Intelligenza Artificiale: Istruzioni per l'uso

Dr.Ir. Jacopo De Stefani

(with the collaboration of Charlotte Nachtegael, MSc)

Lecturer @ TPM-ESS-ICT

Cremona, 13/12/2022

About Me

Academic Background

- BSc in Computer Engineering, Politecnico di Milano, Italy (2011)
- MSc in Computer Science and Engineering, ULB, Belgium (2013)
- MSc in Computer Engineering, Politecnico di Milano, Italy (2015)
- PhD in Machine Learning and Time Series Analysis, ULB, Belgium (2022)

Scientific activity

- 4 international peer-reviewed journal publications
- 6 international peer-reviewed conference proceedings
- **1 international patent**
- Reviewer for International Journal of Forecasting, IEEE Access, Technology and Economics of Smart Grid and Sustainable Energy

TU Delft

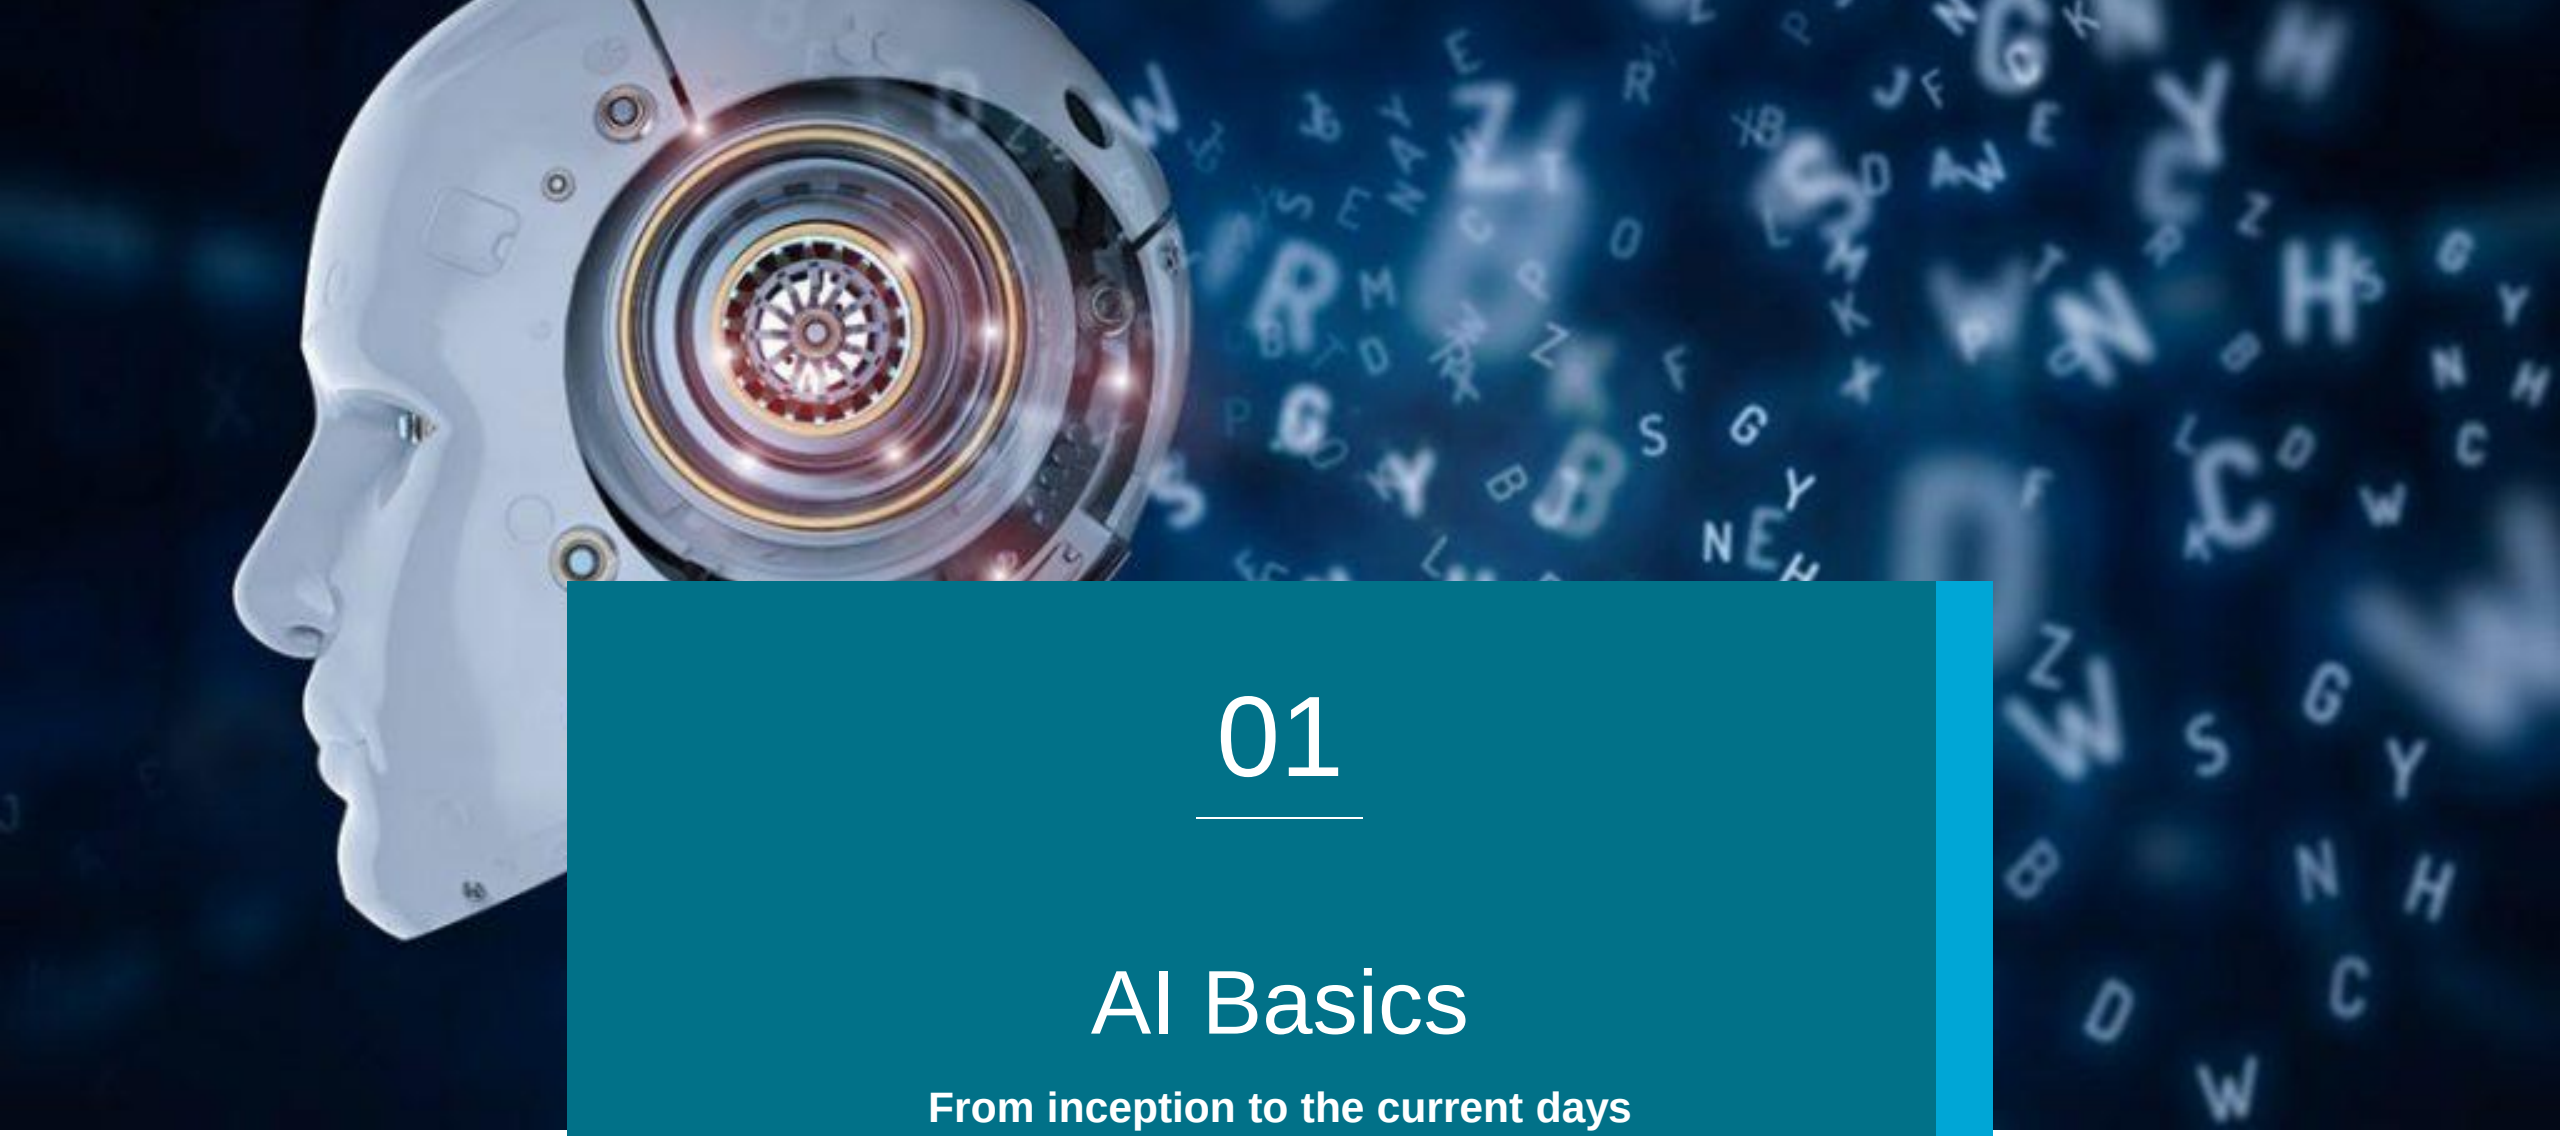

01

AI Basics

From inception to the current days

TU Delft

How would you define Artificial Intelligence?

TU Delft

According to scientific literature

- No consensus on a single definition
- **Thinking Humanly:** Cognitive science/Cognitive modelling
- **Acting Humanly:** Turing test
- **Thinking Rationally:** Logic-based/Deductive Intelligence
- **Acting Rationally:** Rational (trying to achieve the best solution) agents
- Is it more about actual intelligence or perceived intelligence?

TU Delft

Thinking Humanly	**Thinking Rationally**
"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense." (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Acting Humanly	**Acting Rationally**
"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)	"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al.*, 1998) "AI . . . is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

Figure 1.1 Some definitions of artificial intelligence, organized into four categories.

Russell, S. J. (2010). *Artificial intelligence a modern approach*. Pearson Education, Inc..

Chinese Room Argument

- Is it more about actual intelligence or perceived intelligence?
- Does an AI actually understand or does it simply execute an algorithm/set of rules with (super)human capacities?

