ARTIFICIAL INTELLIGENCE 2

Introduction

Francesco Spinnato

* francesco.spinnato@unipi.it
University of Pisa

Artificial Intelligence 1 - Recap

- **Data Understanding and Preprocessing**
  - **data cleaning, integration, reduction, transformation and discretization.**
- Frequent pattern mining:
  - basic concepts, A-priori algorithm, Pattern-Growth approach, vertical data format, pattern evaluation methods, constraint-based frequent pattern mining.
- Clustering:
  - basic concepts, partitioning methods, hierarchical methods, density-based methods, grid-based methods, model evaluation and selection, constraints.
- Outlier detection:
  - statistical, proximity-based, clustering-based and classification-based.
- **Classification**:
  - **basic concepts, decision tree induction, Bayes classification methods, rule-based classification, lazy learners, techniques for improving accuracy, model evaluation and selection.**

What's the "2" in Artificial Intelligence 2?

Beyond tabular data... towards unstructured data:

- time series,
- images,
- text,
- graphs,
- multimodal!

Beyond simple supervised models... towards:

- **deep learning** with neural networks like CNNs, RNNs, transformers, autoencoders...

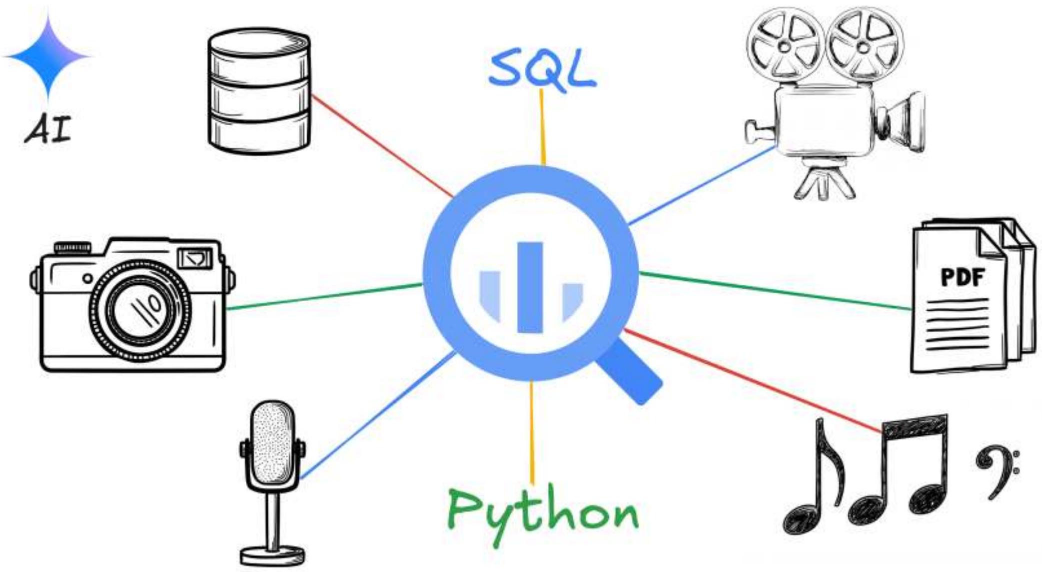

Focus of the Course

Models/Algorithms
Deep learning

Applications/Data
Healthcare

Artificial Intelligence, Machine Learning, Deep Learning

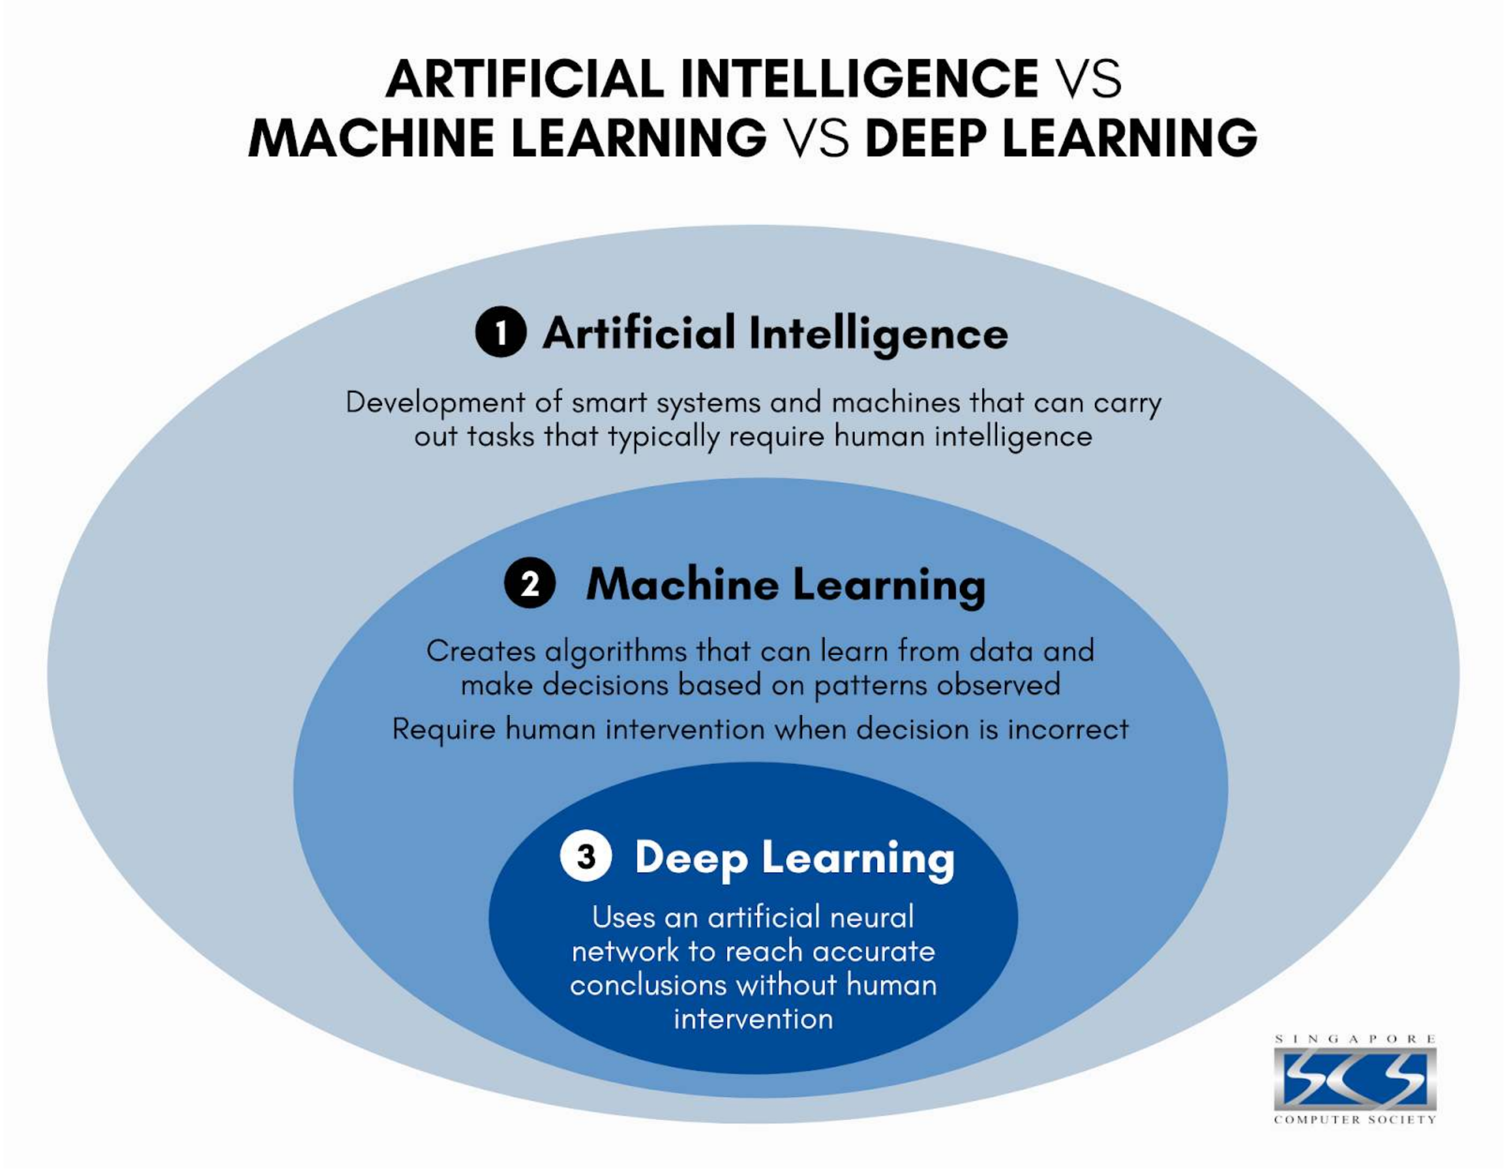

What is Artificial intelligence?

Artificial Intelligence is:

an automated process to carry out tasks
that usually require human intelligence.

What is Machine Learning?

Machine Learning is:

a set of tools/models/algorithms that learn from data
and make decisions and predictions based on the observed patterns.

What is Machine Learning?

