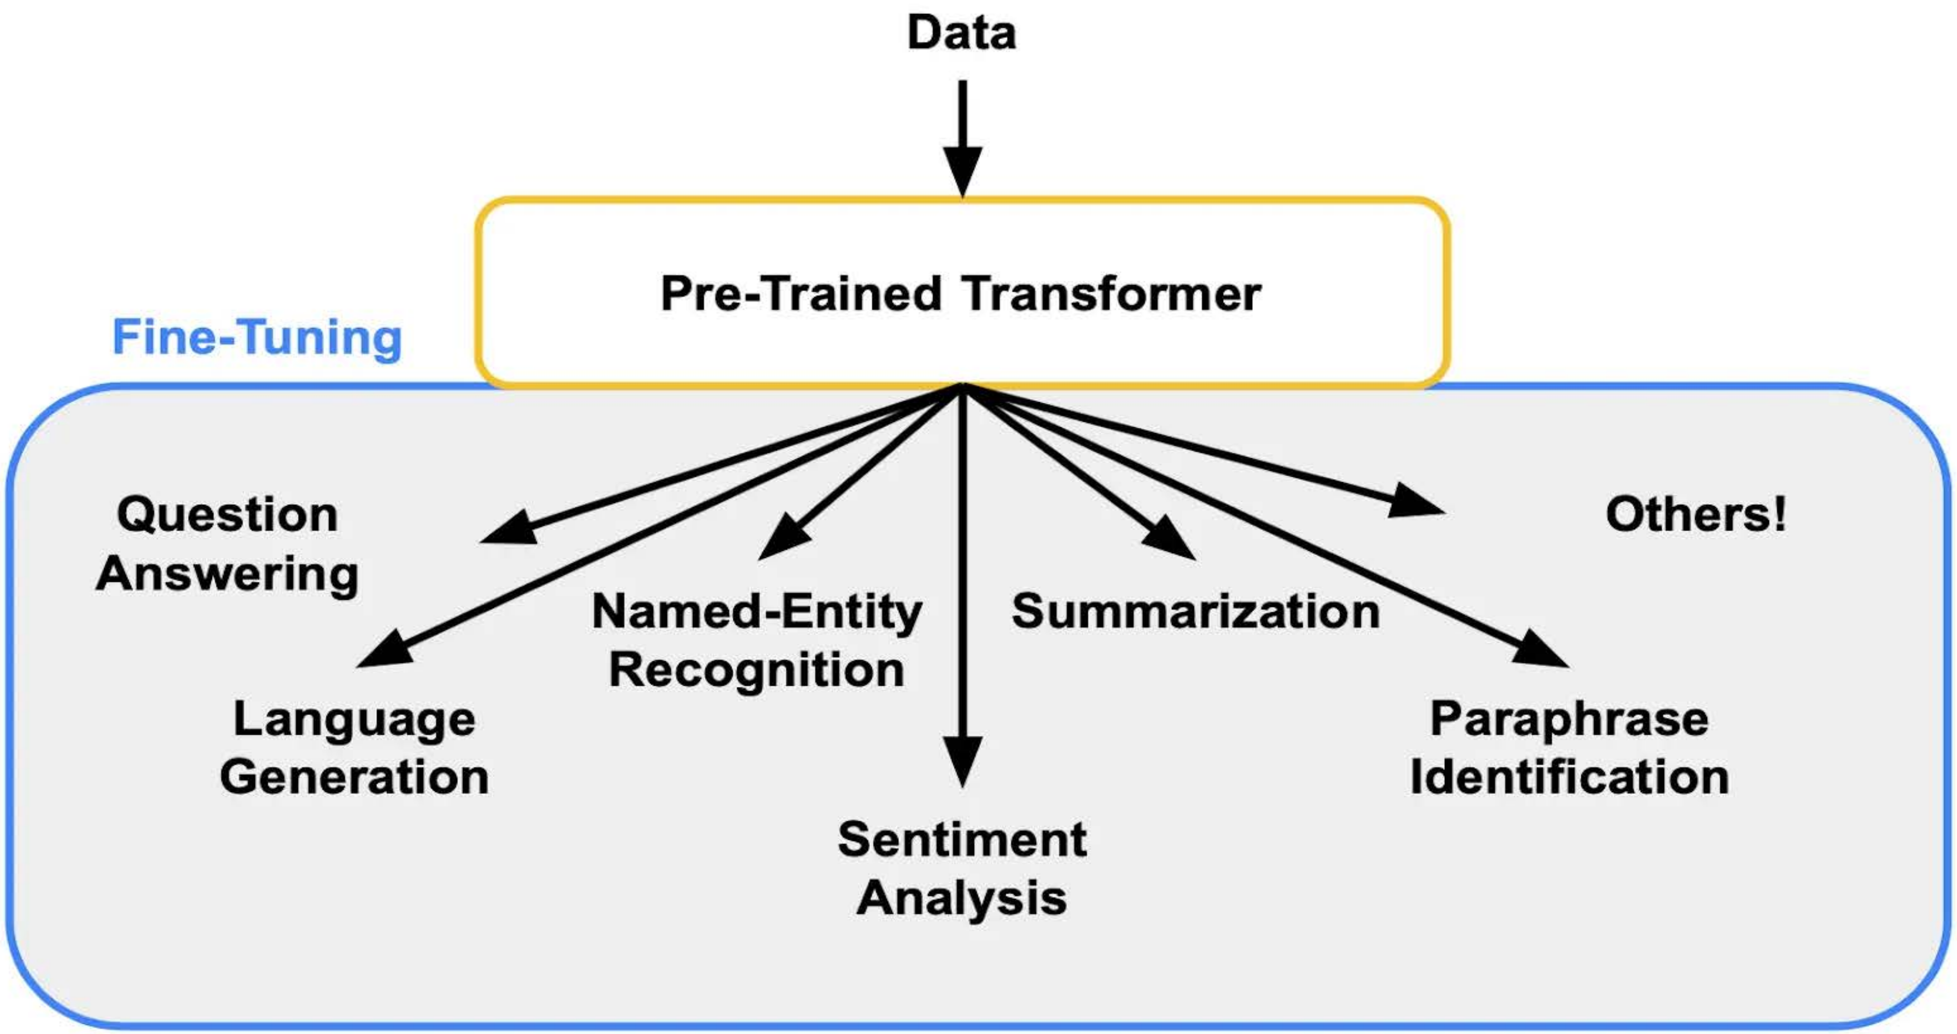

ChatGPT = LLM (GPT) + Conversational Interface

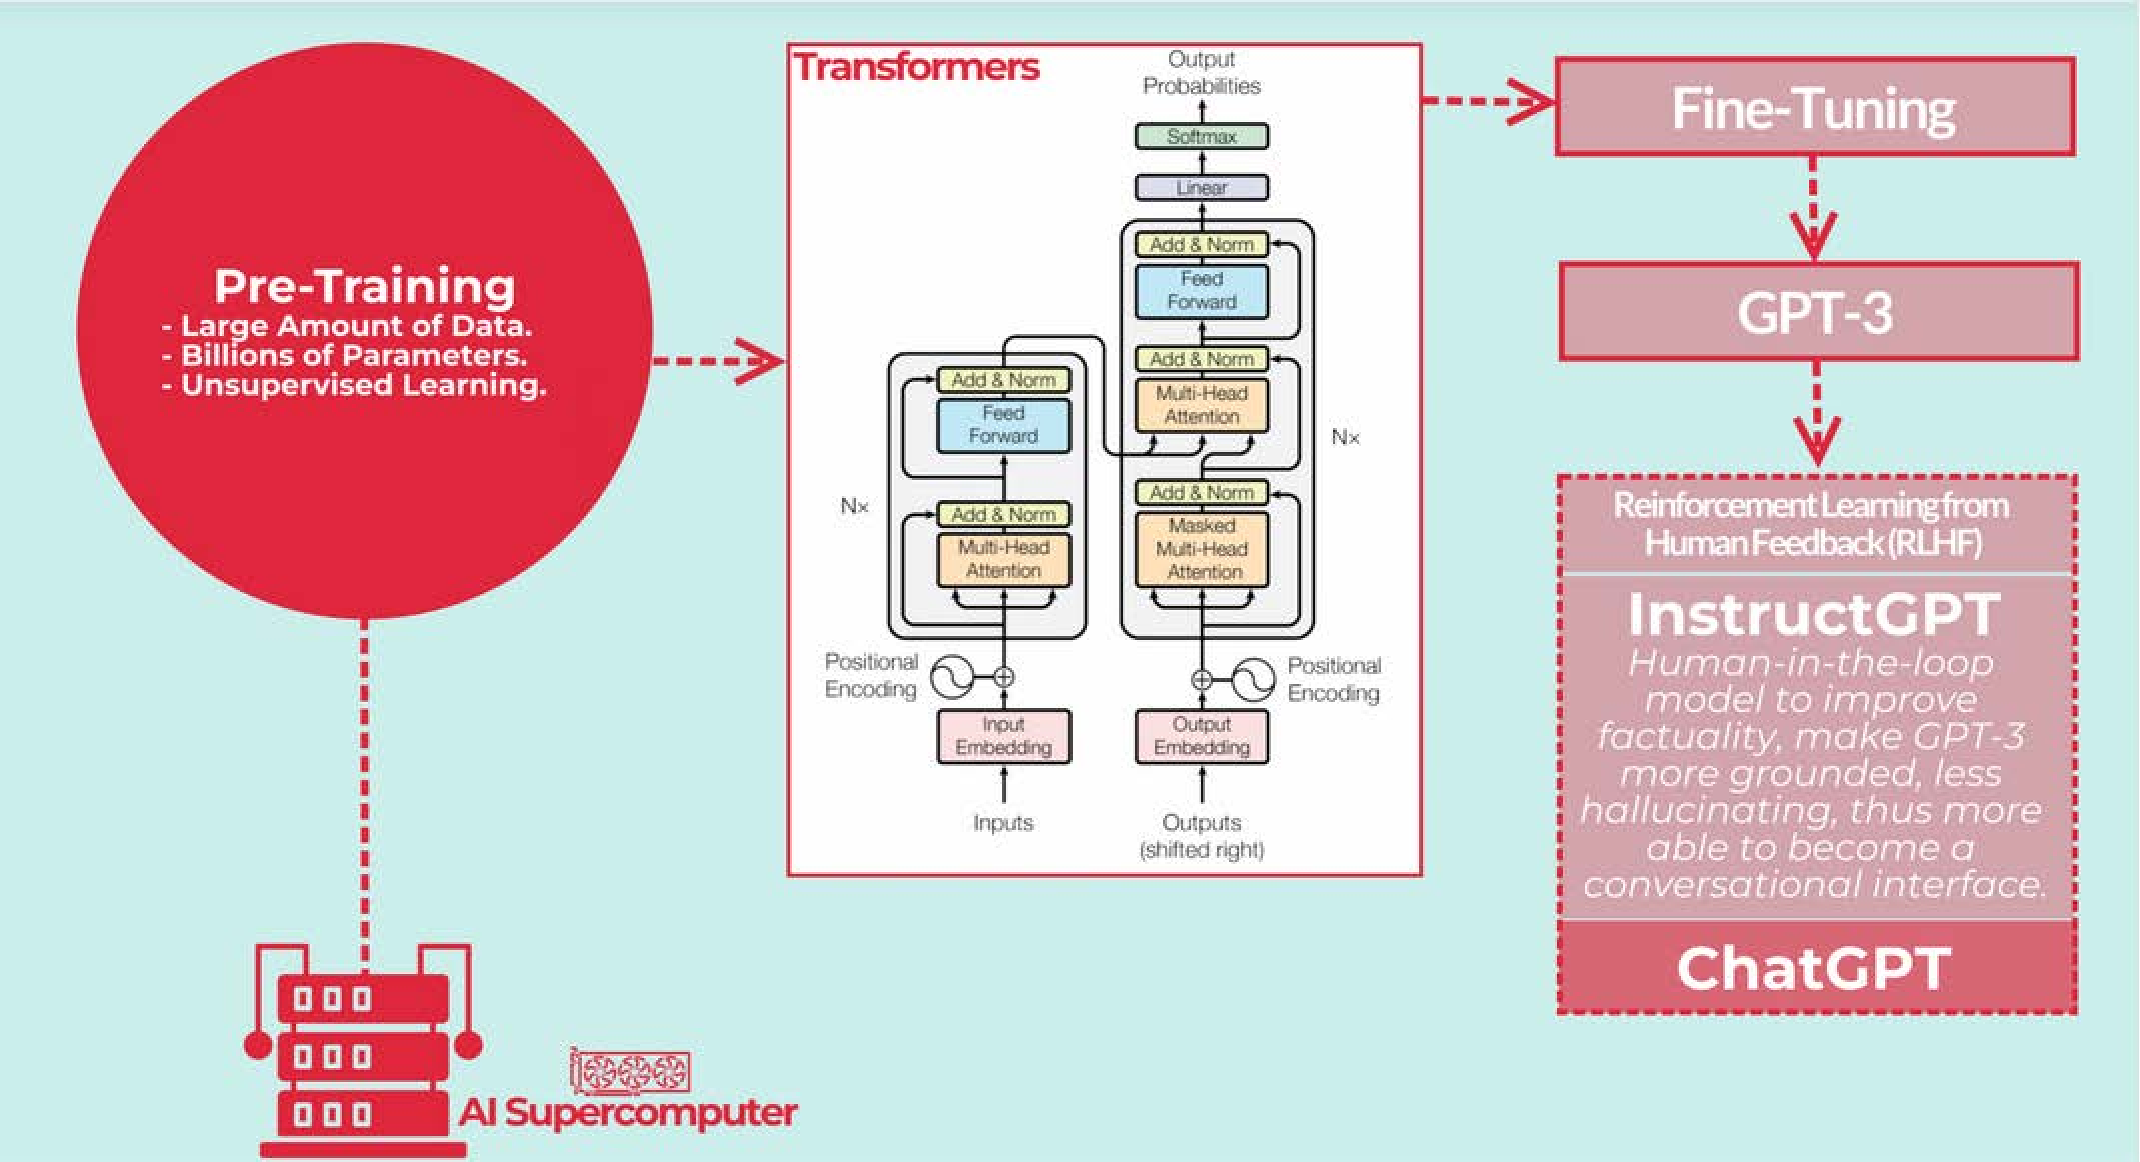

Fine-tuning by reinforcement learning from human feedback (RLHF)