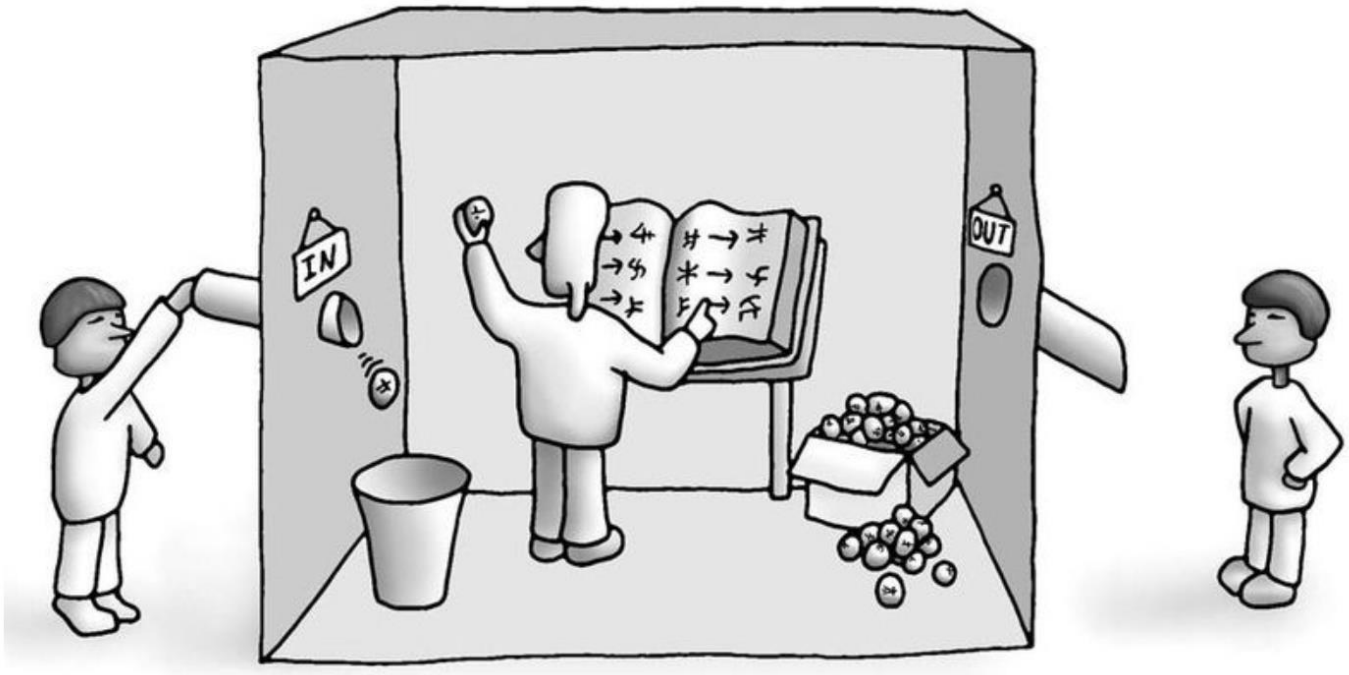

TU Delft

According to the scientific literature

- Multi-faceted field
- Connected and inspired by many different subfields
- In this presentation, focus more on the mathematical and computer science perspective

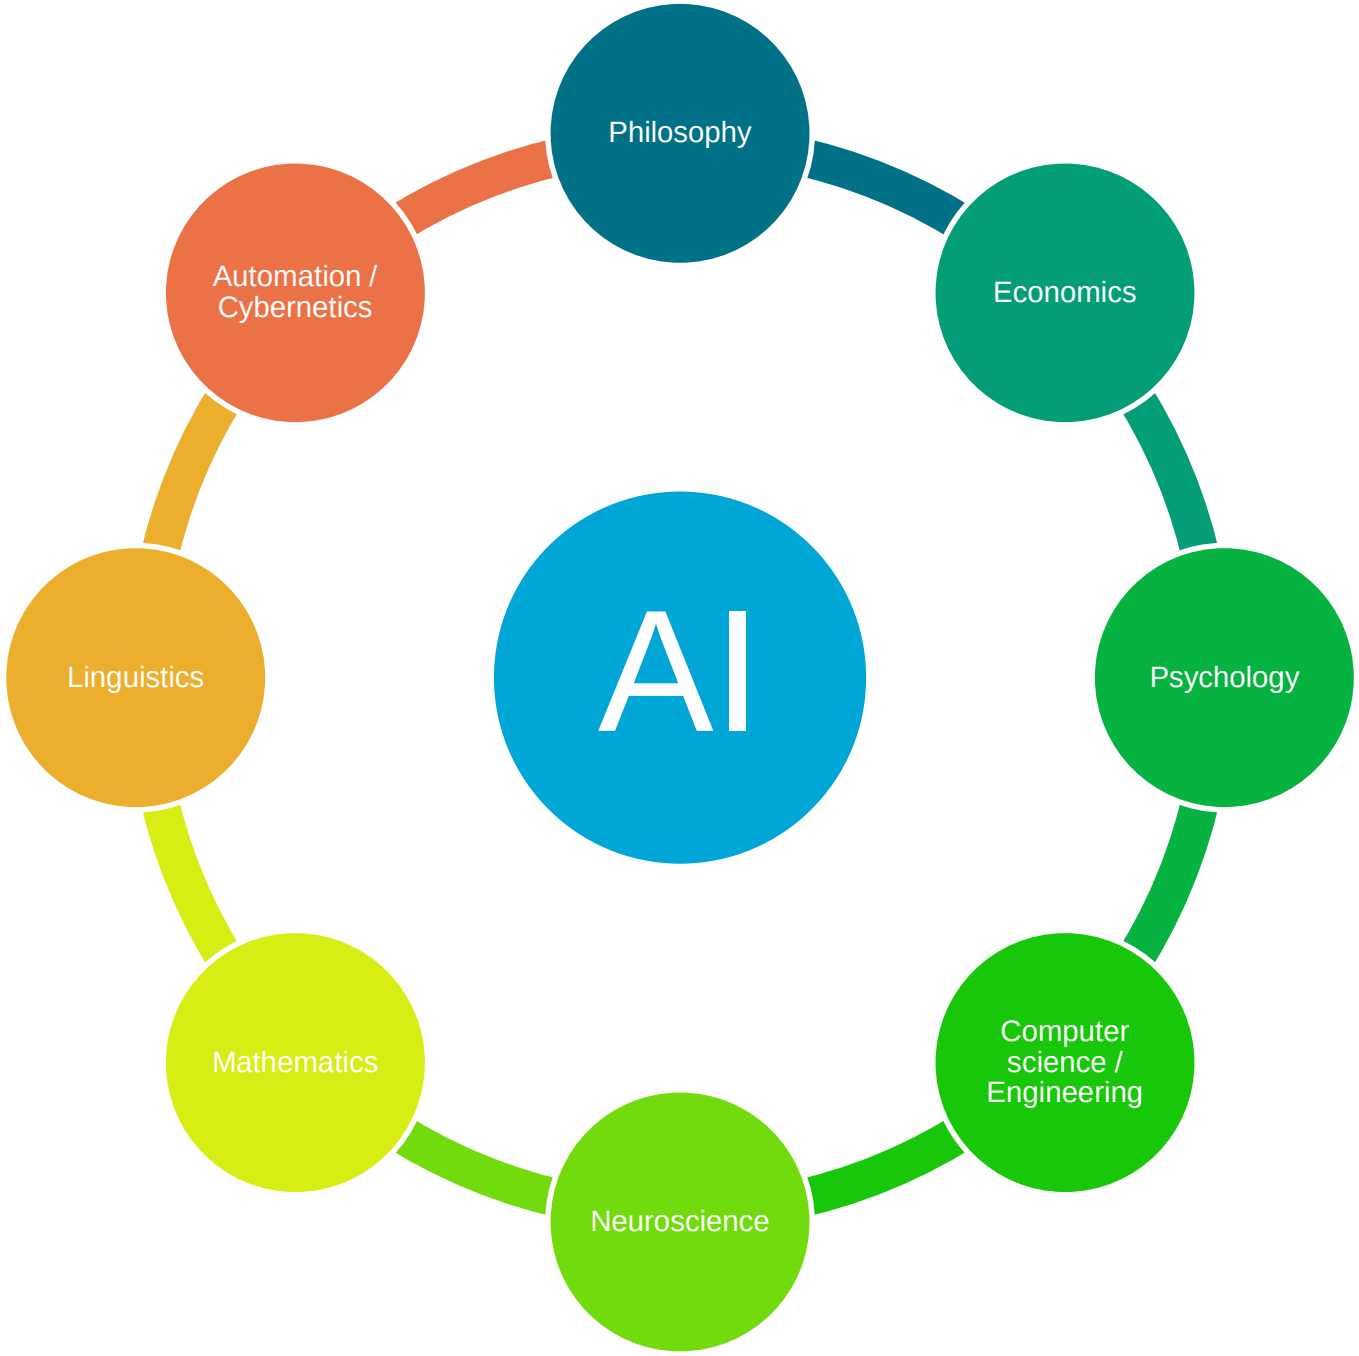

TU Delft

Narrow (weak) vs General (strong) AI

Narrow AI

Dedicated to assist with or take over specific tasks.

General AI

Takes knowledge from one domain, transfers to other domain.

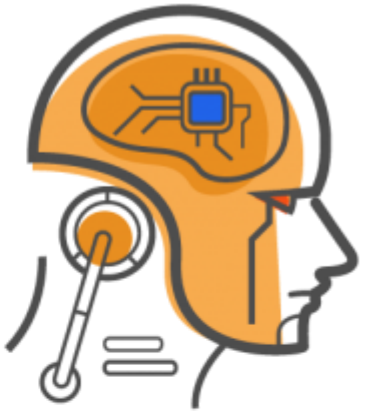

Super AI

Machines that are an order of magnitude smarter than humans.

TU Delft

Source: https://medium.com/@selecticalinfotech/understanding-artificial-intelligence-3bdcde2993c3

When did the development of Artificial Intelligence start?