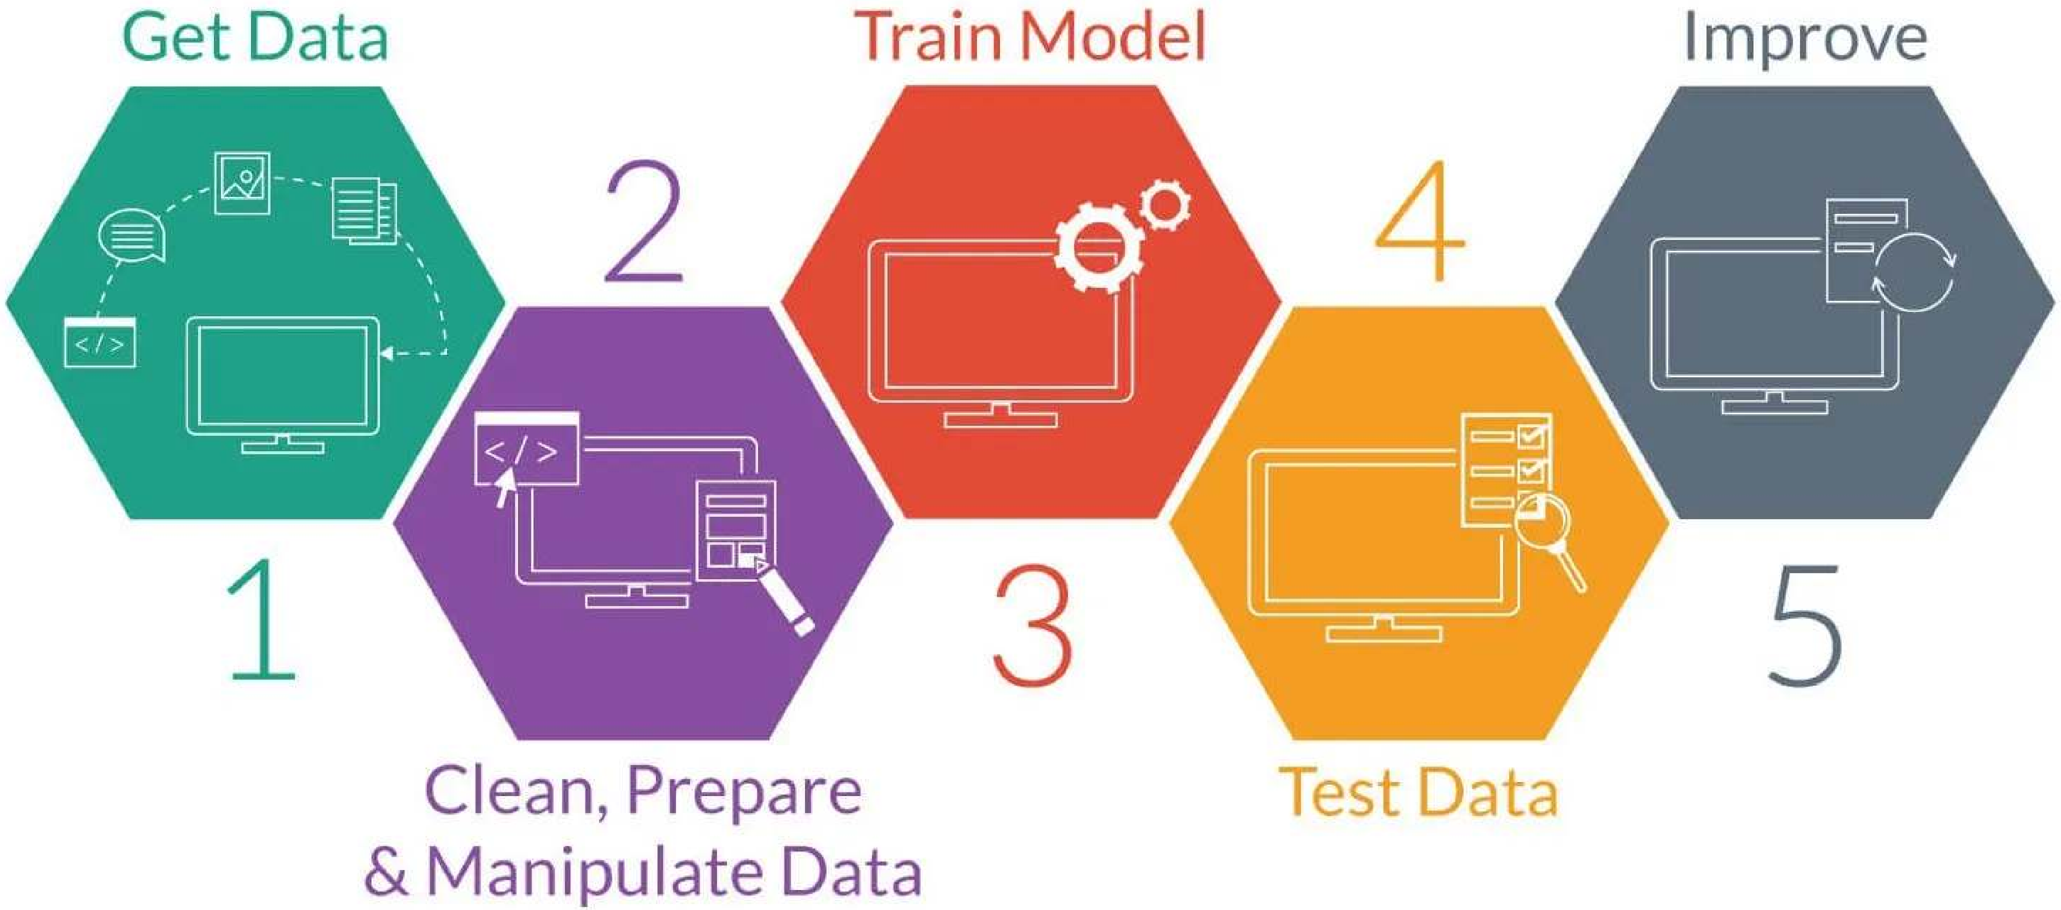

What is Deep Learning?

Deep Learning is:

the subset of machine learning models that use
Artificial Neural Networks
to make decisions without human intervention.

What is Deep Learning?

Deep learning algorithms are inspired by the human brain.

- The human brain is made of billions of neurons connected by synapses.
- Neurons receive signals, combine them, and activate if a threshold is reached.
- Learning occurs by strengthening or weakening synaptic connections.
- Neural networks mimic this idea using nodes, weights, and activation functions.

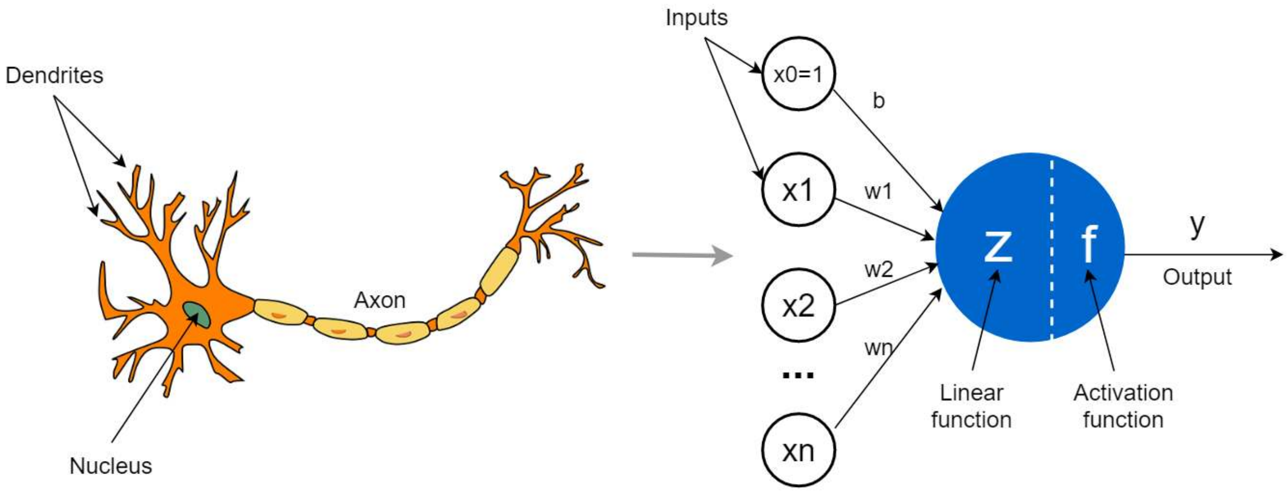

What is Deep Learning?

The brain processes information through many hierarchical layers (e.g., retina → visual cortex → higher cortex).

- Deep neural networks **mimic** this hierarchy using multiple hidden layers.
- Each layer extracts increasingly abstract features.
- Training adjusts weights to learn patterns from data.
- At the end of the day, it's just math + computer science

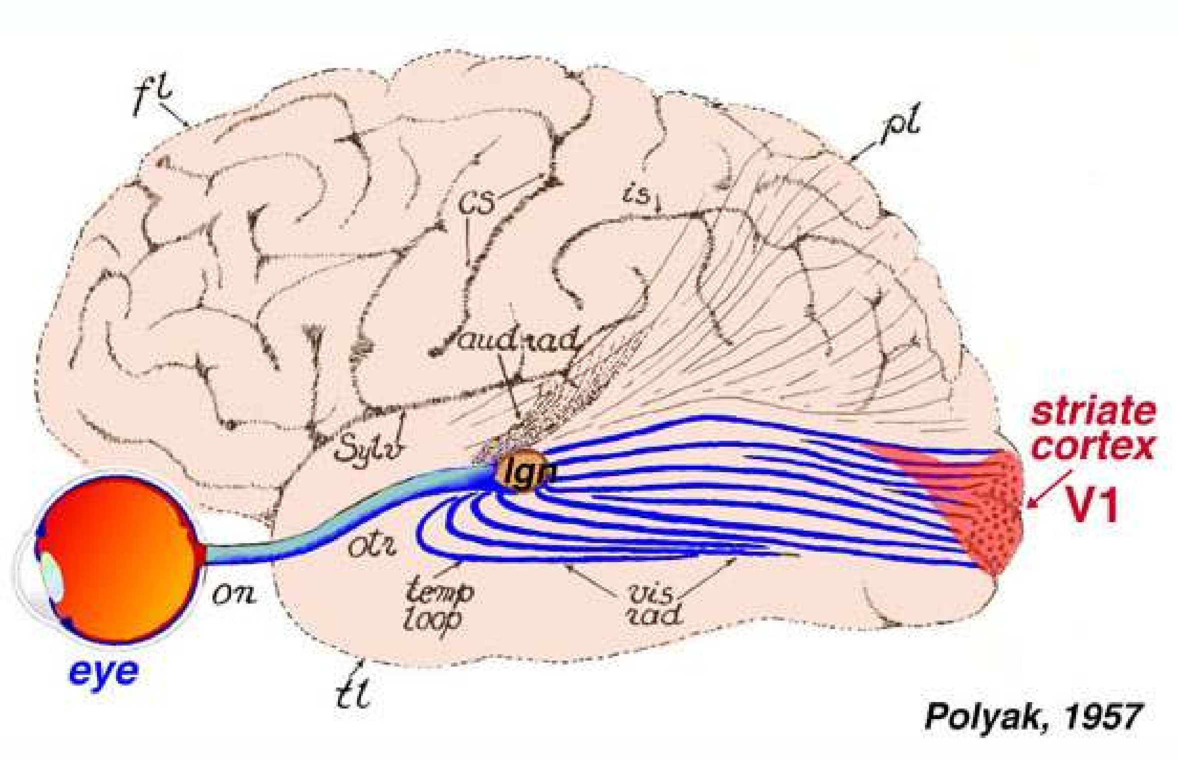

Figure 8. Visual input to the brain goes from eye to LGN and then to primary visual cortex, or area V1, which is located in the posterior of the occipital lobe. Adapted from Polyak (1957).