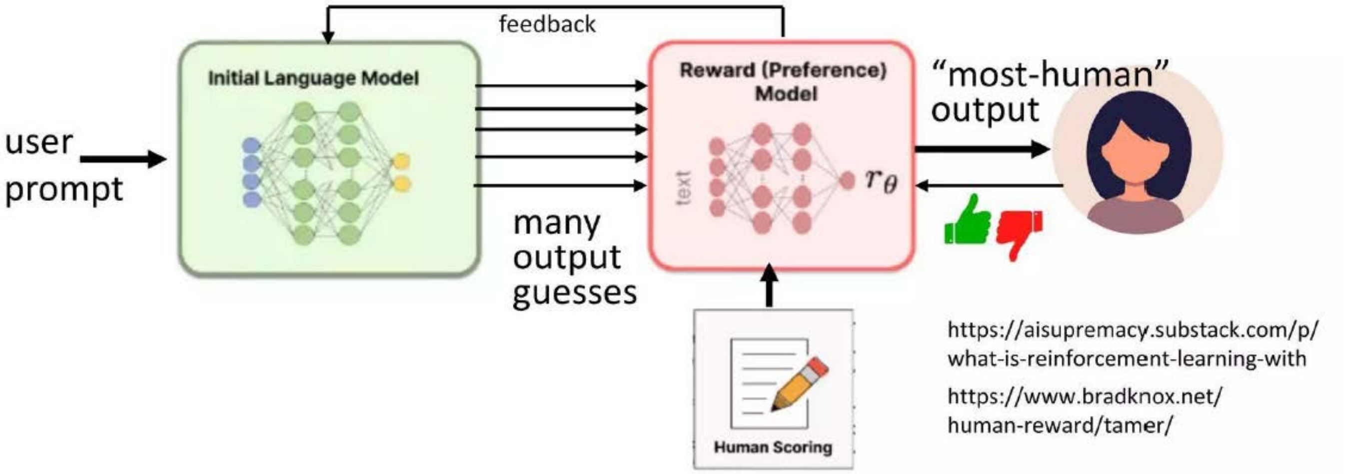

- 1. Humans generate an initial dataset of "typical" queries and "good" responses.
- 2. Humans rank a separate dataset of responses (some good, some bad).
- 3. The reward model is trained on this data.
- 4. The language model trains itself to satisfy the reward model. (How often?)
- 5. The reward model is retrained with (free!) input from hundreds of millions of users. (How often?)