TU Delft

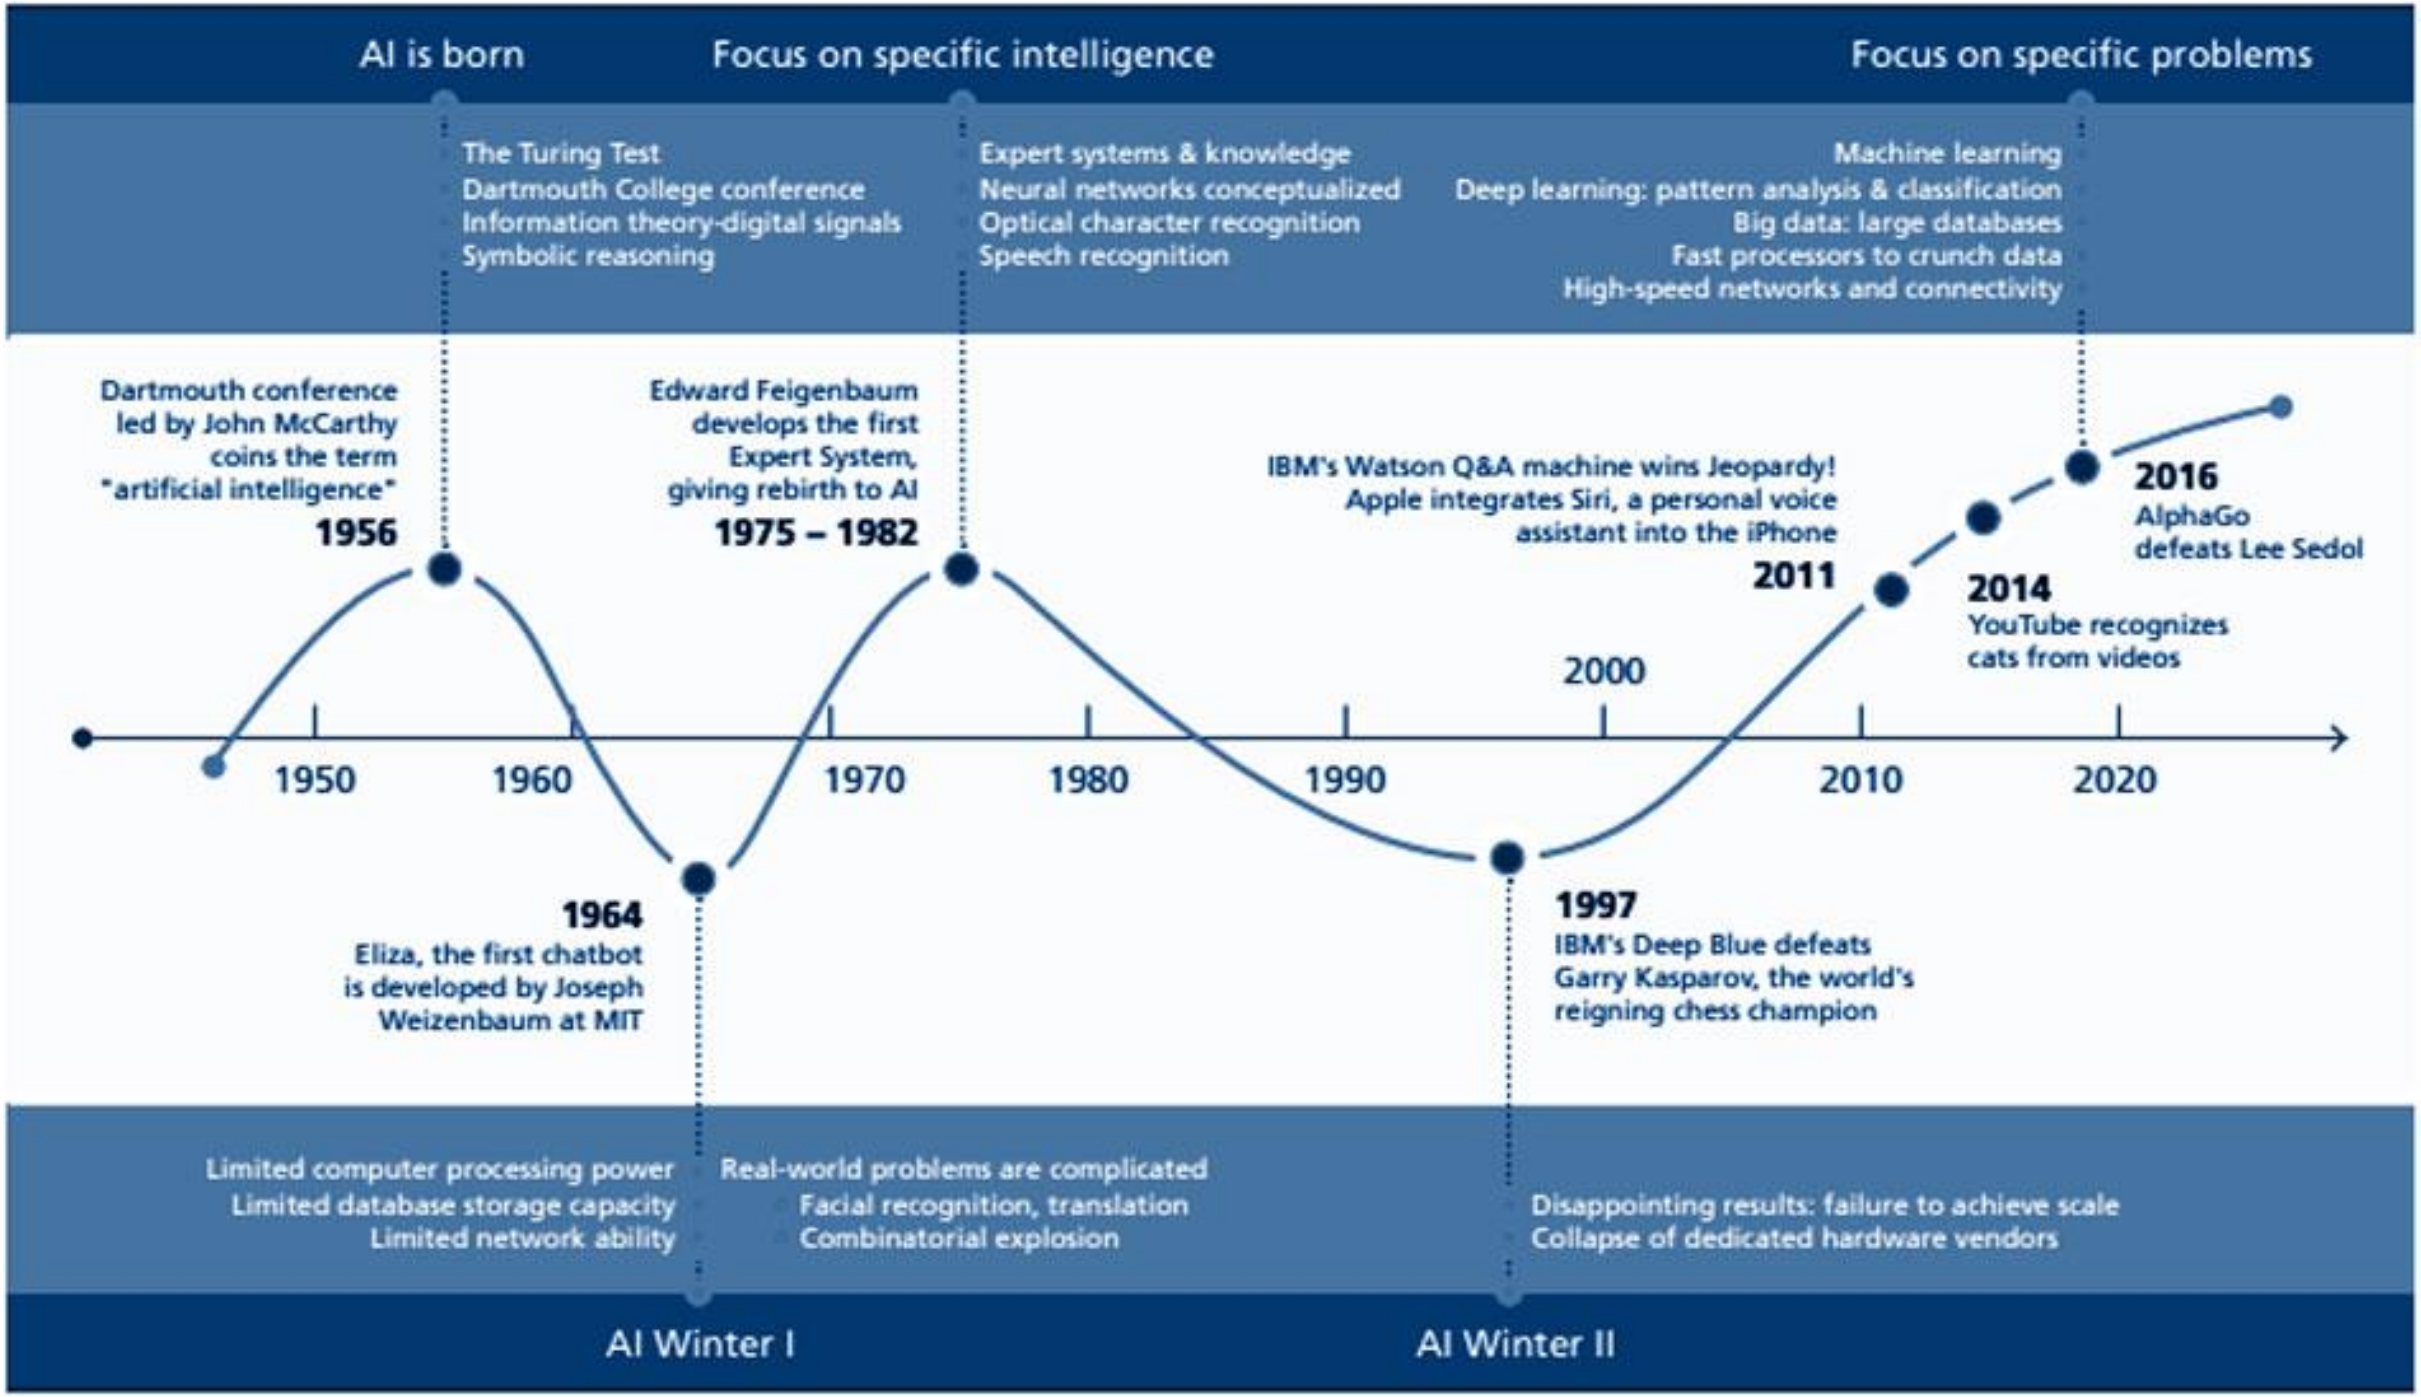
AI is born
Focus on specific intelligence
Focus on specific problems
The Turing Test
Dartmouth College conference
Information theory-digital signals
Symbolic reasoning
Expert systems & knowledge
Neural networks conceptualized
Optical character recognition
Speech recognition
Machine learning
Deep learning: pattern analysis & classification
Big data: large databases
Fast processors to crunch data
High-speed networks and connectivity
Dartmouth conference led by John McCarthy coins the term "artificial intelligence"
1956
Edward Feigenbaum develops the first Expert System, giving rebirth to AI
1975 – 1982
IBM's Watson Q&A machine wins Jeopardy!
Apple integrates Siri, a personal voice assistant into the iPhone
2011
2014
YouTube recognizes cats from videos
2016
AlphaGo defeats Lee Sedol
1950
1960
1970
1980
1990
2000
2010
2020
1964
Eliza, the first chatbot is developed by Joseph Weizenbaum at MIT
1997
IBM's Deep Blue defeats Garry Kasparov, the world's reigning chess champion
Limited computer processing power
Limited database storage capacity
Limited network ability
Real-world problems are complicated
Facial recognition, translation
Combinatorial explosion
Disappointing results: failure to achieve scale
Collapse of dedicated hardware vendors
AI Winter I
AI Winter II