Before Modern Deep Learning

Math/CS tools at the core of deep learning:

- **1676:** Chain rule (mathematical foundation of backpropagation)
  Gottfried Wilhelm Leibniz
- **Circa 1800:** Linear regression as an early neural-like model
  Adrien-Marie Legendre, Carl Friedrich Gauss

https://people.idsia.ch/~juergen/deep-learning-history.html

Early Deep Learning Foundations

- **1920s:** Early recurrent network concepts
  Wilhelm Lenz, Ernst Ising
- **1958:** Multilayer feedforward neural networks
  Frank Rosenblatt
- **1965:** First working deep learning systems (multi-layer learning). *Alexey Ivakhnenko, Valentin Lapa*
- **1967-1968:** Deep learning with stochastic gradient descent. *Alexey Ivakhnenko*
- **1970s:** Early formulations of backpropagation
  Seppo Linnainmaa
- **1986:** Backpropagation popularized for neural networks
  David Rumelhart, Geoffrey Hinton, Ronald Williams

---

https://people.idsia.ch/~juergen/deep-learning-history.html

Modern Deep Learning Emerges

- **1990-1991 (“Annus Mirabilis”):**
  CNNs, RNNs, LSTM, efficient backpropagation variants
  Yann LeCun, Jürgen Schmidhuber, Sepp Hochreiter
- **1990s:**
  - Convolutional Neural Networks (LeNet)
    Yann LeCun
  - Long Short-Term Memory (LSTM)
    Sepp Hochreiter, Jürgen Schmidhuber

https://people.idsia.ch/~juergen/deep-learning-history.html

Representation Learning Revival

- **2000s:** Improved hardware and larger datasets
- **2006:** Deep belief networks and unsupervised pretraining
  Geoffrey Hinton, Simon Osindero, Yee-Whye Teh
- Renewed interest in deep architectures

https://people.idsia.ch/~juergen/deep-learning-history.html

Hardware Revolution

Deep Learning can solve complex problems **accurately** and **efficiently**.

Data

Algorithms

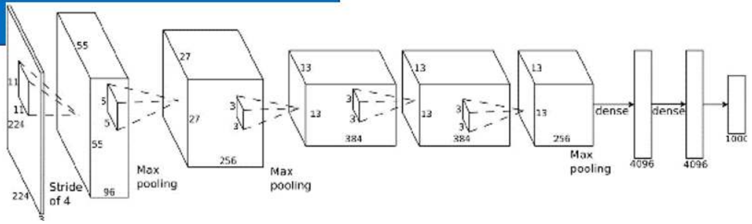

Hardware

Software

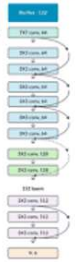

- **Data:** availability of large, rich datasets.

Deep Learning can solve complex problems **accurately** and **efficiently**.

Data

Algorithms

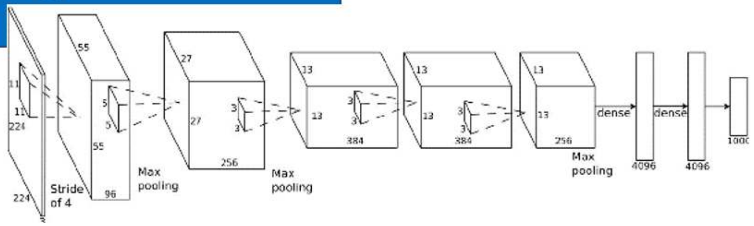

Hardware

Software

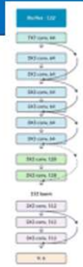

Keras

TensorFlow 2.0

- **Algorithms:** improved neural network methods.

Hardware Revolution

Deep Learning can solve complex problems **accurately** and **efficiently**.

Data

Algorithms

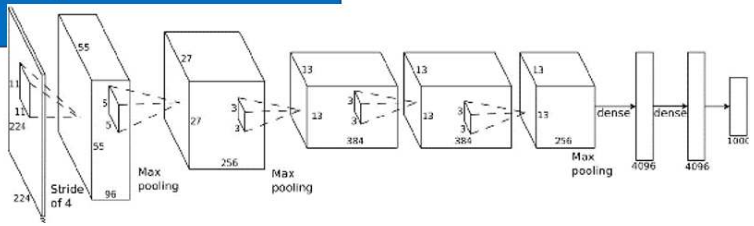

Hardware

Software

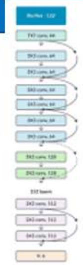

Keras

TensorFlow 2.0

- **Hardware:** GPU-based parallel computing.

Deep Learning can solve complex problems **accurately** and **efficiently**.

Data

Algorithms

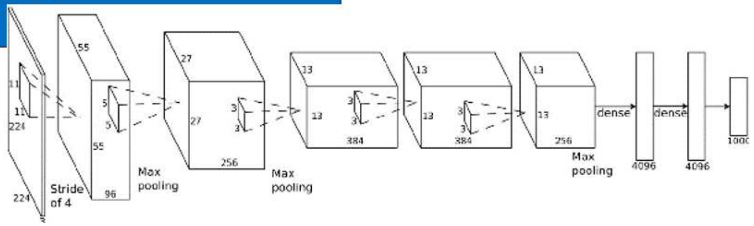

Hardware

Software

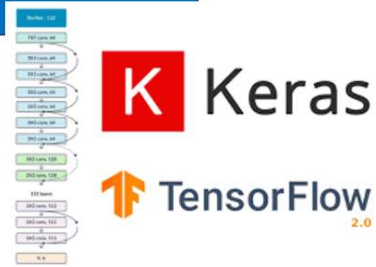

- **Software:** accessible Python frameworks such as TensorFlow and PyTorch.

Breakthrough Era

- **2012:** ImageNet breakthrough (AlexNet)
  Alex Krizhevsky, Ilya Sutskever, Geoffrey Hinton
- **2014-2015:** Very deep networks (ResNet)
  Kaiming He et al.
- **2010s:**
  - CNNs dominate vision
  - RNNs/LSTMs dominate speech and time series

---

https://people.idsia.ch/~juergen/deep-learning-history.html

Transformers and Foundation Models

- **2017:** Transformer architecture
  Ashish Vaswani et al.
- **2018-Present:** Large-scale pretraining and fine-tuning
  OpenAI, Google Brain, DeepMind
- **2020s:** Multimodal and foundation models
  Focus on scalability, efficiency, explainability, and ethics

https://people.idsia.ch/~juergen/deep-learning-history.html

Healthcare

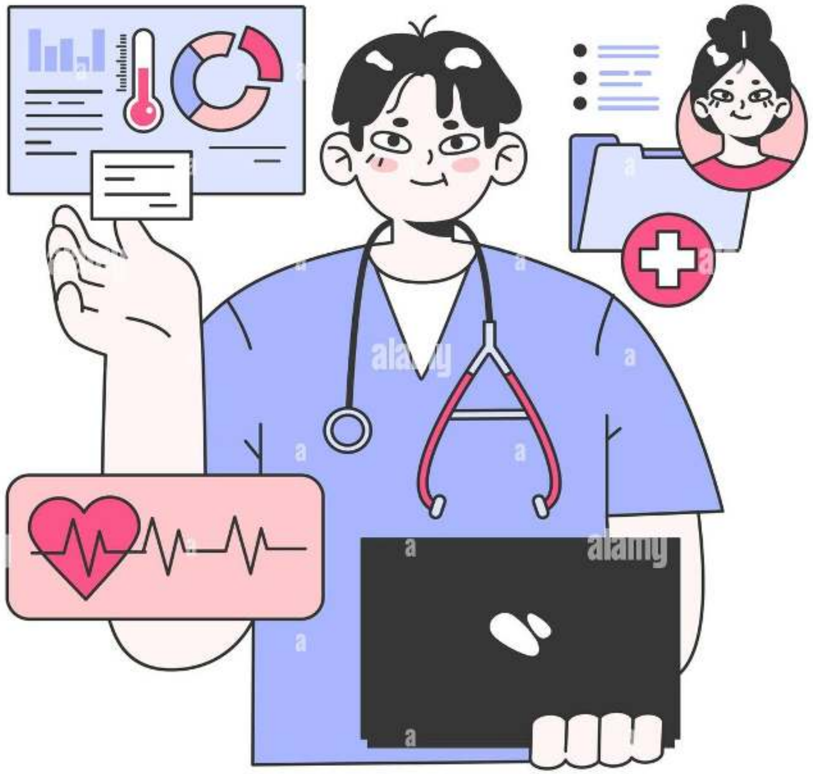

- Humans are the only species able to systematically improve health through technology.
- Medical progress has been driven by accumulated human knowledge.
- New technologies include diagnostics, drugs, procedures, and devices.

What is Healthcare Data?