THE NEURAL NETWORKS BEHIND THE FOUNDATION MODELS: THE TRANSFORMER ARCHITECTURE

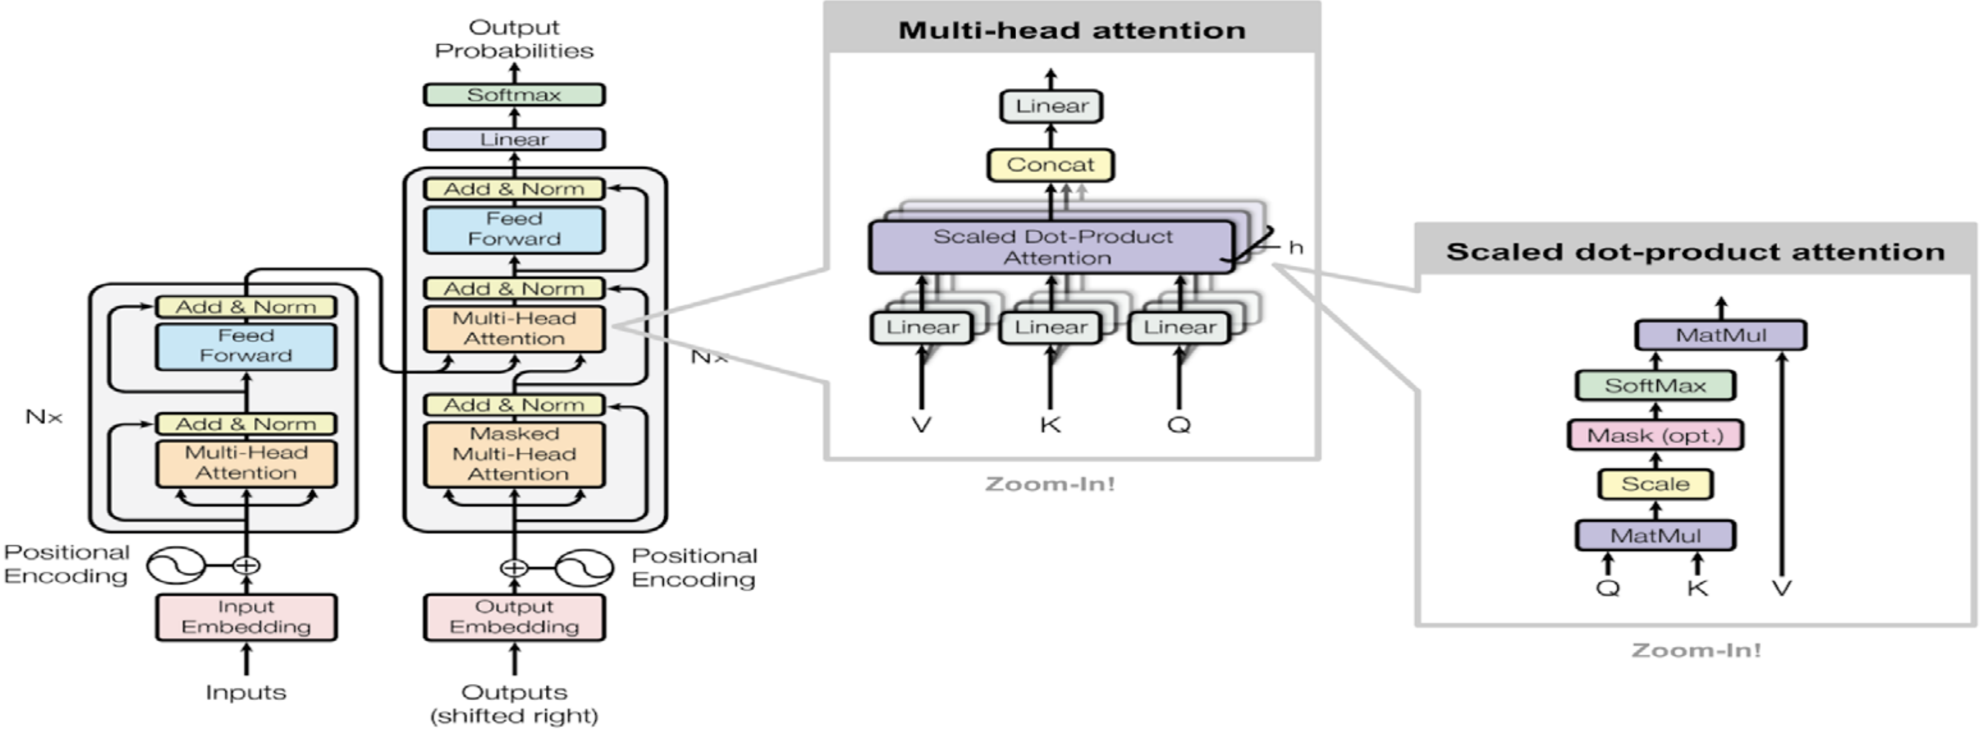

Very deep neural networks, with billions of parameters, require intensive computation (GPUs) for learning and inference

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

GENERATIVE AI: FROM RESPONSIBLE USE OF AI TO AI AS INFRASTRUCTURE

Sašo Džeroski

Head of Department of knowledge technologies, Jožef Stefan Institute, Ljubljana, Slovenia

AI IS A TOOL.
VERY POWERFUL TOOL.

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

USE IT RESPONSIBLY!

BUT WAIT. HOW DO I USE GEN-AI RESPONSIBLY?

PROBLEMS WITH LLMs

Existing LLMS have many flaws

- They can hallucinate
- They are expensive to update with new/ changing factual knowledge
- They can produce socially and ethically unacceptable outputs

Chanel's CEO went to Microsoft HQ and asked ChatGPT to show her a picture of her company's leadership. They were all men in suits

PROBLEMS WITH LLMs

Your data goes to the chatbot servers

- US for OpenAI/ChatGPT
- China for DeepSeek

It can be used for many purposes, incl. further training of LLMs

Do not send sensitive and personal data

Pentagon Staff Used DeepSeek's Chatbot for Days Before Block

- Defense employees connected to Chinese server, official says
- Popularity of DeepSeek in US has raised security concerns

The Pentagon building in Arlington, Virginia. *Photographer: Tom Brenner/Bloomberg*

RESPONSIBLE USE OF LLMs/GEN-AI

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

Data protection

Verification

Transparency

Principle #1: Protect data and information

Many popular GenAI tools collect and store information from users as part of their learning process. While some GenAI services offer a “no training use” policy as part of enterprise agreements, by default, researchers should first assume that any data or information entered into web-based GenAI could become part of its training data and thus be recoverable by third parties from future iterations of the model. Terms for the tool’s use of information entered by the user should first be verified by the researcher in advance: do not assume that data entered is private, that output is already in public domain, nor that one can share information that is private, export-controlled, sensitive, or proprietary.

Principle #2: Verify

Researchers have a duty to verify the output of GenAI systems to ensure that they have not introduced inaccuracies or biases into the research process. GenAI can produce extremely convincing text, imagery, audio and video (and increasingly other types of data); but GenAI tools may not stay faithful to some ground truth or fairly represent the world.

Use of GenAI carries real risks regarding the reproducibility and reliability of research. This is particularly true in the context of research involving cultural practices or human diversity, where the risk of bias is high.

Principle #3: Be transparent

Consistent with disciplinary and publishing norms, the use of GenAI should be disclosed/identified at the time of proposing, reporting, publishing or when engaging in other dissemination activities. In reporting such use, the following existing policies should be adhered to:

- **Research Misconduct Policy**
- **Intellectual Property Policy**

While norms around the reporting of GenAI use are evolving, we urge researchers to err on the side of acknowledging and **reporting all use of GenAI tools in the research process**, including use in idea development, assistance with composition, code assistance, or generation of presentations and figures.