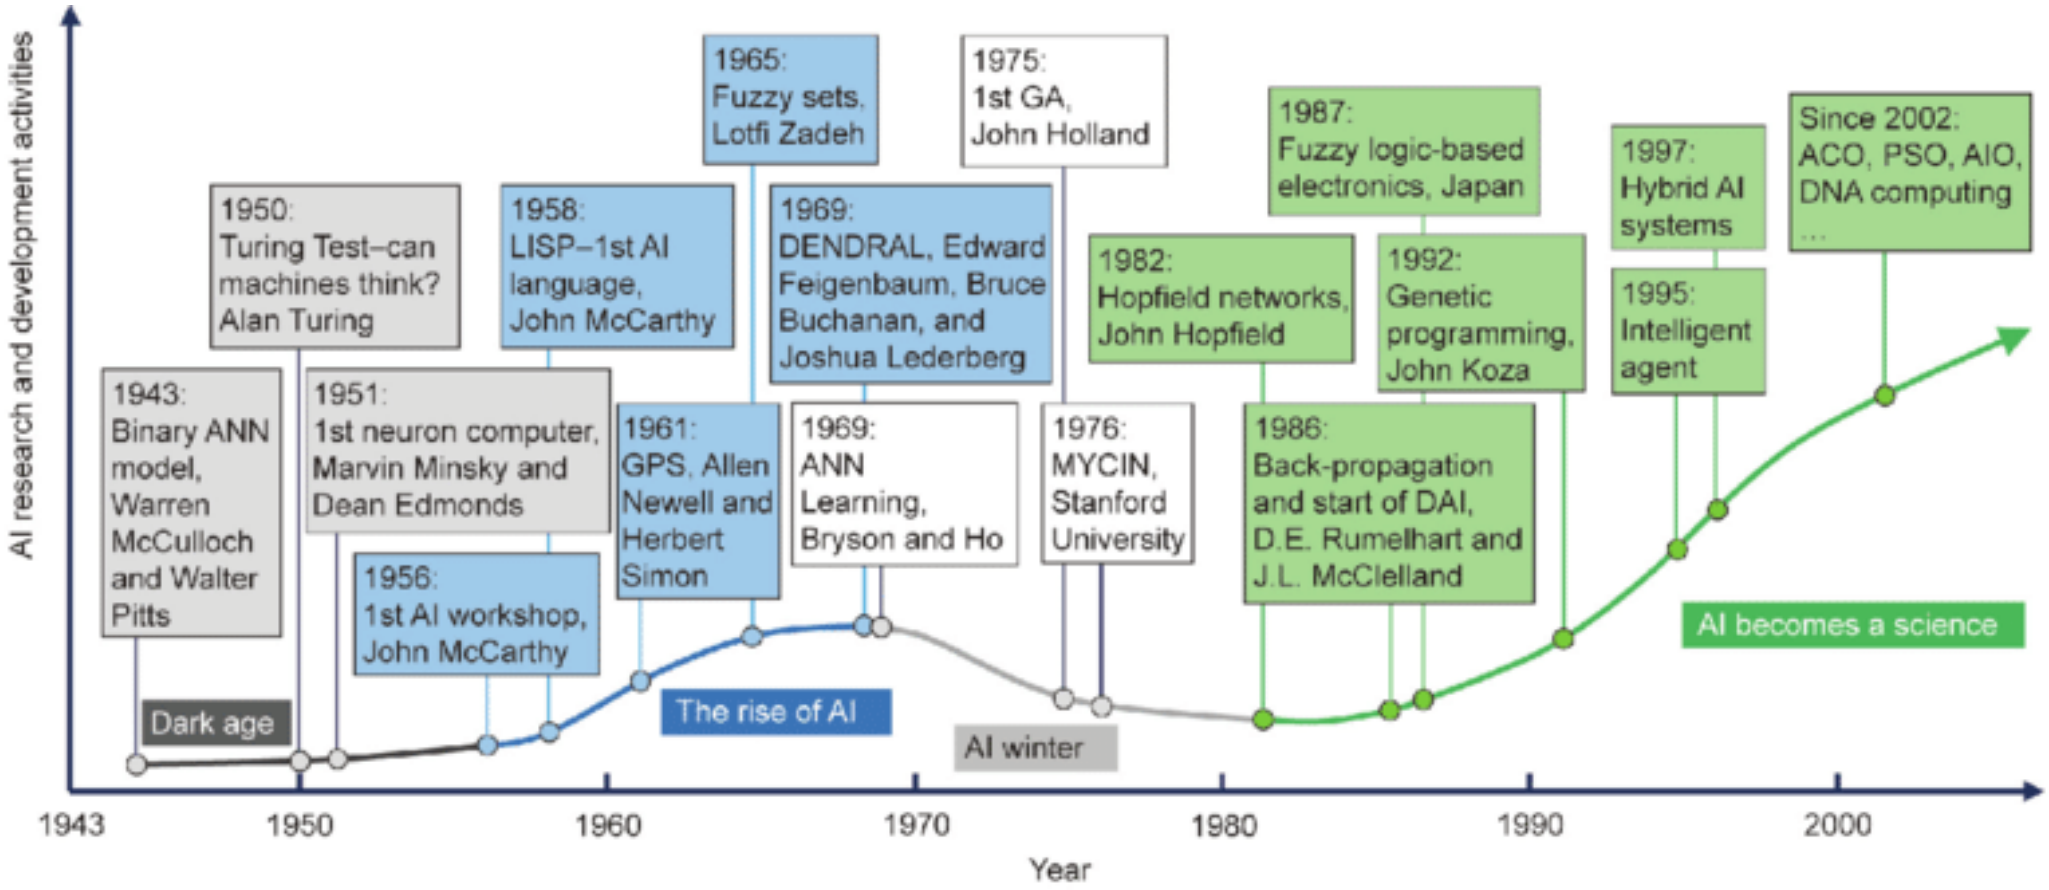
AI research and development activities
1943:
Binary ANN model, Warren McCulloch and Walter Pitts
1950:
Turing Test–can machines think? Alan Turing
1951:
1st neuron computer, Marvin Minsky and Dean Edmonds
1956:
1st AI workshop, John McCarthy
1958:
LISP–1st AI language, John McCarthy
1961:
GPS, Allen Newell and Herbert Simon
1965:
Fuzzy sets, Lotfi Zadeh
1969:
DENDRAL, Edward Feigenbaum, Bruce Buchanan, and Joshua Lederberg
1969:
ANN Learning, Bryson and Ho
1975:
1st GA, John Holland
1976:
MYCIN, Stanford University
1982:
Hopfield networks, John Hopfield
1986:
Back-propagation and start of DAI, D.E. Rumelhart and J.L. McClelland
1987:
Fuzzy logic-based electronics, Japan
1992:
Genetic programming, John Koza
1995:
Intelligent agent
1997:
Hybrid AI systems
Since 2002:
ACO, PSO, AIO, DNA computing
...
Dark age
The rise of AI
AI winter
AI becomes a science
1943
1950
1960
1970
1980
1990
2000
Year

TU Delft

02

Machine Learning

From AI to Learning Machines

TU Delft

AI, Machine Learning, Deep Learning

AI

Computing systems which are capable of performing tasks that humans are very good at, for example recognising objects.

Machine Learning

The field of AI that applies statistical methods to enable computer systems to learn from the data towards an end goal.

Deep Learning

Neural Networks with several hidden layers.

Neural Networks

Are biologically inspired networks that extract abstract features from the data in a hierarchical fashion.

Narrow AI

Designed to perform a single task. What we have today.

Artificial General Intelligence (AGI)

Can accomplish any intellectual task that a human can do. Remains an aspiration.

Artificial Super Intelligence (ASI)

A form of intelligence that exceeds the performance of humans in all domains.

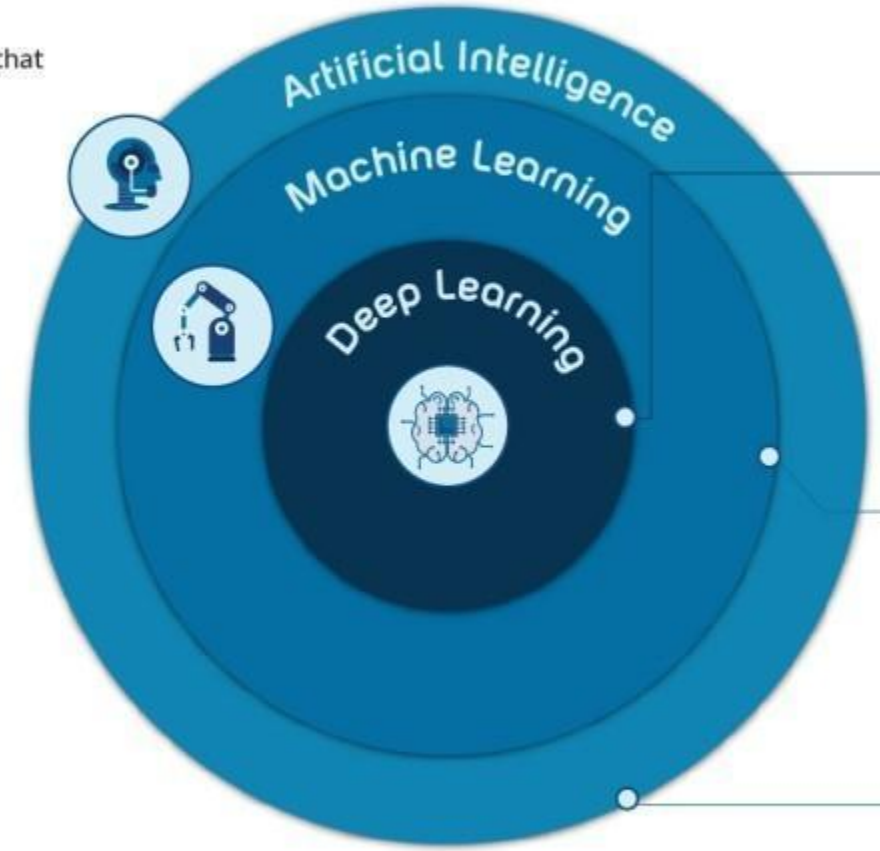

Non-exhaustive list

Deep Learning

Recurrent Neural Networks, Convolutional Neural Networks, Deep Reinforcement Learning with Deep-Q Learning, Capsules & GANs.

Machine Learning

Support Vector Machine, Decision Trees, Gradient Boosting, Principal Component Analysis, Logistic Regression, Linear Regression, K-means Clustering.

Classical Artificial Intelligence

Rule Based Systems, Search Algorithms Depth First, Breadth First, A* algorithm, Propositional Calculus, Predicate Calculus Logic.

www.dls.ltd

Imtiaz Adam
@deeplearn007

Machine Learning Definition

A computer program is said to **learn** from experience *E* with respect to some class of tasks *T* and performance measure *P*, if its performance at tasks in *T*, as measured by *P*, improves with experience *E*.

Mitchell, T. M., & Mitchell, T. M. (1997). *Machine learning* (Vol. 1, No. 9). New York: McGraw-hill.