Information generated across clinical and biomedical settings:

- Structured data: diagnoses, procedures, lab tests, medication details
- Unstructured data: clinical notes, medical images, physiological signals, patient-generated data from sensors and wearables
- Molecular and omics data: genomics, proteomics, metabolomics
- Biomedical research data: molecules, chemistry, drug screening results

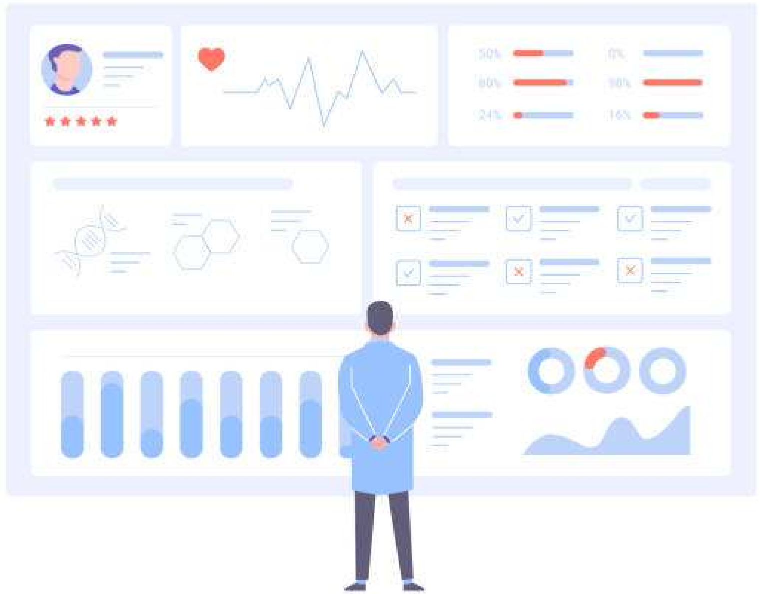

Healthcare Data

Health-related data generation is projected to surpass 100 zettabytes.

The potential for artificial intelligence applications is basically endless.

Rising Healthcare Costs

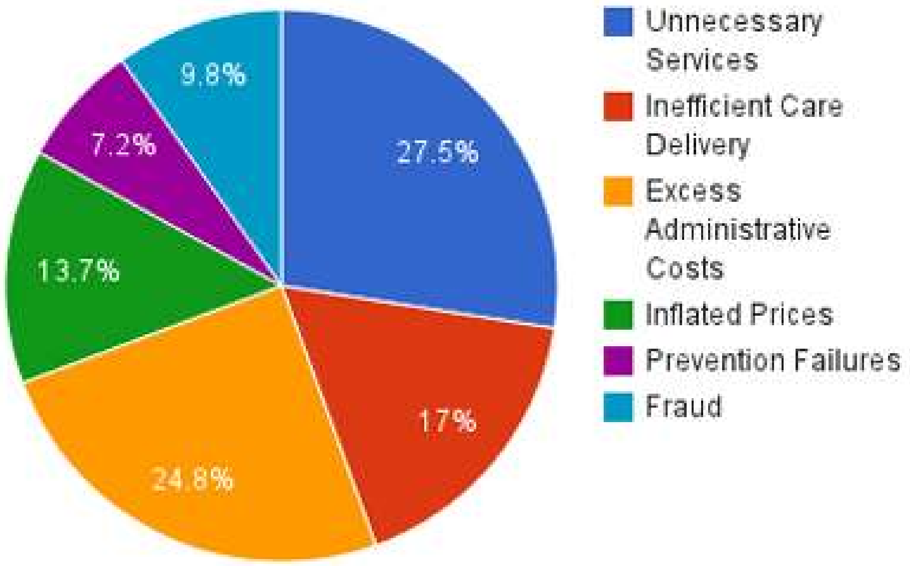

- Life expectancy increases → escalating healthcare spending.
- US healthcare cost in 2019: over $3.6 trillion (17.8% of GDP).
- Estimated annual waste: $760–935 billion.

Challenges and Opportunities

- Medical knowledge and health data are expanding rapidly in volume, diversity, and complexity, causing expertise to become outdated.
- Rising patient loads and administrative demands are increasing clinician workload and contributing to burnout.
- Data overload and complex workflows reduce meaningful doctor–patient interaction.
- System strain contributes to high rates of preventable medical errors.

AI can help in several healthcare domains, and for several applications!

Bioinformatics

Combines biology and computer science to study biological information

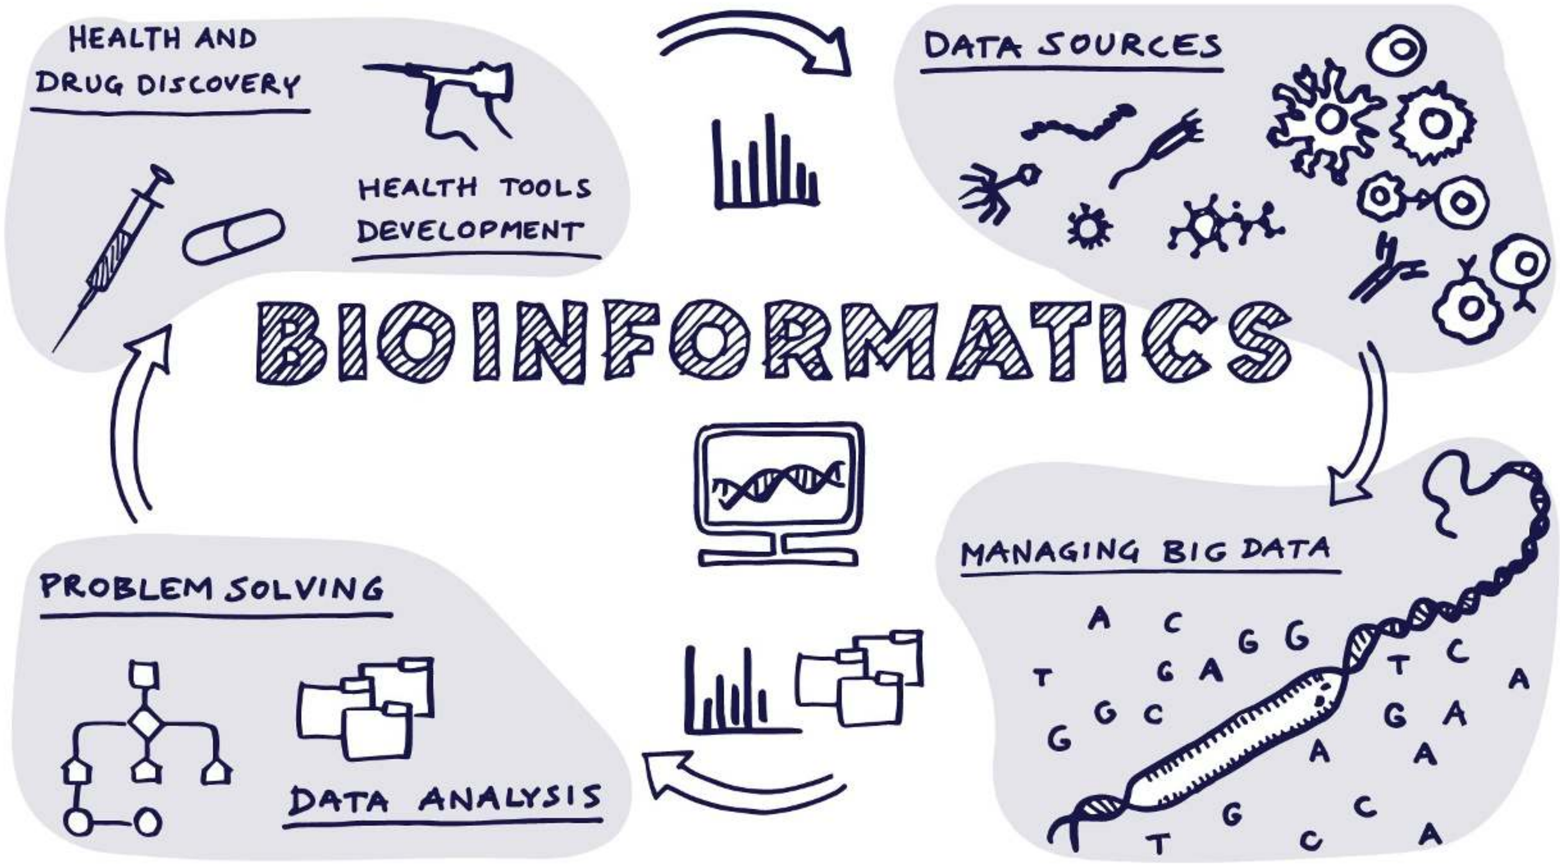

Bioinformatics - The Protein Folding Problem

DREAMING UP PROTEINS

Researchers used deep neural networks to invent, or 'hallucinate', sequences of amino acids that could fold into proteins; in some cases they have synthesized these proteins to compare their actual structures with predictions.

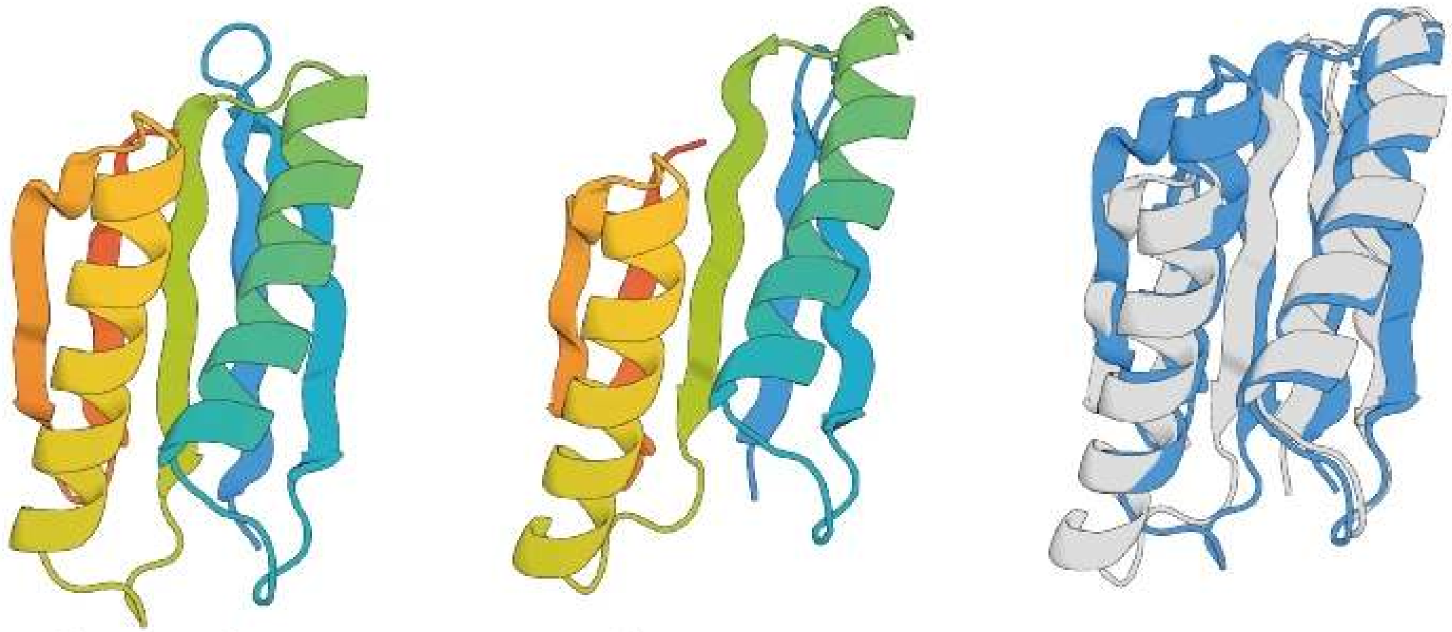

'Hallucinated' protein (software prediction)