UNIVERSITY of ROCHESTER

UNESCO (2023) — Guidance for Generative AI in Education and Research

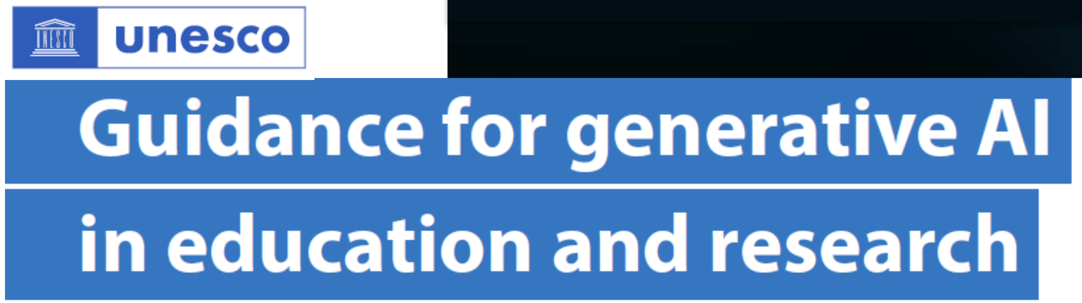

Key risks for researchers

- **Hallucination**: GenAI produces confident but false outputs — including fake citations
- **Data privacy:** your inputs go to servers in the US (OpenAI) or China (DeepSeek) and may be used for further training
- **Black box**: outputs cannot be explained or verified
- **Copyright:** training data scraped without consent; outputs may violate IP rights

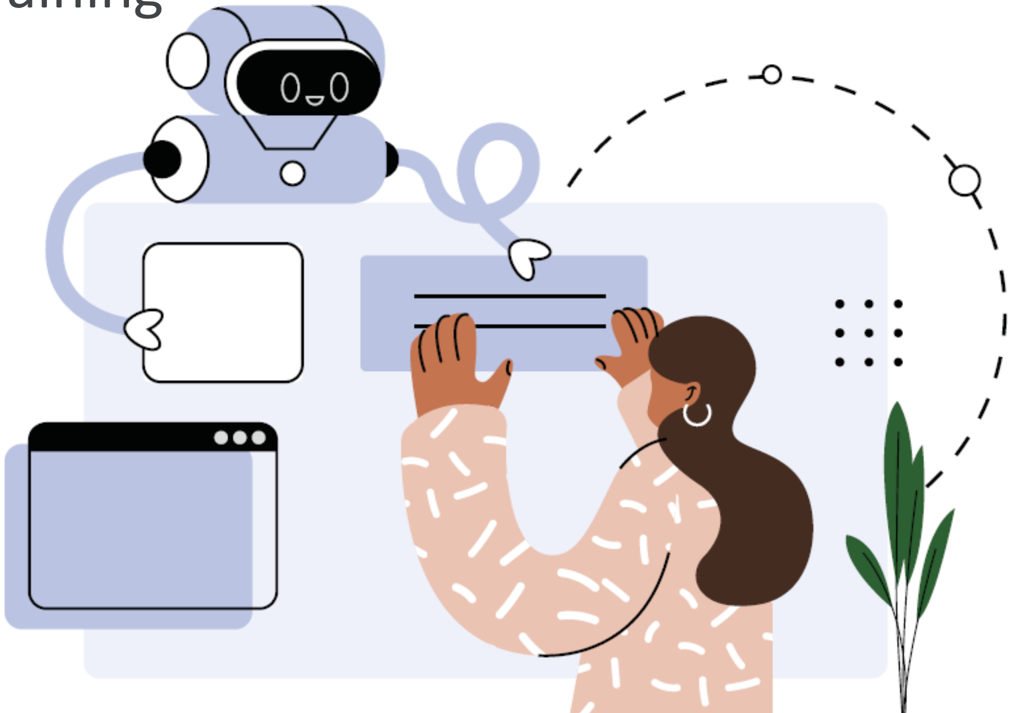

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

UNESCO (2023) — Guidance for Generative AI in Education and Research - Recommendations

Researchers should:

- Follow research integrity principles; remain ultimately responsible for scientific output
- Use GenAI transparently and maintain a critical approach to all outputs
- Preserve privacy and confidentiality of research data — do not send sensitive data to commercial tools
- Refrain from using GenAI for peer review or grant evaluation

Research organisations should:

- Guide responsible use and monitor how tools are developed and used
- Deploy their own GenAI tools to ensure data protection and confidentiality

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

Living guidelines on the RESPONSIBLE USE OF GENERATIVE AI IN RESEARCH

AI is transforming every stage of the research process. GenAI tools can facilitate scientific work and accelerate discovery — **when used in the right way.**

The European Commission, countries and ERA stakeholders have collaboratively developed recommendations to support the responsible integration of GenAI in research.

The guidelines follow the principles of **research integrity** and address the main challenges researchers face when using generative AI.

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ERA Living guidelines on the Responsible Use of Generative AI in Research

RESEARCHERS should...