Machine Learning Example

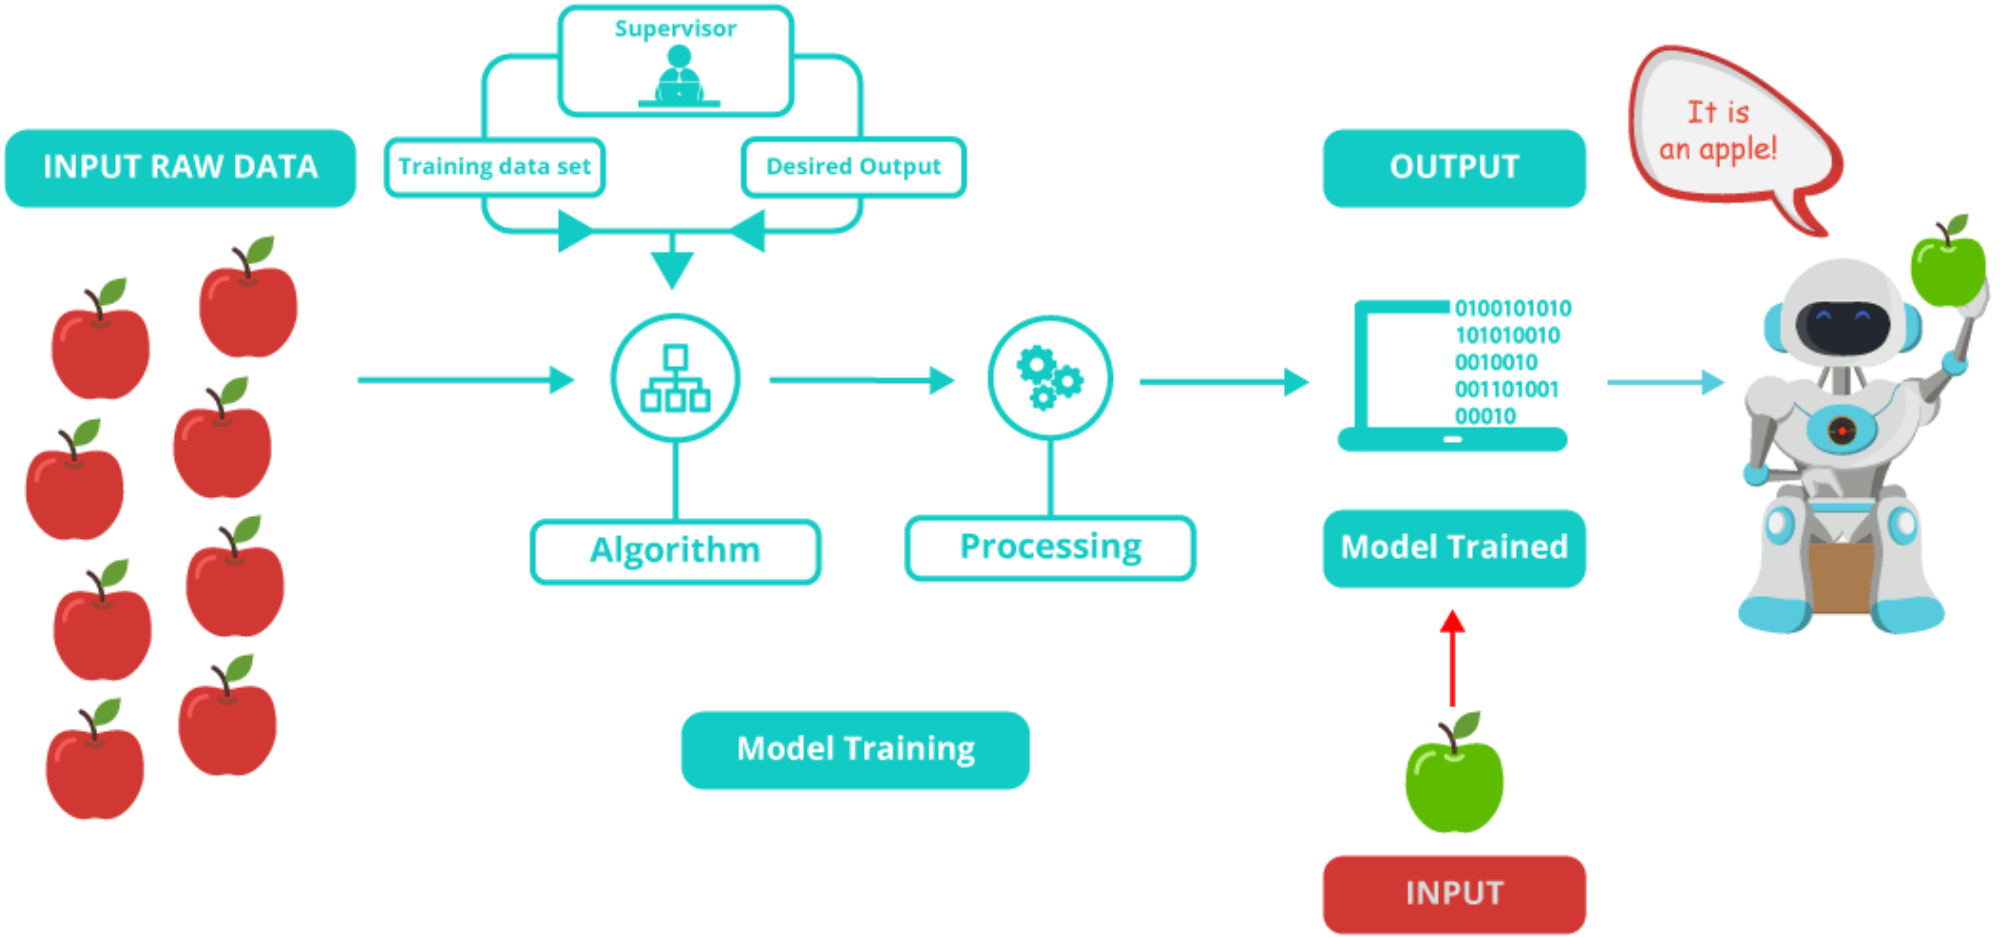

Mitchell, T. M., & Mitchell, T. M. (1997). *Machine learning* (Vol. 1, No. 9). New York: McGraw-hill.