Actual structure (experimentally determined)

Superposition of hallucinated (blue) and actual (grey) structures

Bioinformatics - The Protein Folding Problem

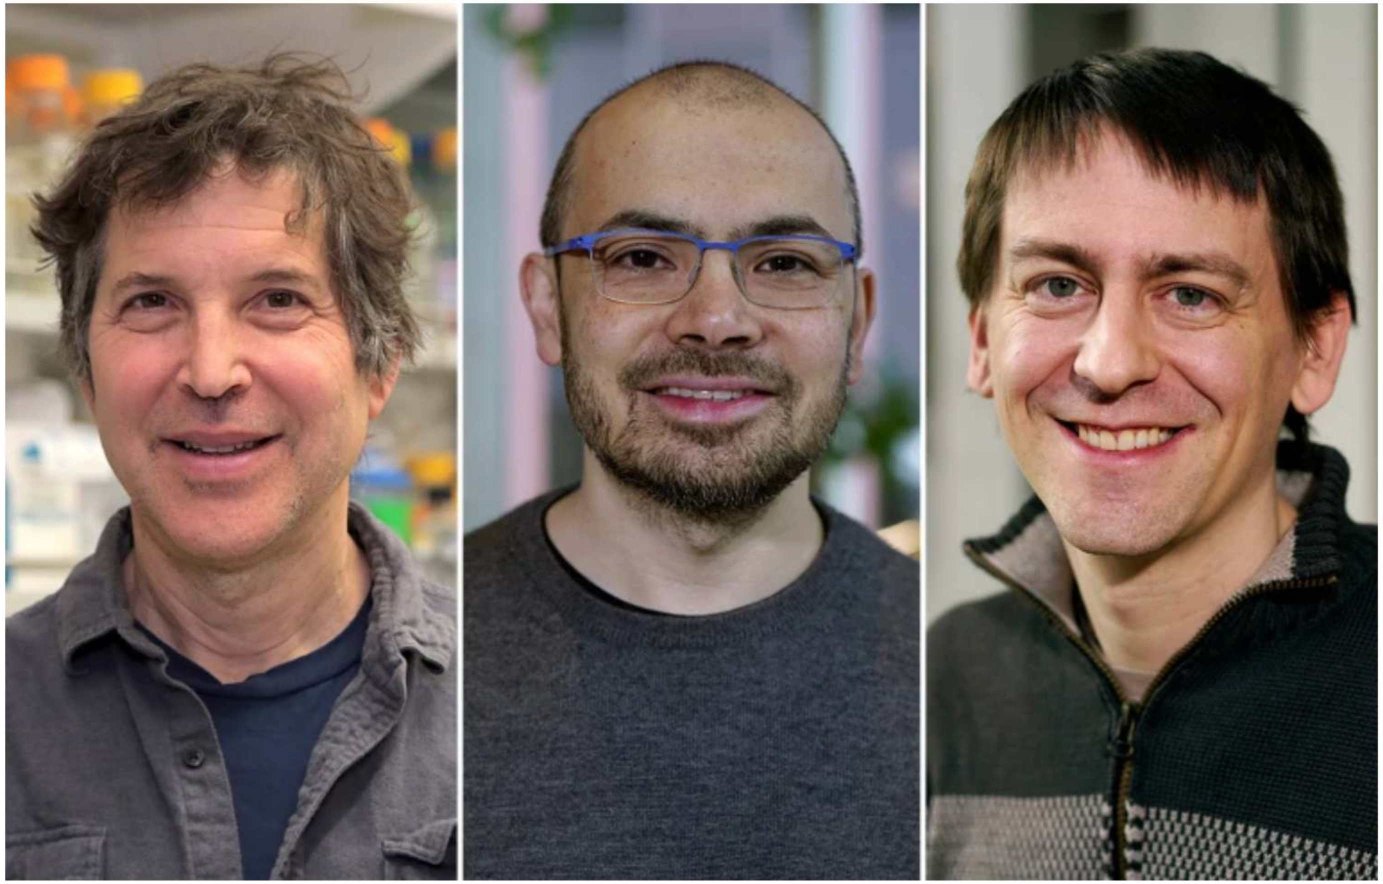

David Baker, Demis Hassabis and John Jumper (left to right) won the chemistry Nobel for developing computational tools that can predict and design protein structures. Credit: BBVA Foundation