- Follow key principles of research integrity, use GenAI **transparently** and remain **ultimately responsible** for scientific output
- Use GenAI preserving **privacy, confidentiality, and intellectual property rights** on both inputs and outputs
- Maintain a **critical approach** to using GenAI and ***continuously learn how to use it responsibly*** — AI literacy
- **Refrain** from using GenAI tools in sensitive activities, e.g. peer reviews or evaluations

ERA Living guidelines on the Responsible Use of Generative AI in Research

RESEARCH ORGANISATIONS should…

- Guide the responsible use of GenAI by their researchers and actively monitor how they develop and use such tools
- Integrate and apply these guidelines, adapting or expanding them when needed
- **Deploy their own GenAI tools** to ensure data protection and confidentiality

FUNDING ORGANISATIONS should…

- Support the responsible use of GenAI in research
- Use GenAI transparently, ensuring confidentiality and fairness
- Facilitate the transparent use of GenAI by applicants

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

ARTIFICIAL INTELLIGENCE AS INFRASTRUCTURE FOR SCIENCE

A slightly different view of infrastructure

Infrastructure is not only physical equipment

Infrastructure is also

- *(SW) tools*
- *services*
- *other resources (e.g. data)*

An excellent example is CLARIN & CLARIN.si

Language resources as infrastructure

- *Corpora (text, speech)*
- *Lexica*
- *Language tools (e.g., concordancers)*

CLARIN.SI

http://www.clarin.si/

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

llm.ijs.si AS INFRASTRUCTURE FOR SCIENCE

Local installation(s) of
Large Language Models

- For use in science/ by scientists, but also supporting personnel

- Chat assistant
- Different LLMs

- Programatic access (API)
- Fine tuning of LLMs with own data

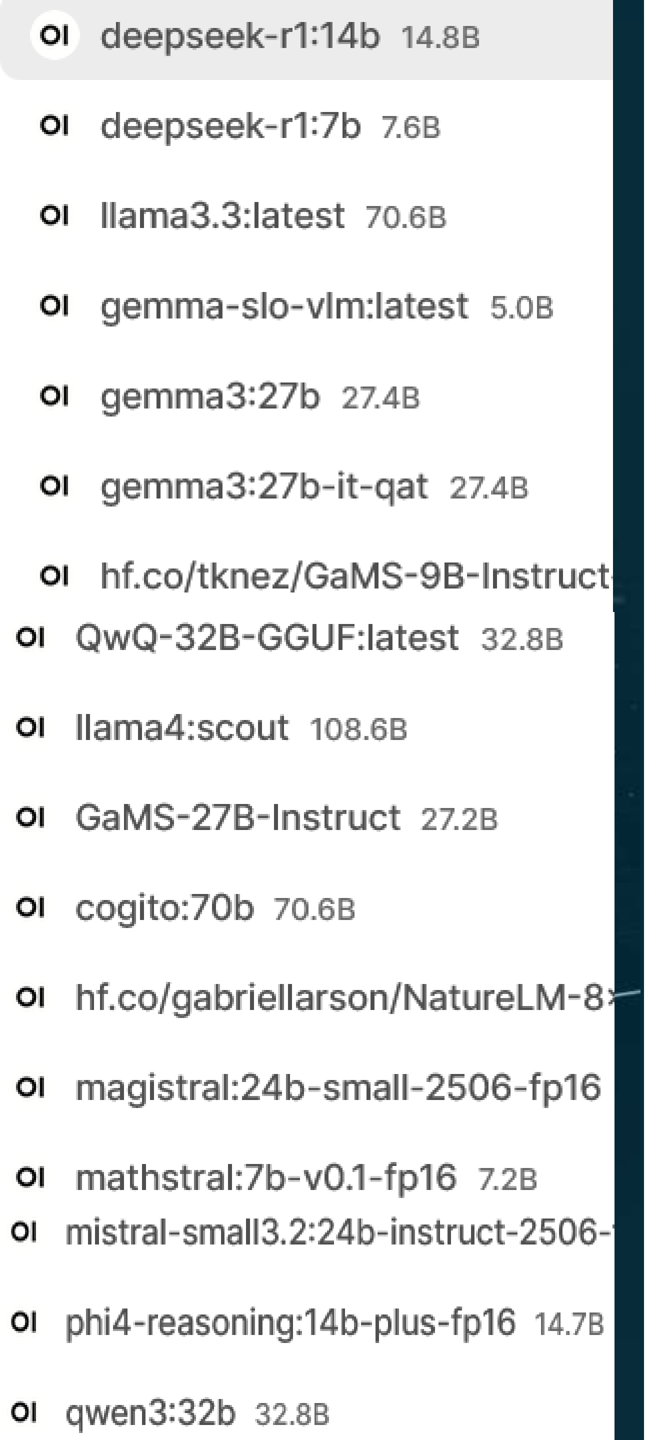

AI AS INFRASTRUCTURE FOR SCIENCE (llm.ijs.si)

Local infrastructure at JSI, serving LLMs to researchers at JSI

At the end of March 2024, we released a preliminary service, based on the LLM model Mistral7B

By October 2024, we had added a few additional LLMs

We have a local setup of open-sourced LLMs, hosted on JSI infrastructure, available via a chat assistant & programmatic API

- Fast and scalable implementation using Paged Attention
- Up to few thousand concurrent connections

Different LLMs available at llm.ijs.si for different purposes

- Multiple models are implemented concurrently, including DeepSeek and MathModels – Mathstral, which are specialized in custom domains.
- Current list (left) includes GAMS, a LLM for Slovene
- This specialization allows them to deliver tailored solutions for specific applications, enhancing their effectiveness and performance in niche areas.

llm.ijs.si: Simultaneous query to multiple LLMs

deepseek-r1:14b ⌄ +
llama3.3:latest ⌄ −
gemma3:27b ⌄ −

How can AI help scientists?

deepseek-...