Machine Learning as Engineering Problem

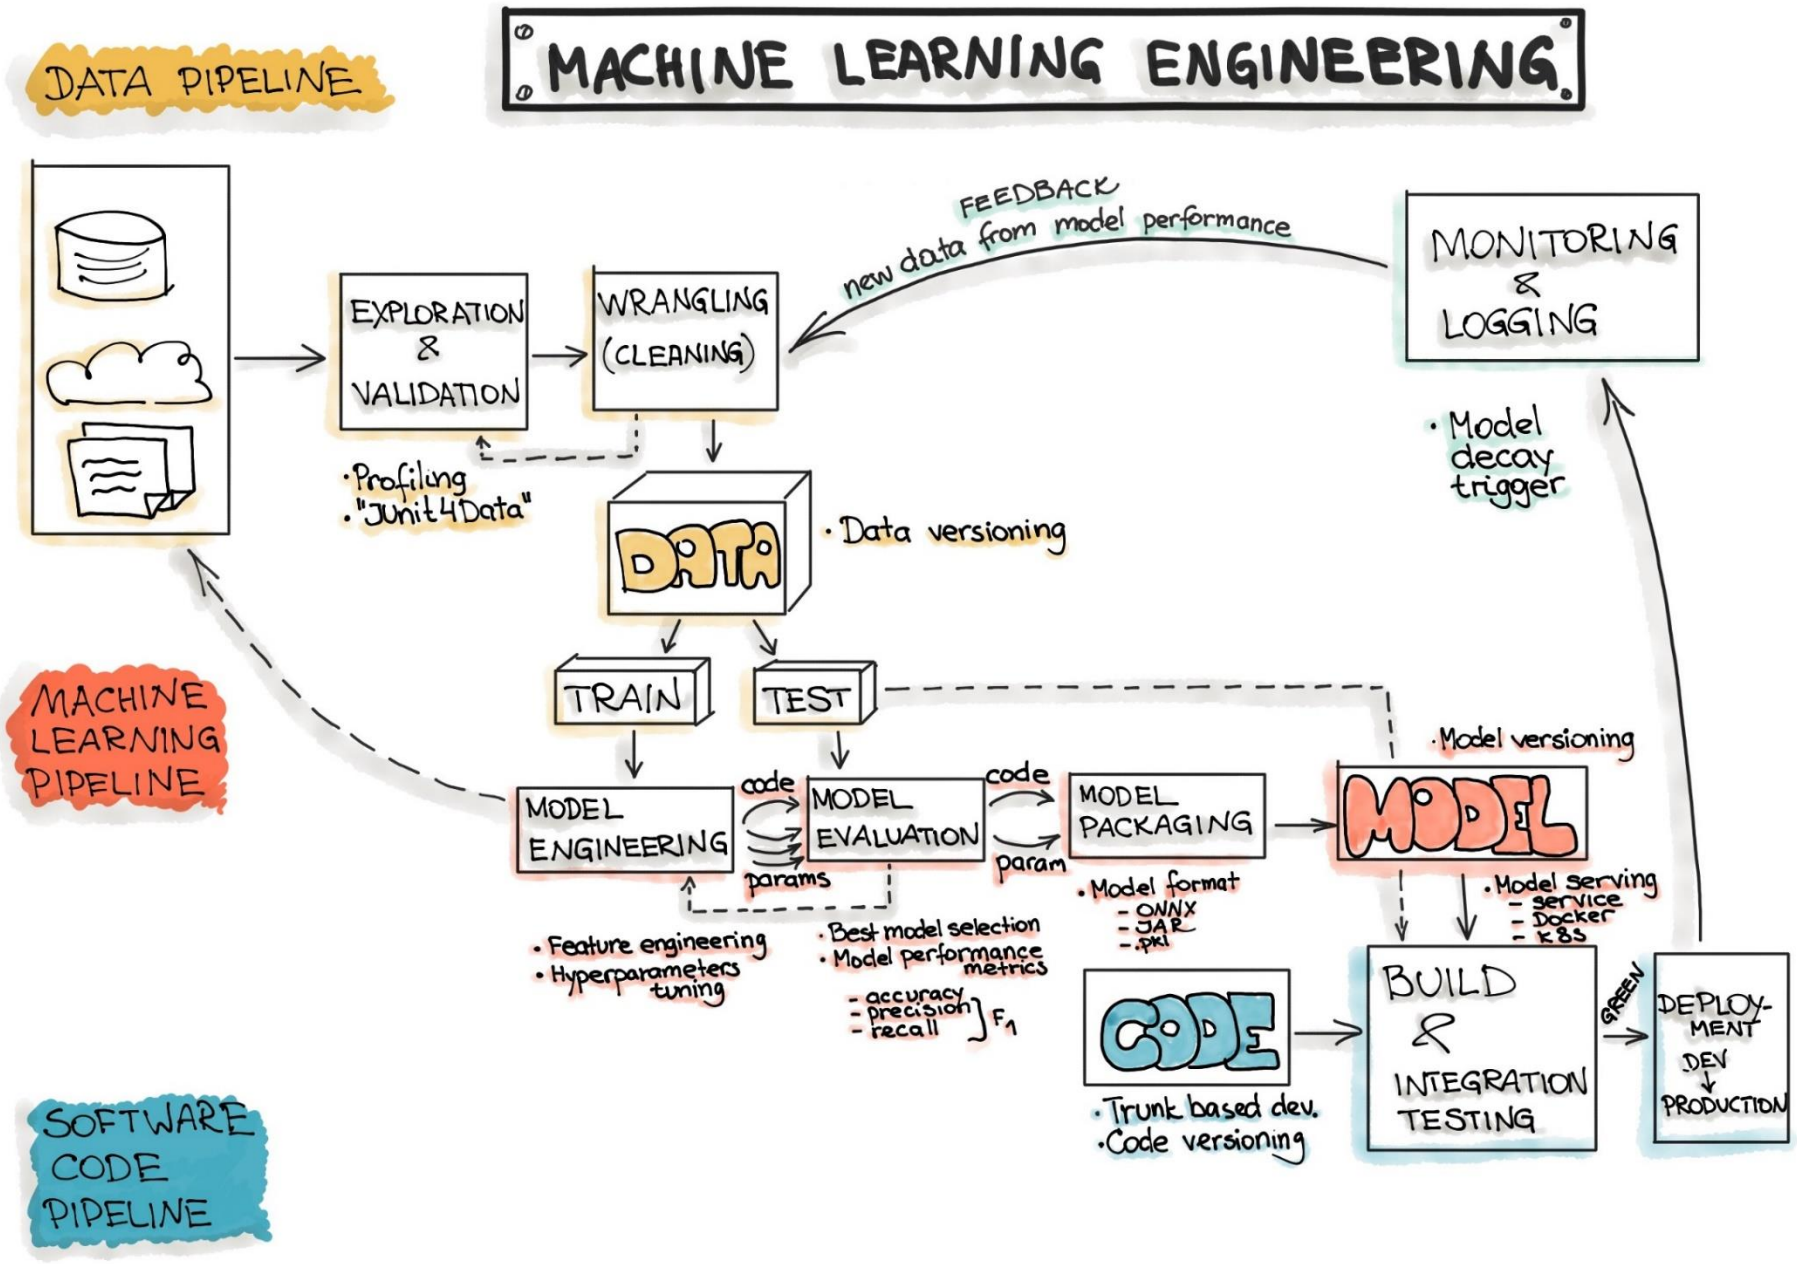

Source: https://ml-ops.org/content/end-to-end-ml-workflow

Machine Learning Taxonomy

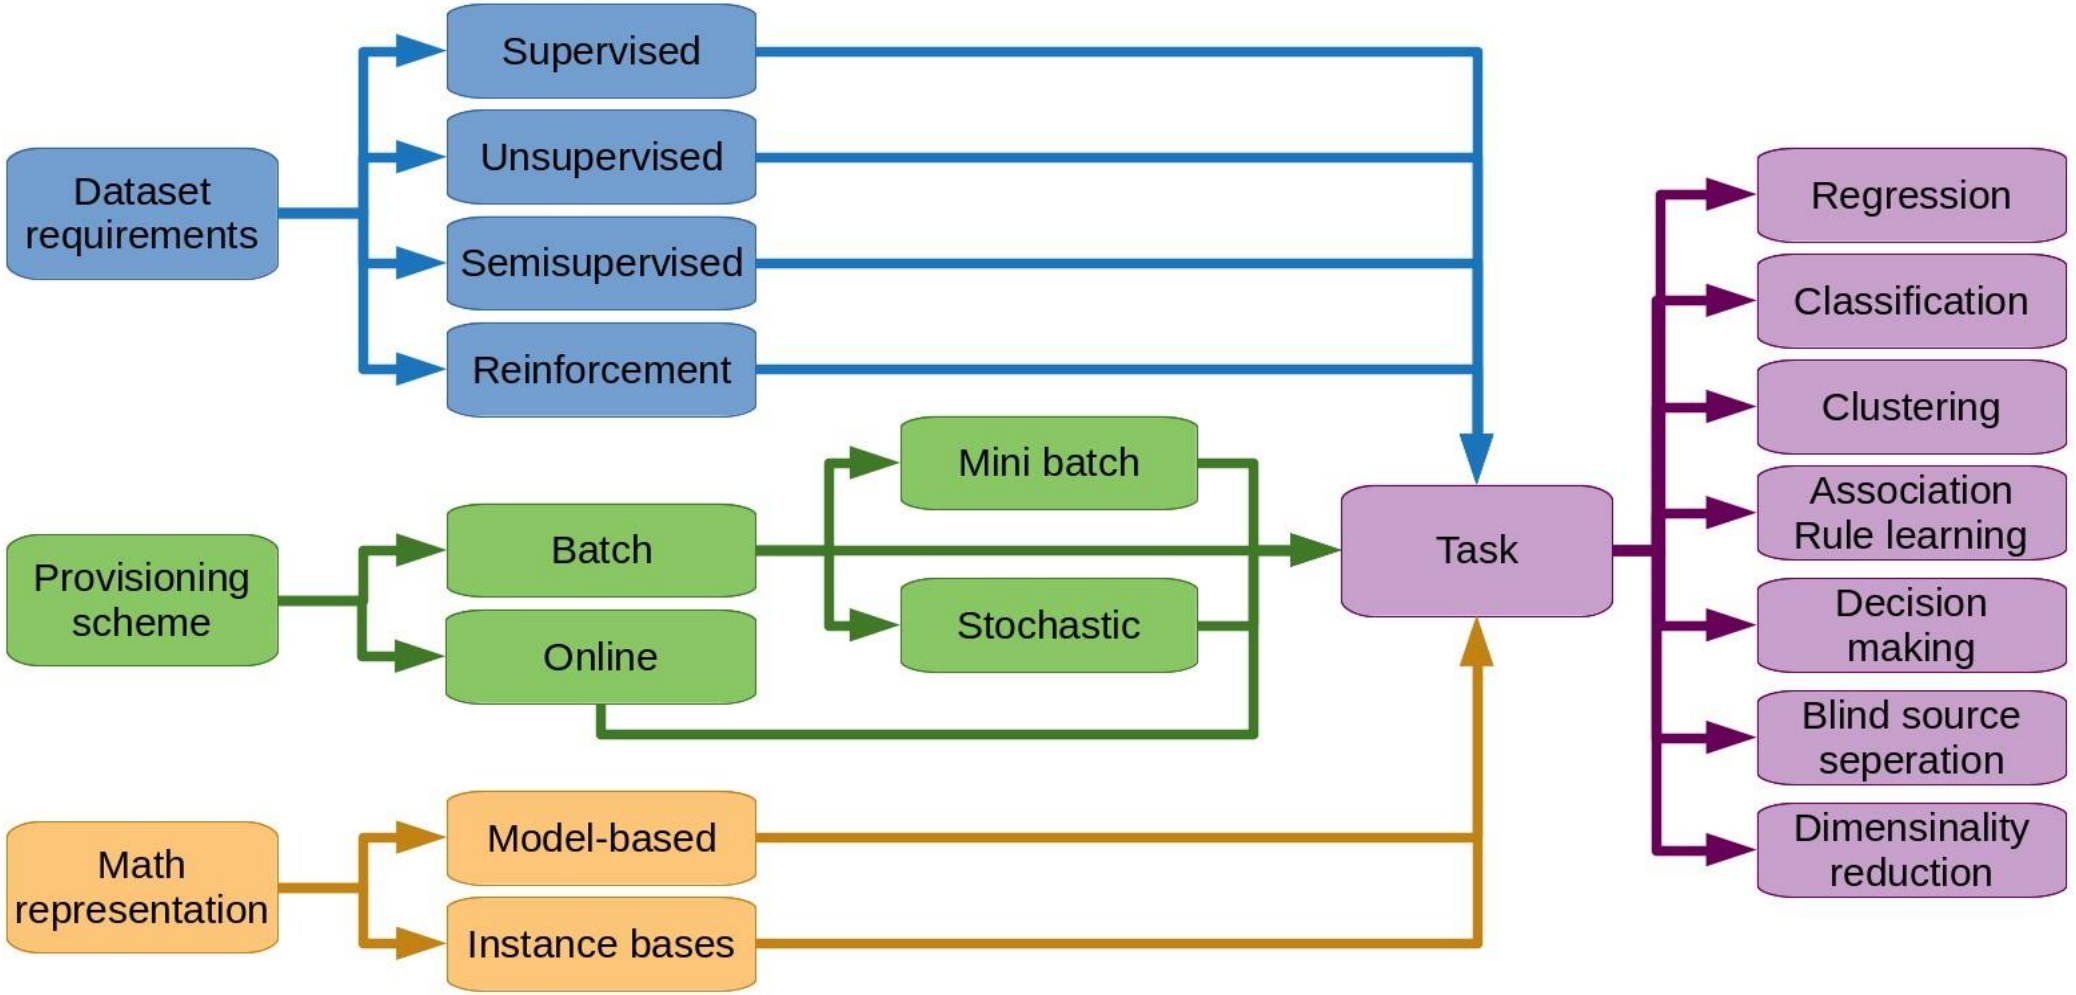

TU Delft

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Classification by Dataset Requirements