Learning to predict functional properties from chemical structures

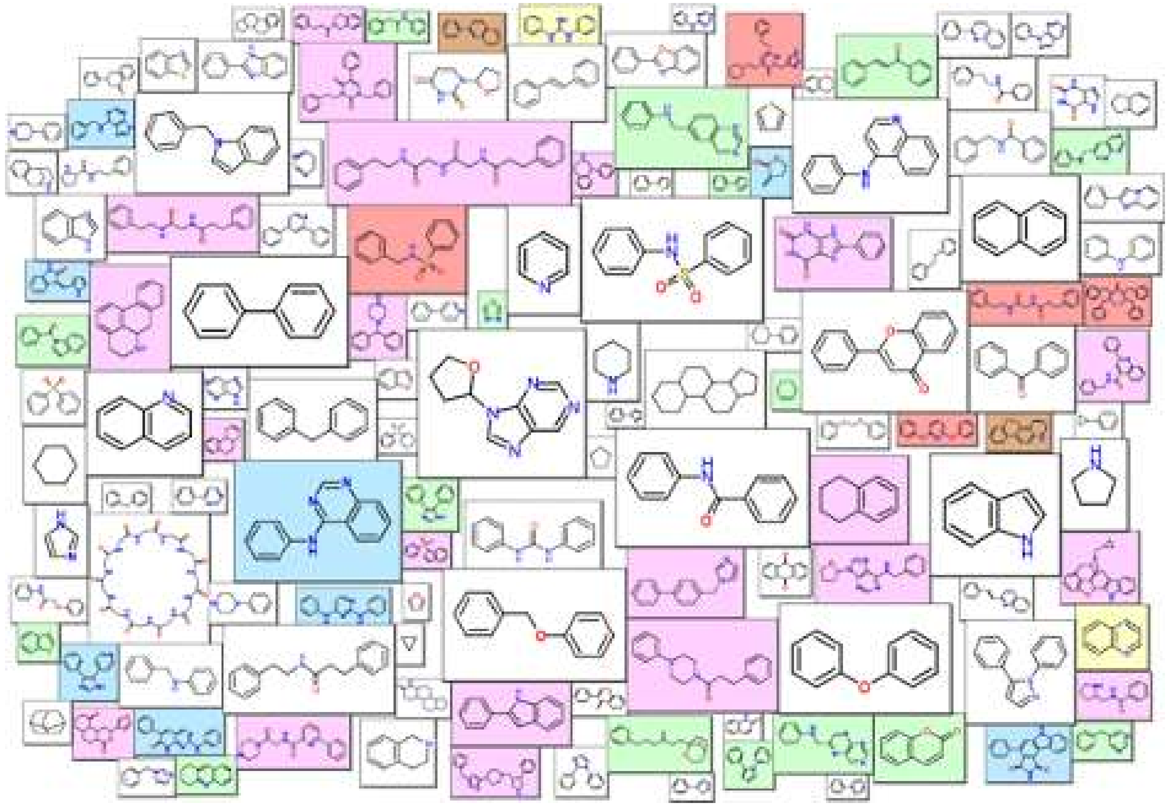

Drug discovery

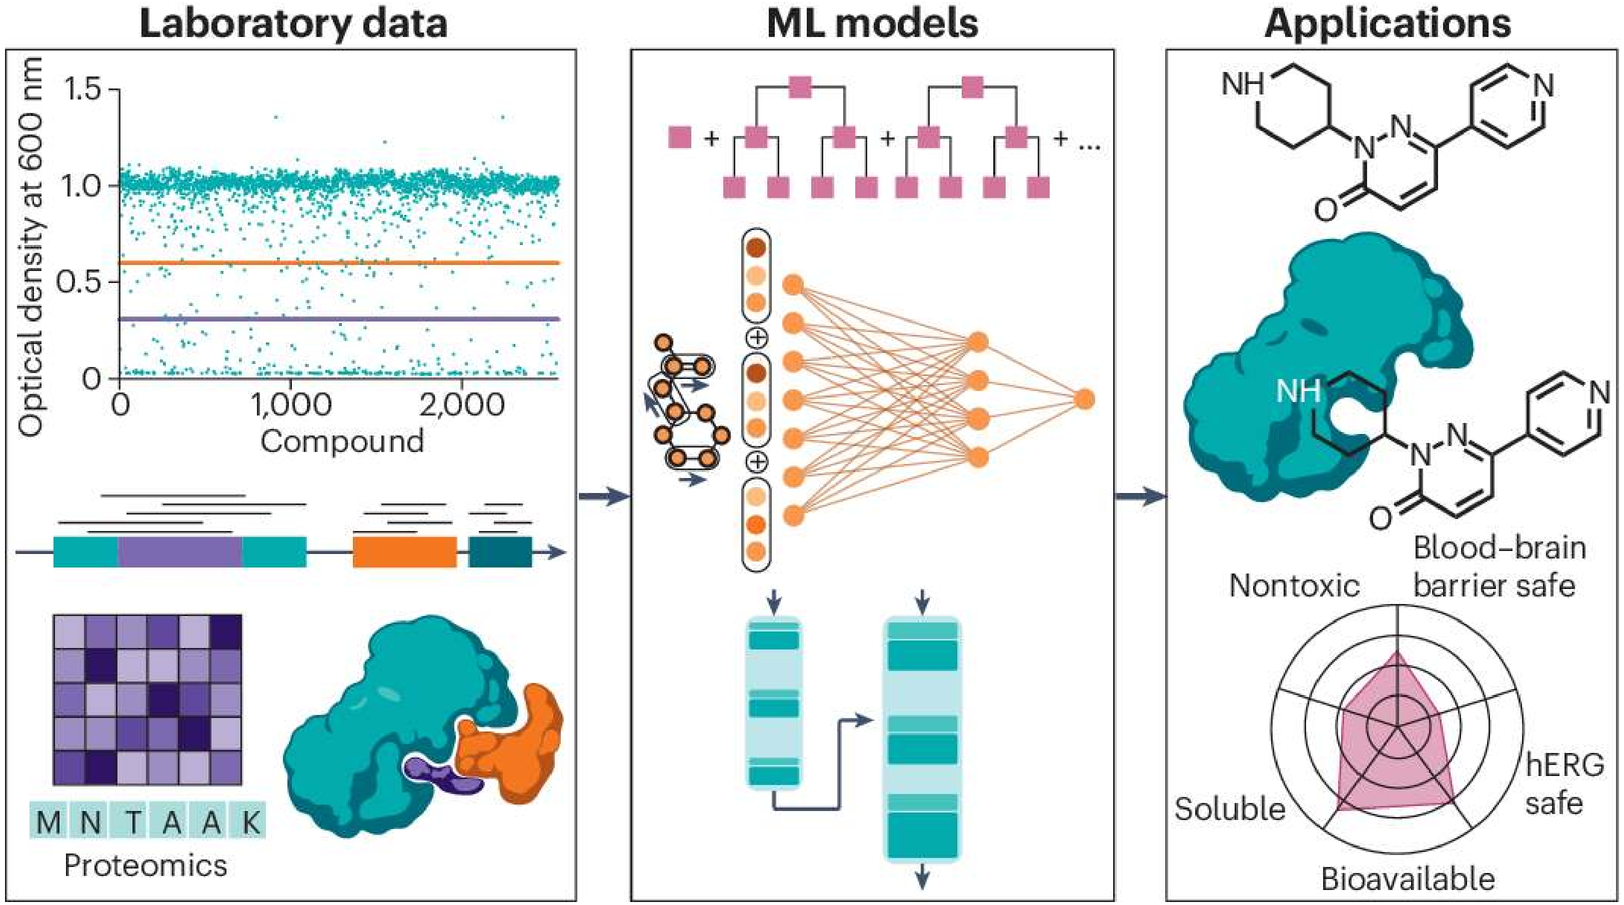

Machine Vision - Diabetic Retinopathy Detection

Process and model image data.