AI, or Artificial Intelligence, can significantly assist scientists in various ways, enhancing their research capabilities and efficiency. Here are some key areas where AI contributes to scientific advancements:

1. **Data Analysis**: AI excels at processing large datasets quickly and efficiently. Scientists can use AI algorithms to analyze complex data from experiments,

llama3.3:latest

AI can help scientists in numerous ways, transforming the way they conduct research, analyze data, and make discoveries. Here are some examples:

1. **Data analysis**: AI algorithms can quickly process large amounts of data, identifying patterns and correlations that may be difficult for humans to detect. This is particularly useful in fields like genom astronomy,

gemma3:27b

How AI is Helping Scientists: A Deep Dive

AI is rapidly becoming an indispensable tool for scientists across *all* disciplines. It's not about replacing scientists, but *augmenting* their abilities, speeding up research, and uncovering insights previously impossible to find. Here's a breakdown of how AI is helping, categorized for

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

The llm.ijs.si chat assistant

- Custom-built chat assistant
- Powered by four RTX 3090 GPUs, each with 29GB of memory, for enhanced performance
- If you have an @ijs.si email address, get your account NOW :-)

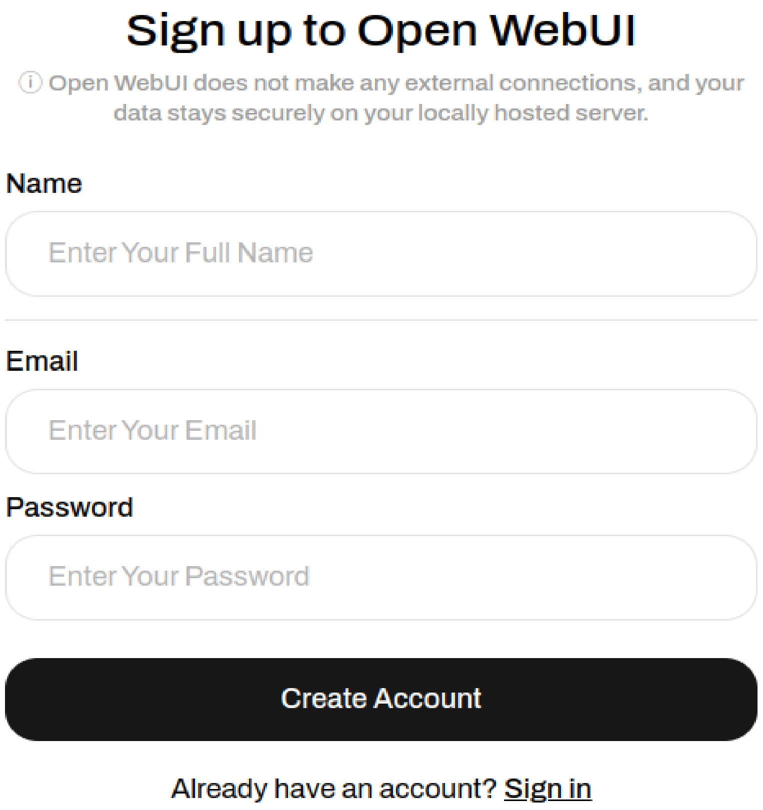

DEPARTMENT OF KNOWLEDGE TECHNOLOGIES
Jožef Stefan Institute

The benefits of local installation(s) of Large Language Models

No data leakage *- enables secure and efficient inference of pretrained models without the need to share or transmit the data out of JSI*

Reproducibility *– In proprietary large language models, we cannot control all variables/parameters. With local instantiations, we can control the flow of information and ensure reproducible results, enabling us to contribute to open and reproducible science*

Lower costs *- we only need to pay for the compute time on our own infrastructure, bearing no additional costs per call (as opposed to the proprietary models such as ChatGPT…)*

AI AS INFRASTRUCTURE: SUPPORTING RESPONSIBLE USE OF AI

The problem: Commercial LLMs send your data to servers in the US or China — unacceptable for sensitive research data

The ERA guideline: Research organisations should *deploy their own GenAI tools* to ensure data protection and confidentiality

JSI's answer — llm.ijs.si

- Open-source LLMs hosted on JSI infrastructure
- Available as chat interface and programmatic API
- Multiple models: Llama, DeepSeek, Mathstral and others
- Data stays inside JSI — no third-party exposure
- Available to all JSI researchers today

Slovenia's answer — SLAIF

- The Slovenian AI Factory: national AI infrastructure
- Same philosophy, national scale
- AI-for-Science vertical: dedicated support for researchers
- Sovereign AI compute: data stays in Slovenia

llm.ijs.si is JSI's implementation of ERA best practice.

SLAIF extends this to the national level.