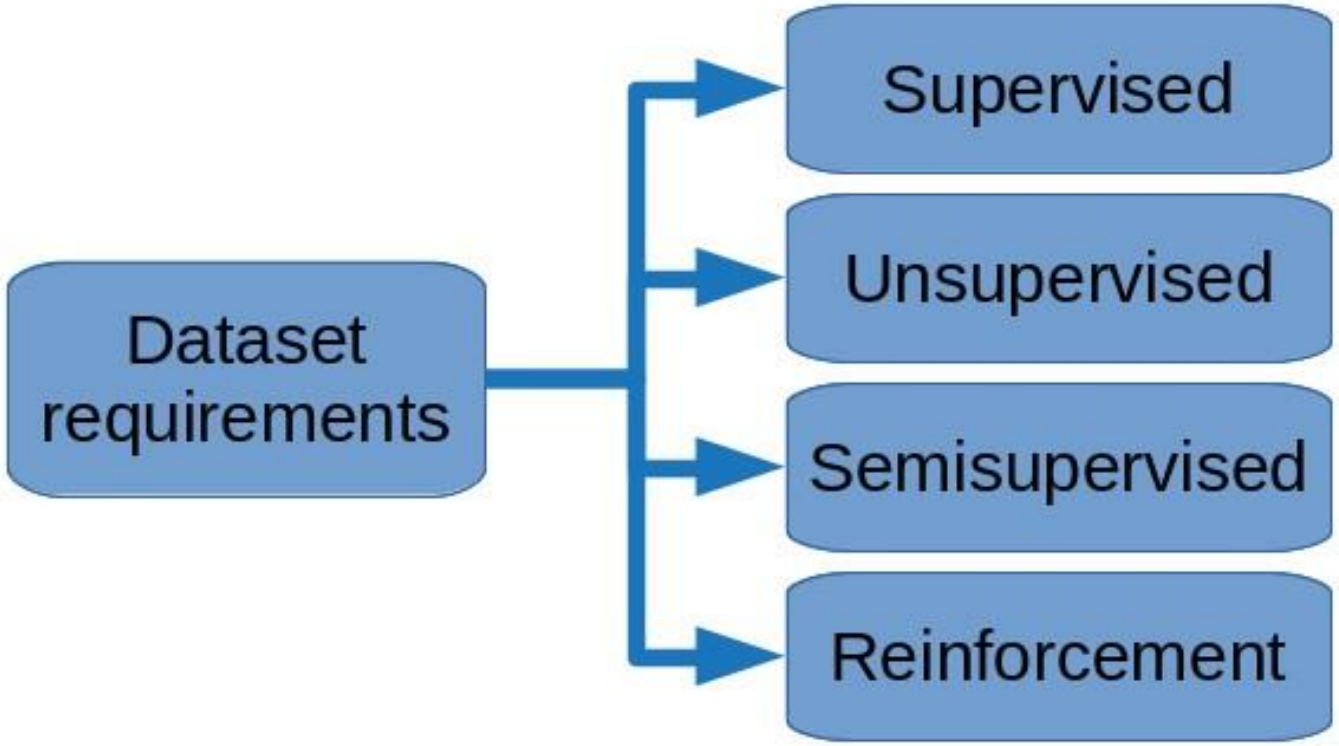

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

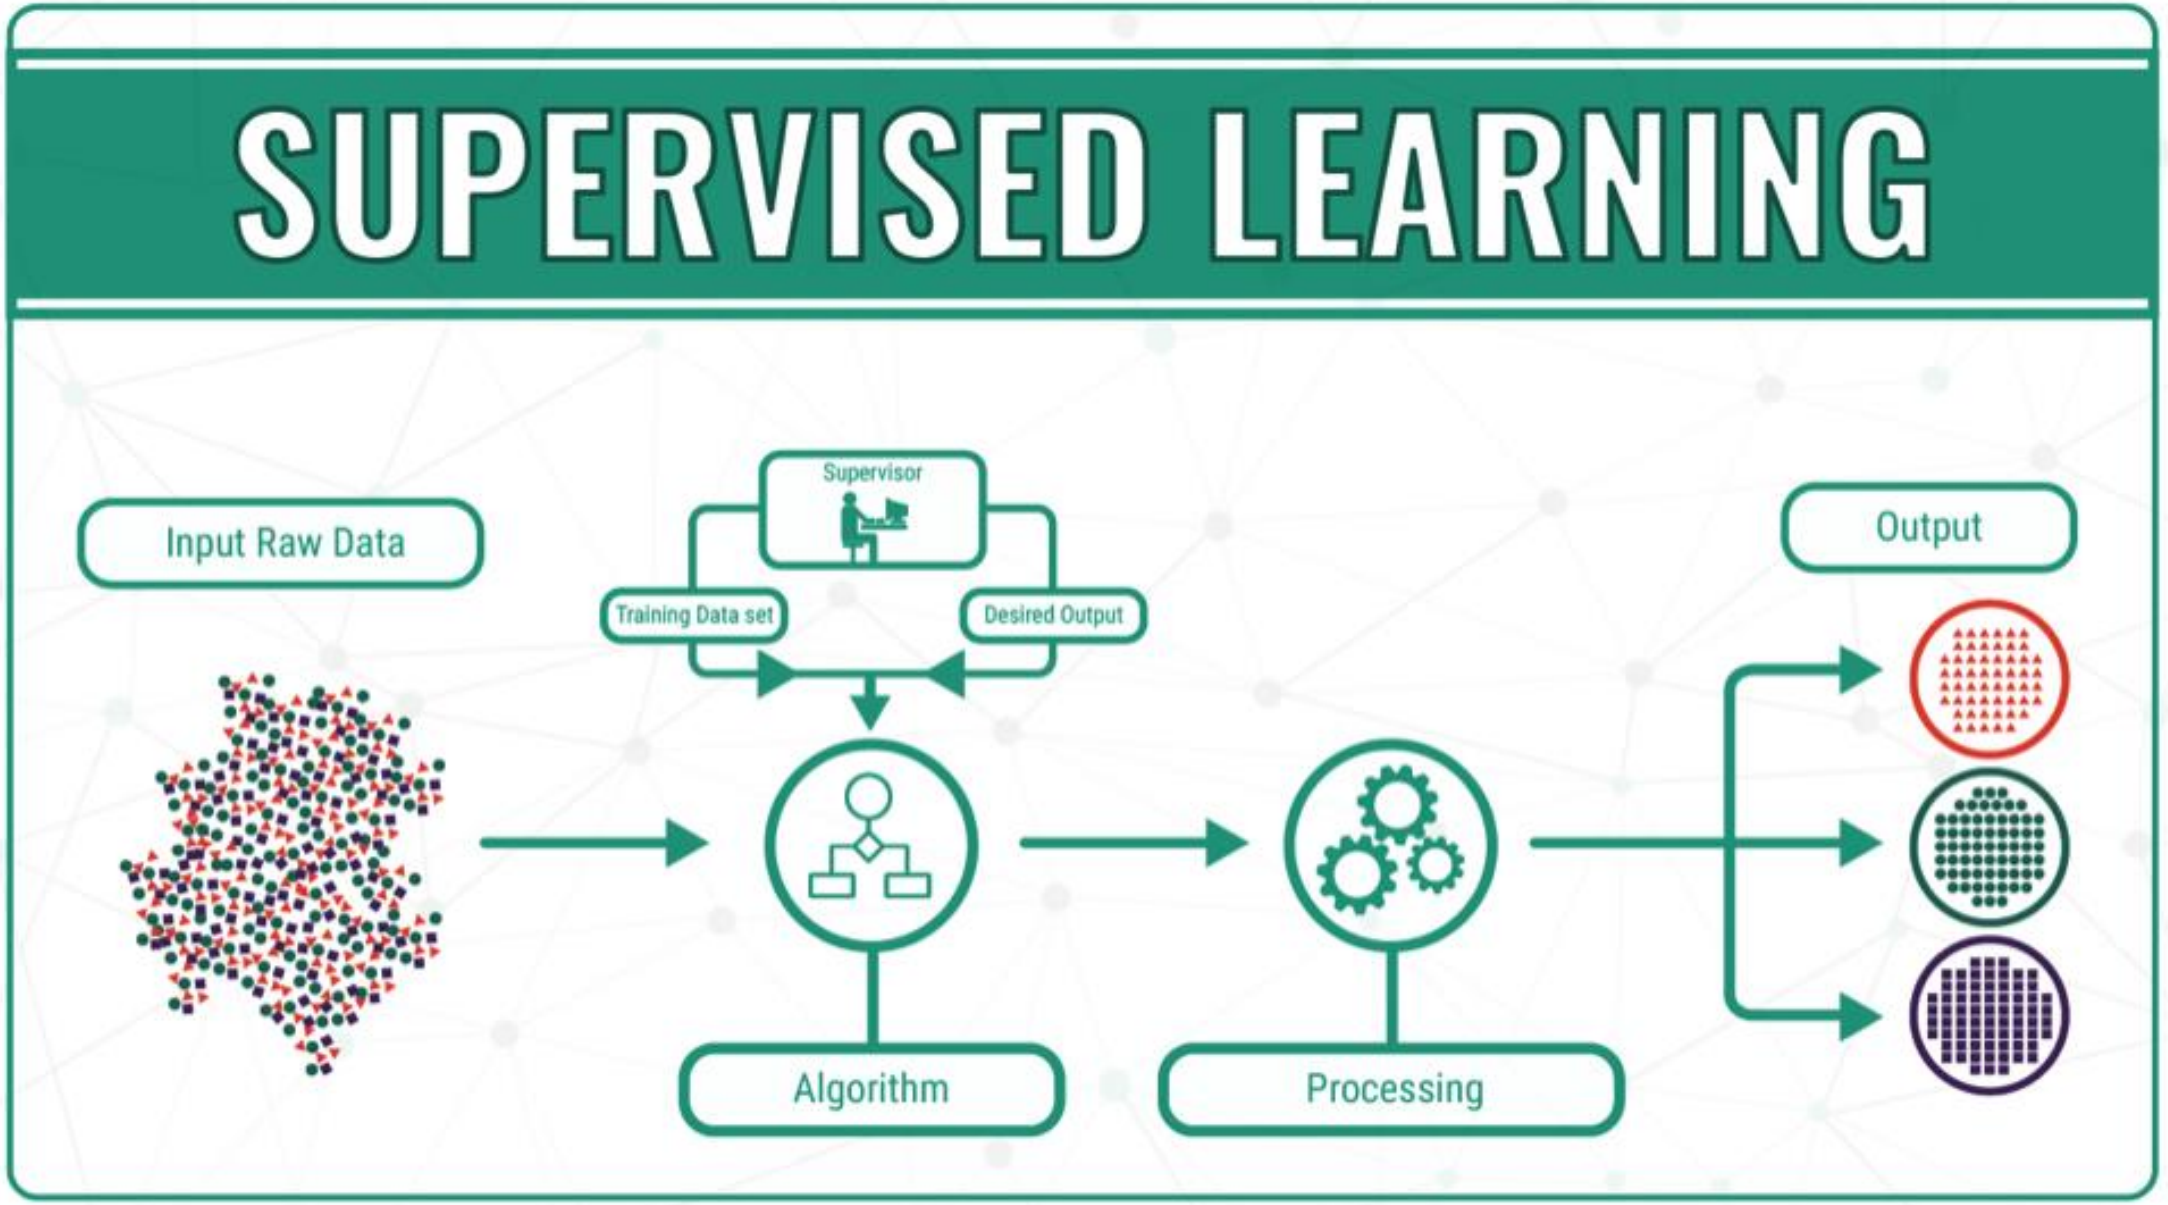

TU Delft

Source: https://www.crayondata.com/machine-learning-explained-understanding-supervised-unsupervised-and-reinforcement-learning/

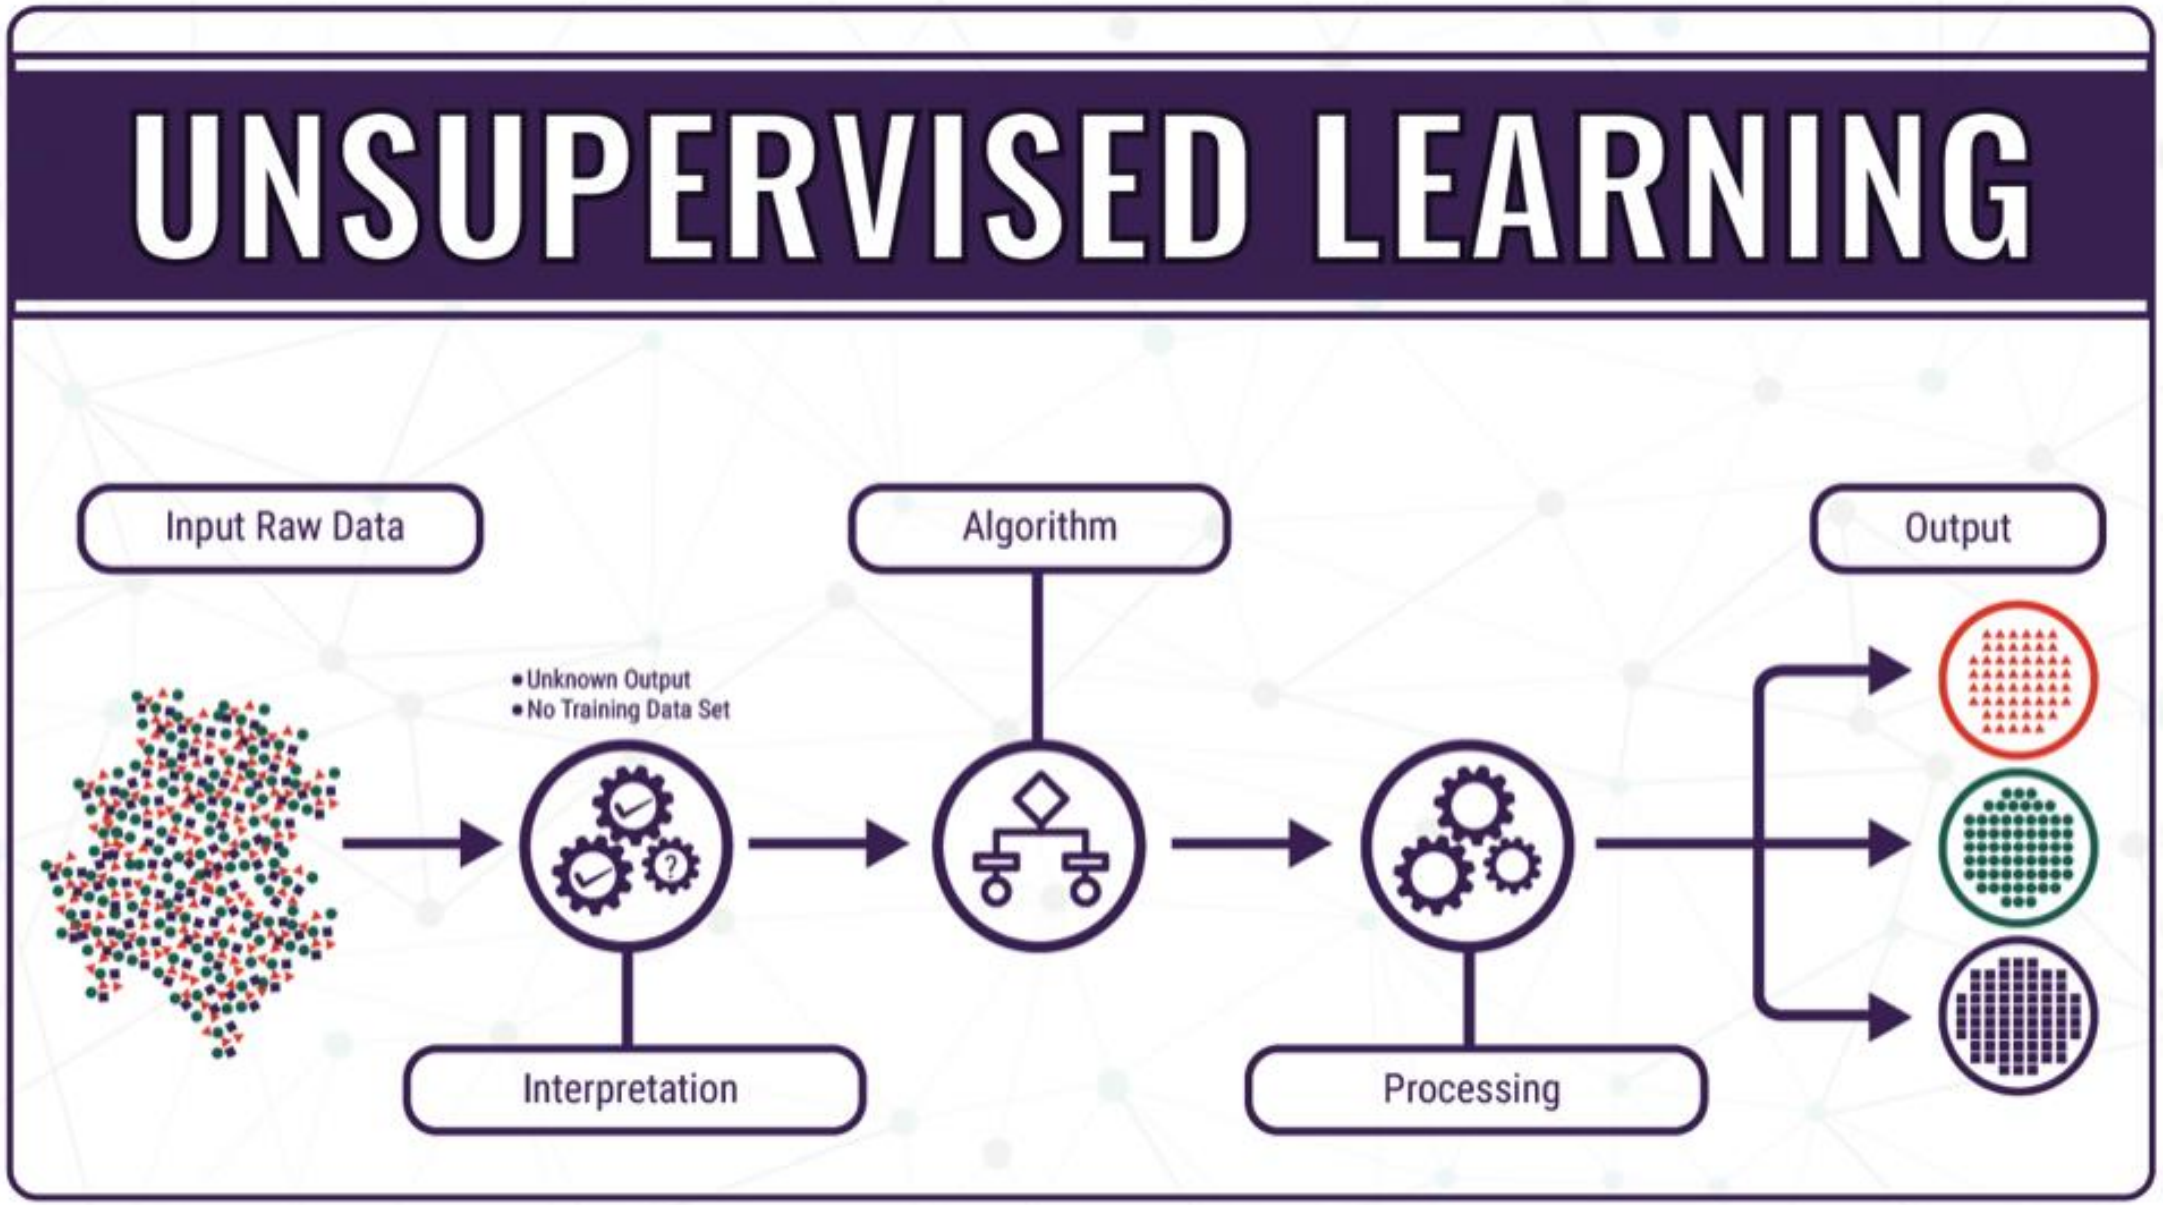

TU Delft

Source: https://www.crayondata.com/machine-learning-explained-understanding-supervised-unsupervised-and-reinforcement-learning/