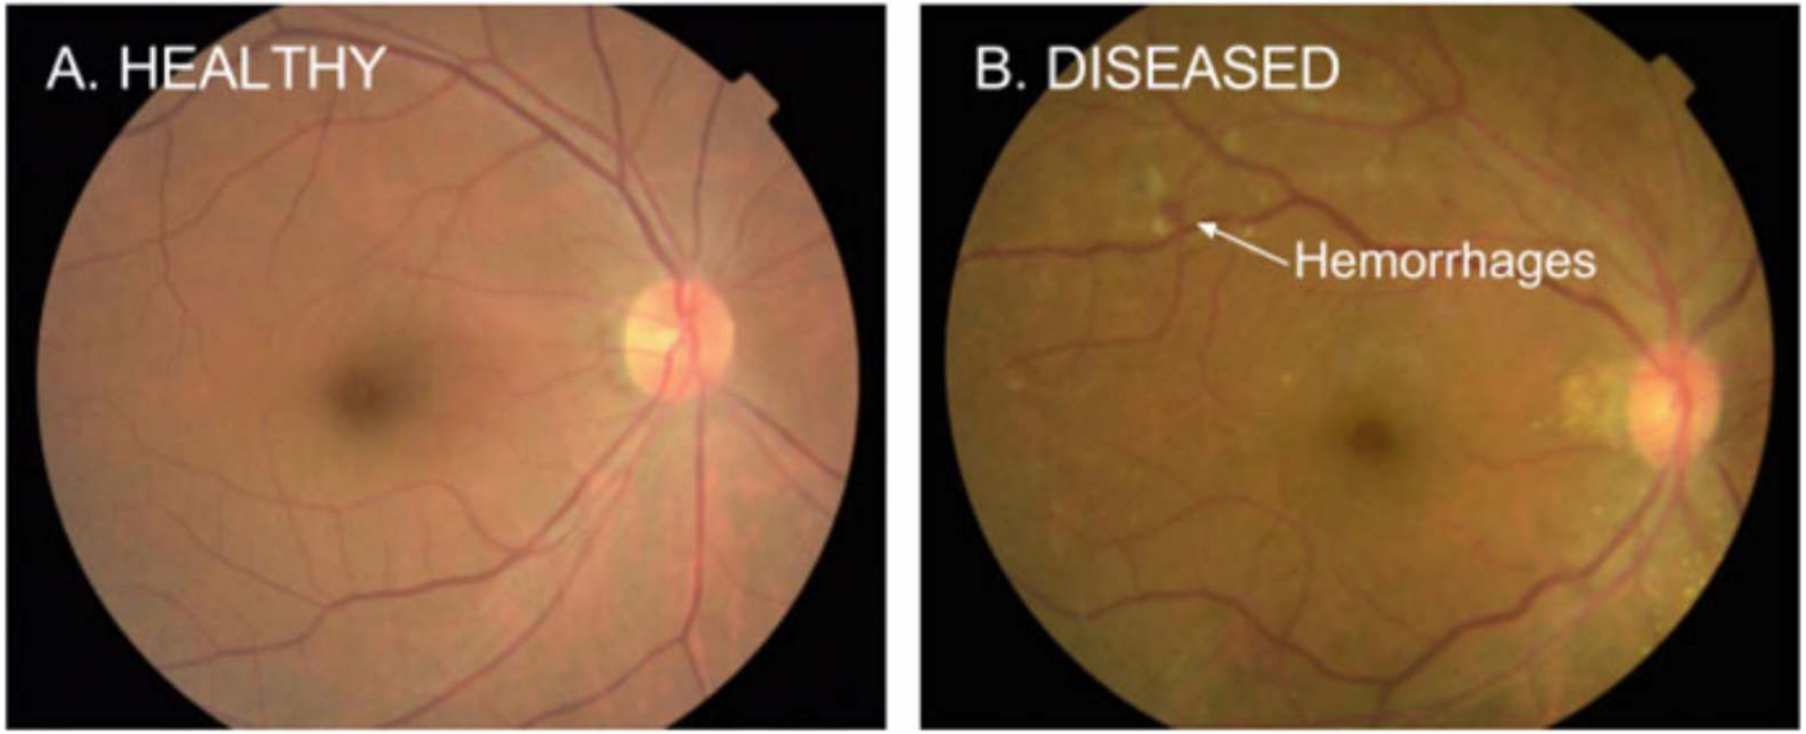

Machine Vision - Skin Lesions

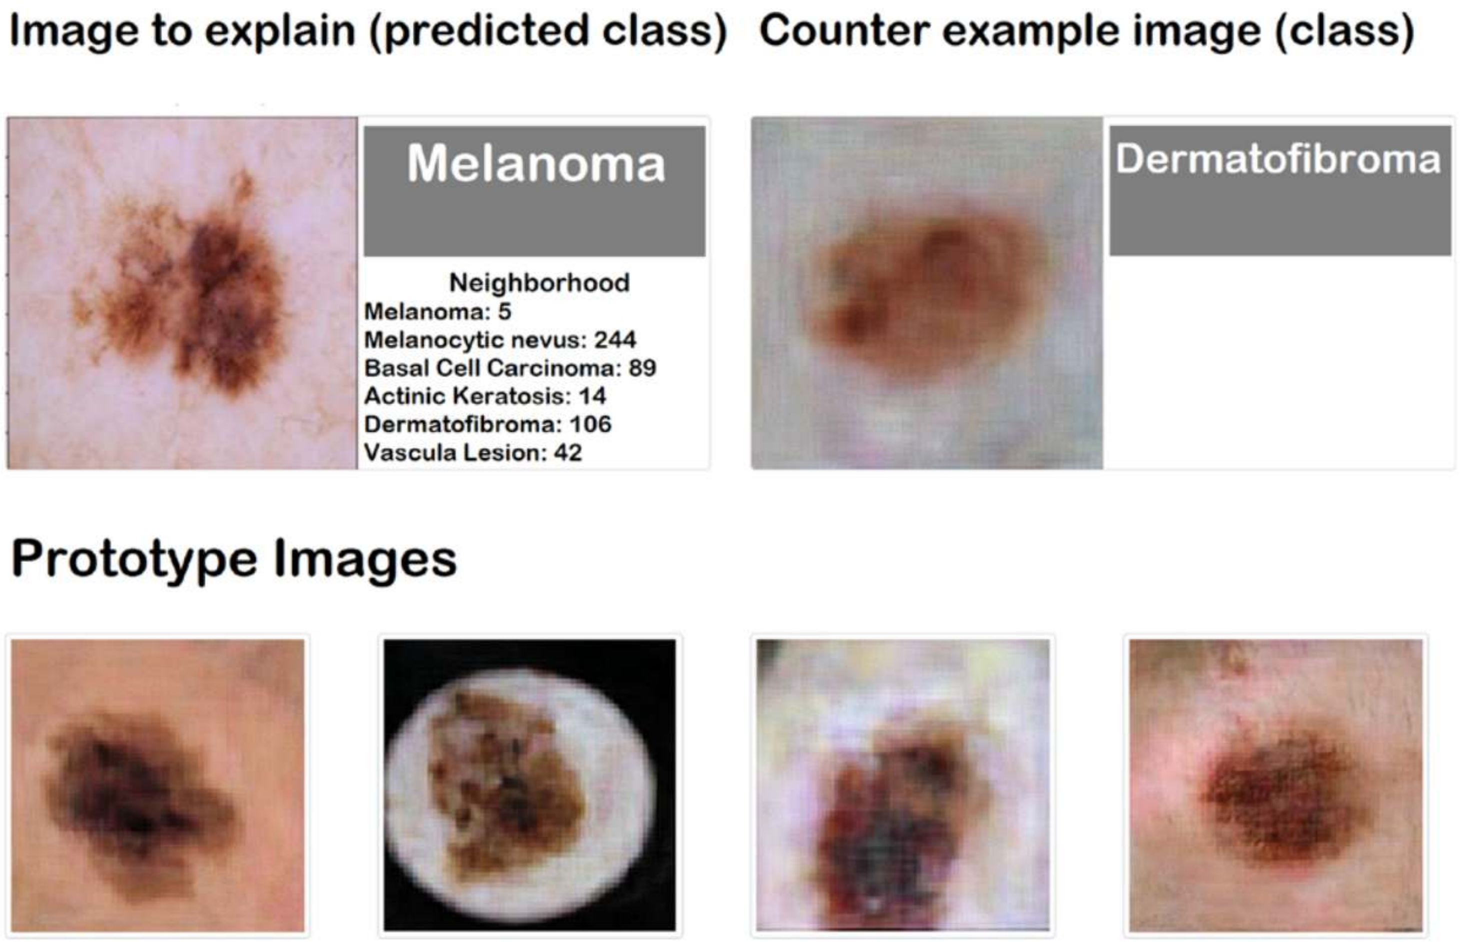

Metta, Carlo, et al. "Improving trust and confidence in medical skin lesion diagnosis through explainable deep learning." International Journal of Data Science and Analytics 20.1 (2025): 183-195.

Signal Processing - ECG

Analyze temporal/sequential data for predicting diseases, e.g., Early Detection of Heart Failure

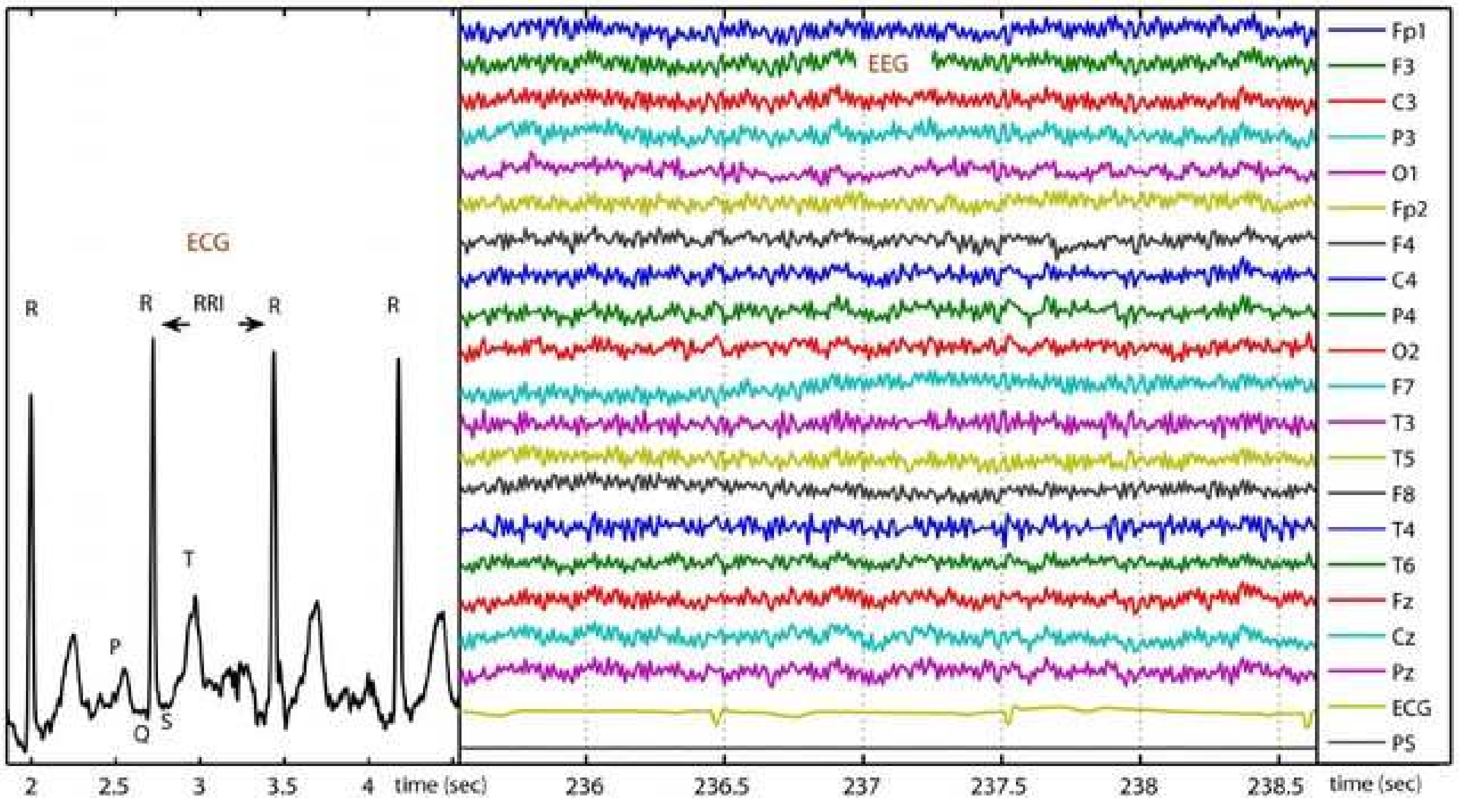

Scoring laboratory polysomnography (PSG) data is still a manual task...

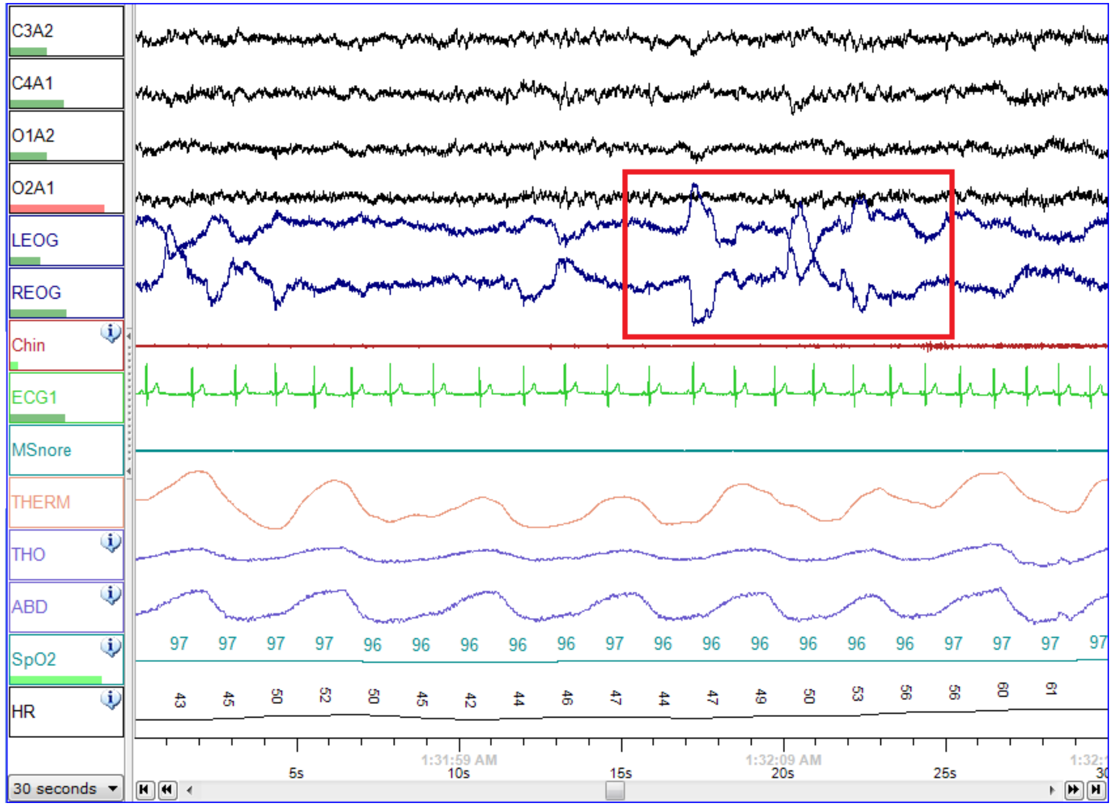

Patient outcome prediction

Readmission prediction, Length of stay prediction, Mortality prediction...

PatientID	Gender	Age	Zip code	Test
55998	M	19	15723	Negative
88557	F	35	15674	Positive
55868	F	35	15674	Positive
44551	M	45	15623	Negative
58524	M	45	15623	Negative
25584	F	61	15633	Negative
58744	F	61	15643	Positive
87524	M	19	15762	Positive
87384	M	19	15762	Negative
17583	F	19	15762	Positive

M: male; F: female

Predicting disease spreading, hospitalization, death.

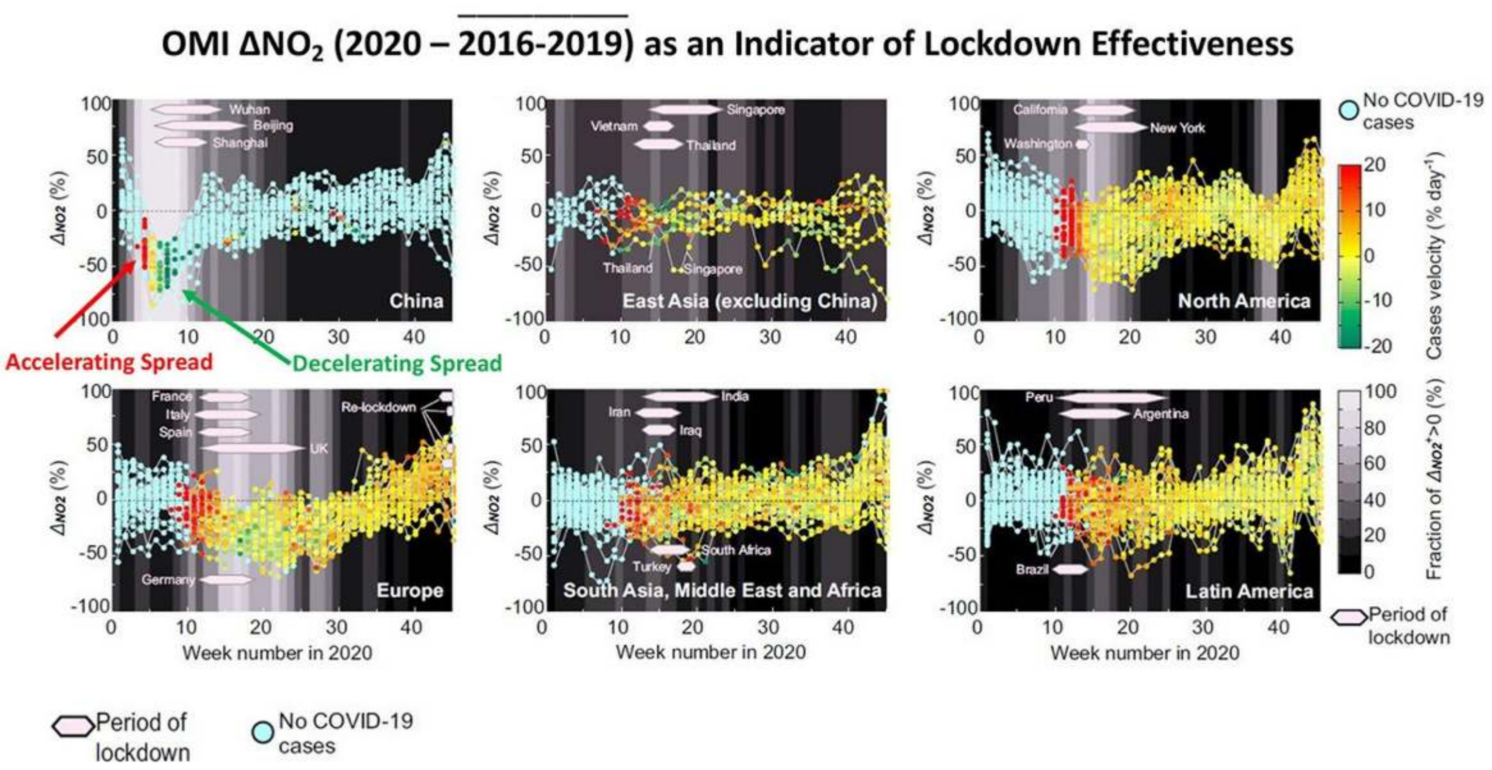

Fraud detection, Cost estimation

TRADITIONAL RULE-BASED APPROACH

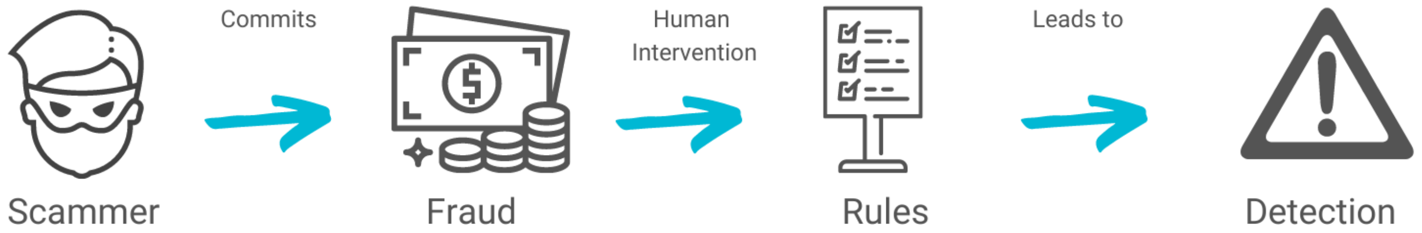

MACHINE LEARNING APPROACH

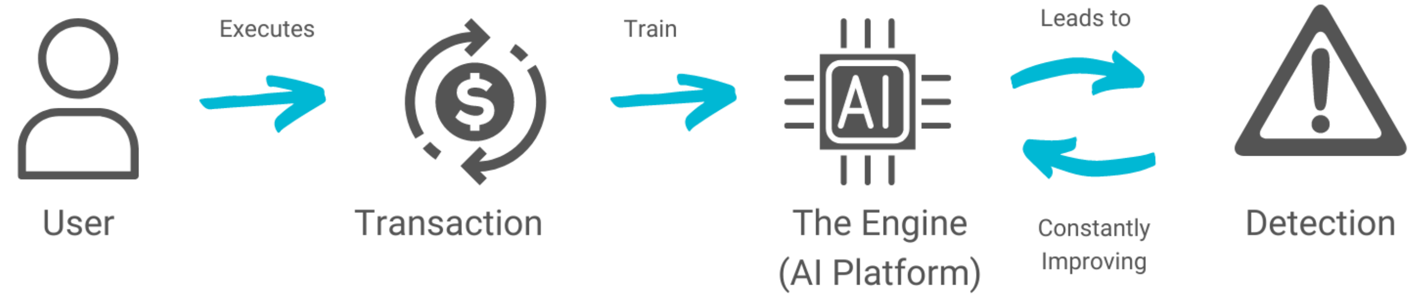

Natural Language Processing

Convert natural language into machine-processable data.

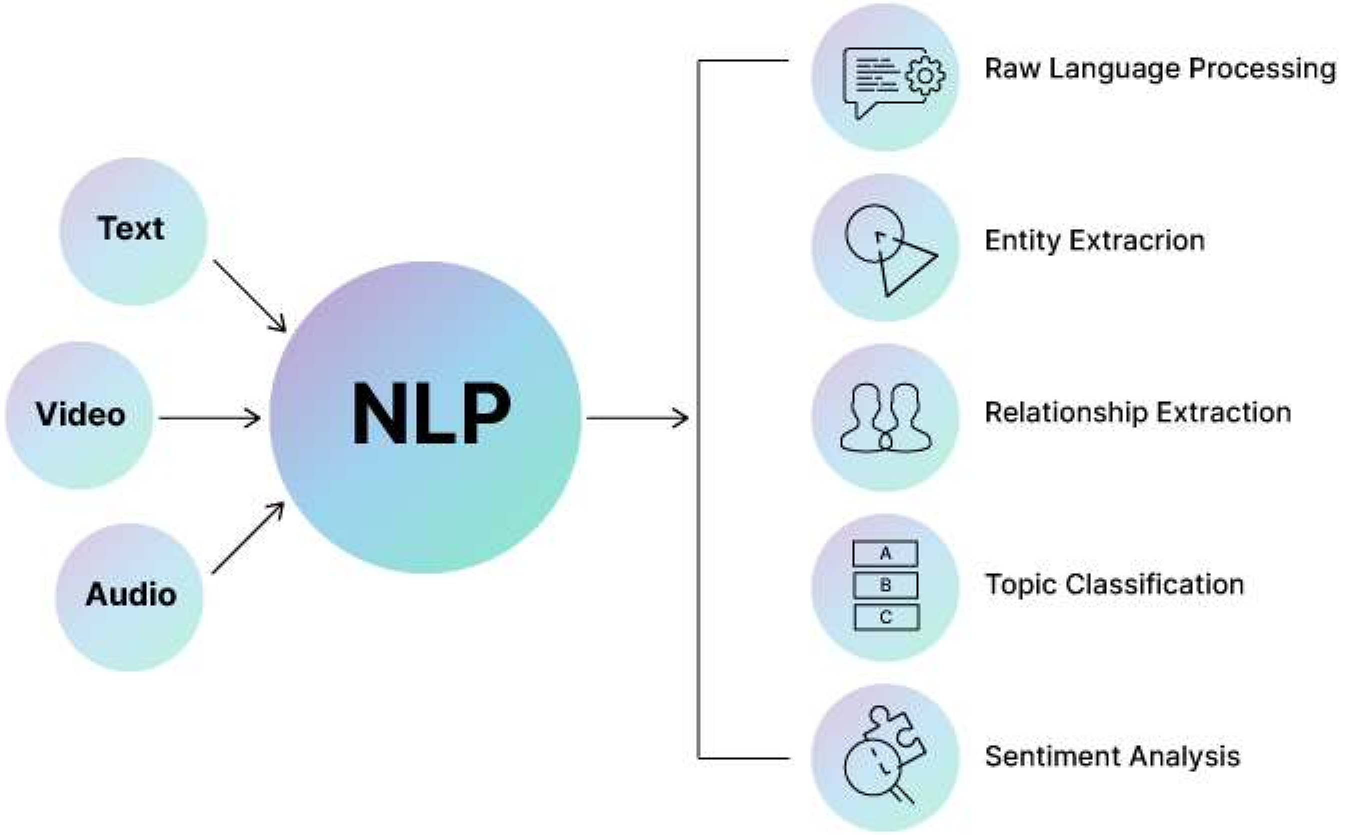

Goals

By the end of this course, you will be able to:

- understand healthcare data;
- choose the right tool for your own tasks;
- apply deep learning models
- dialogue with machine learning experts