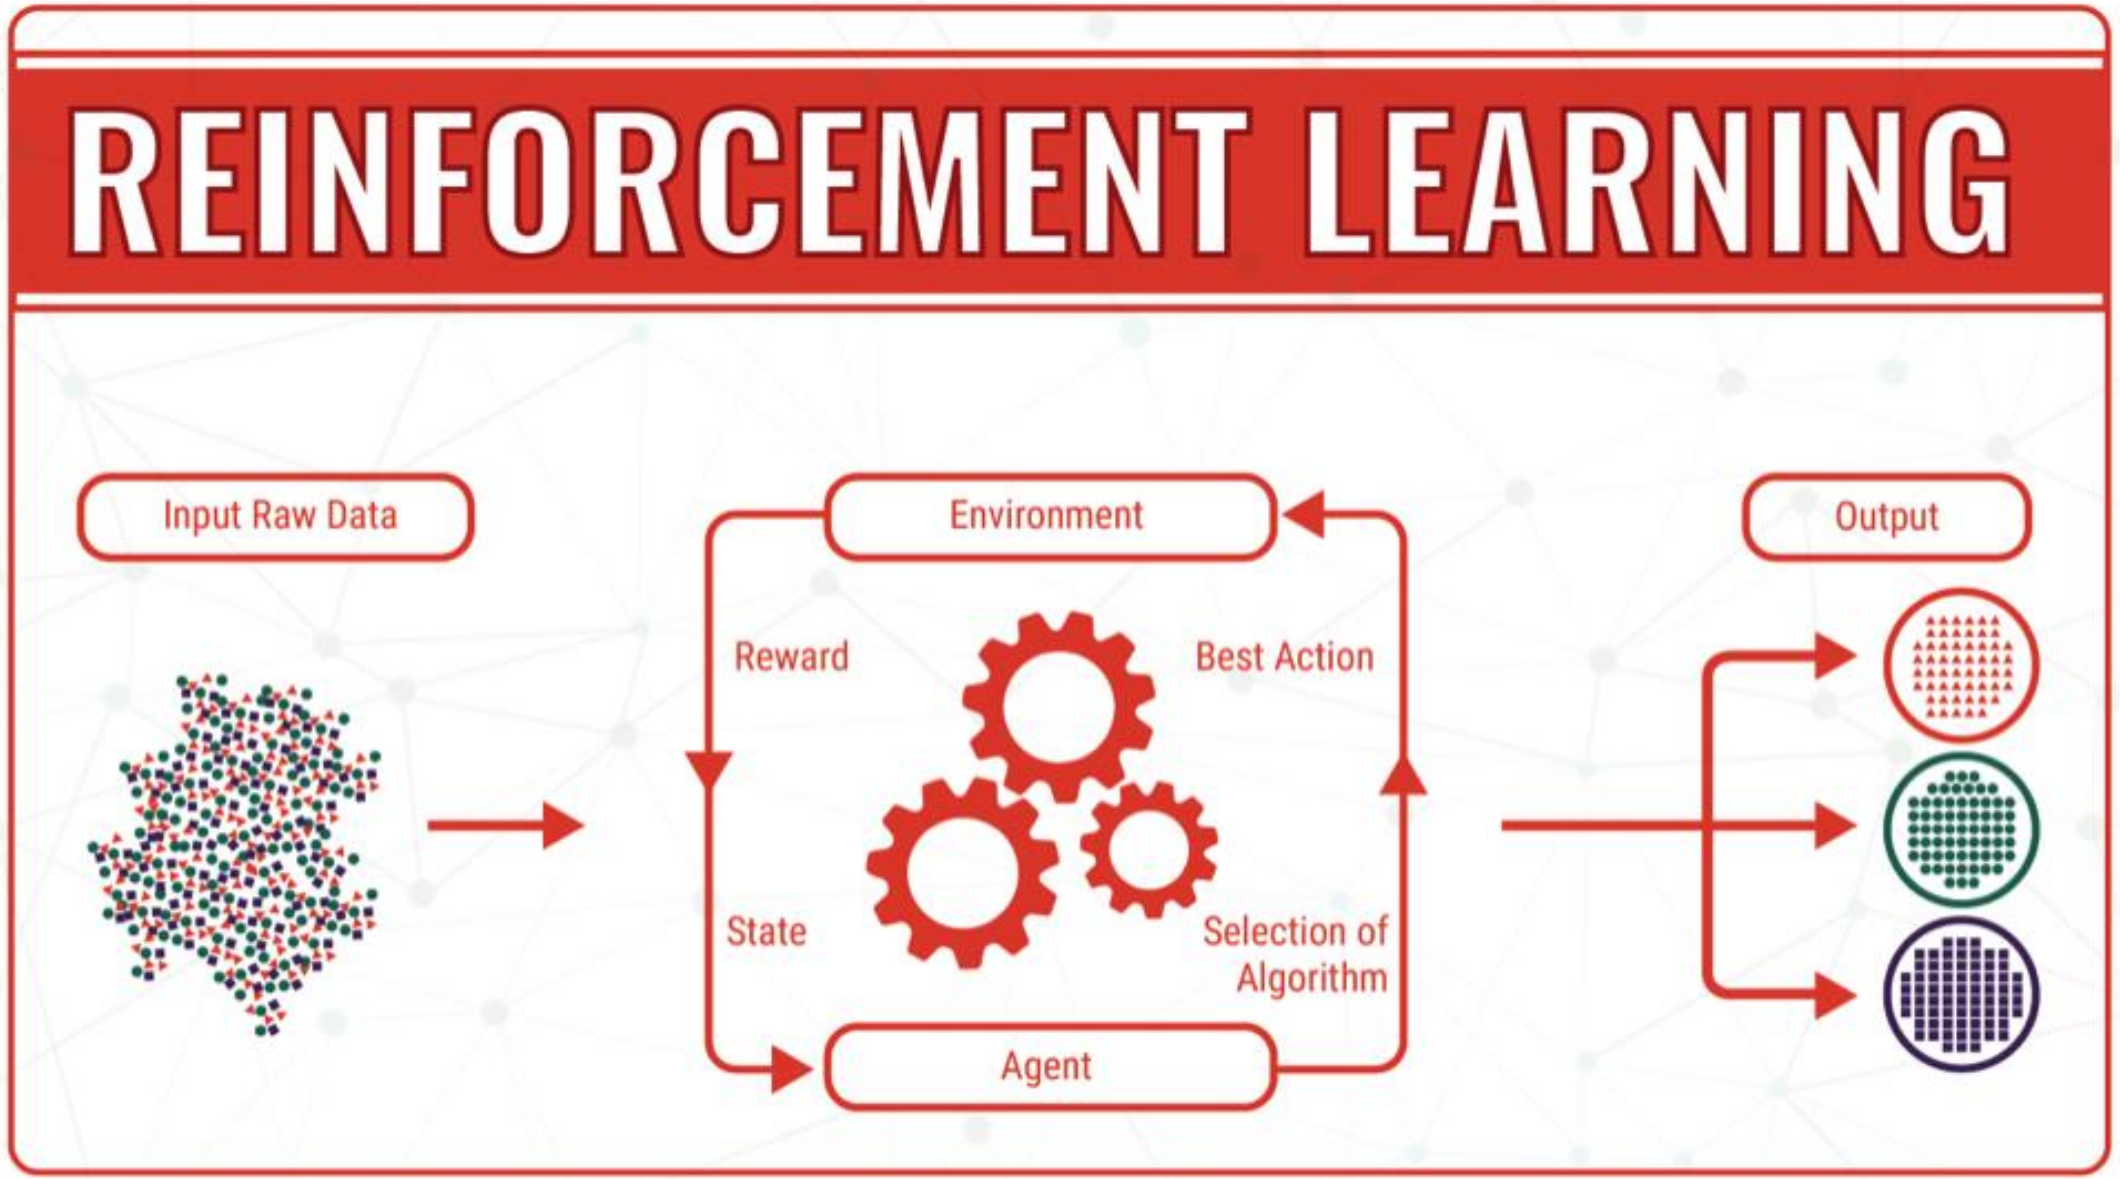

TU Delft

Source: https://www.crayondata.com/machine-learning-explained-understanding-supervised-unsupervised-and-reinforcement-learning/

Classification by Model Representation

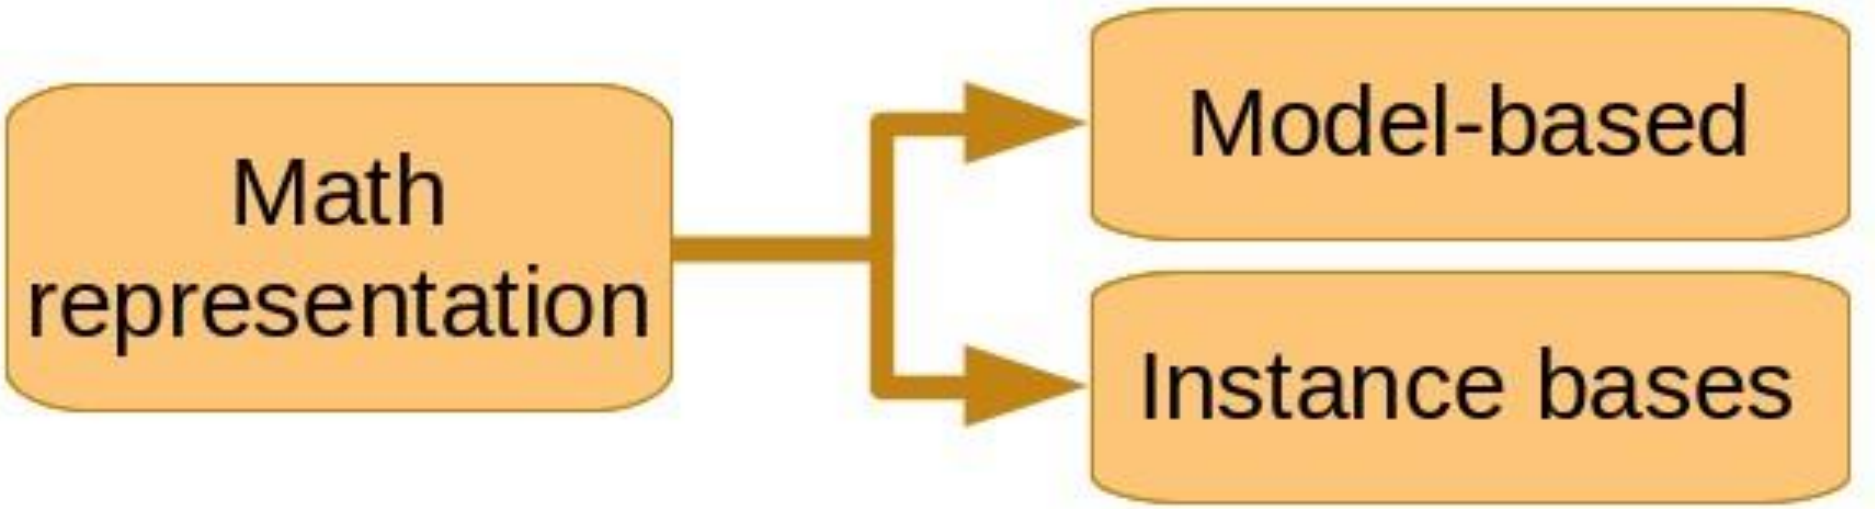

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Instance-based learning

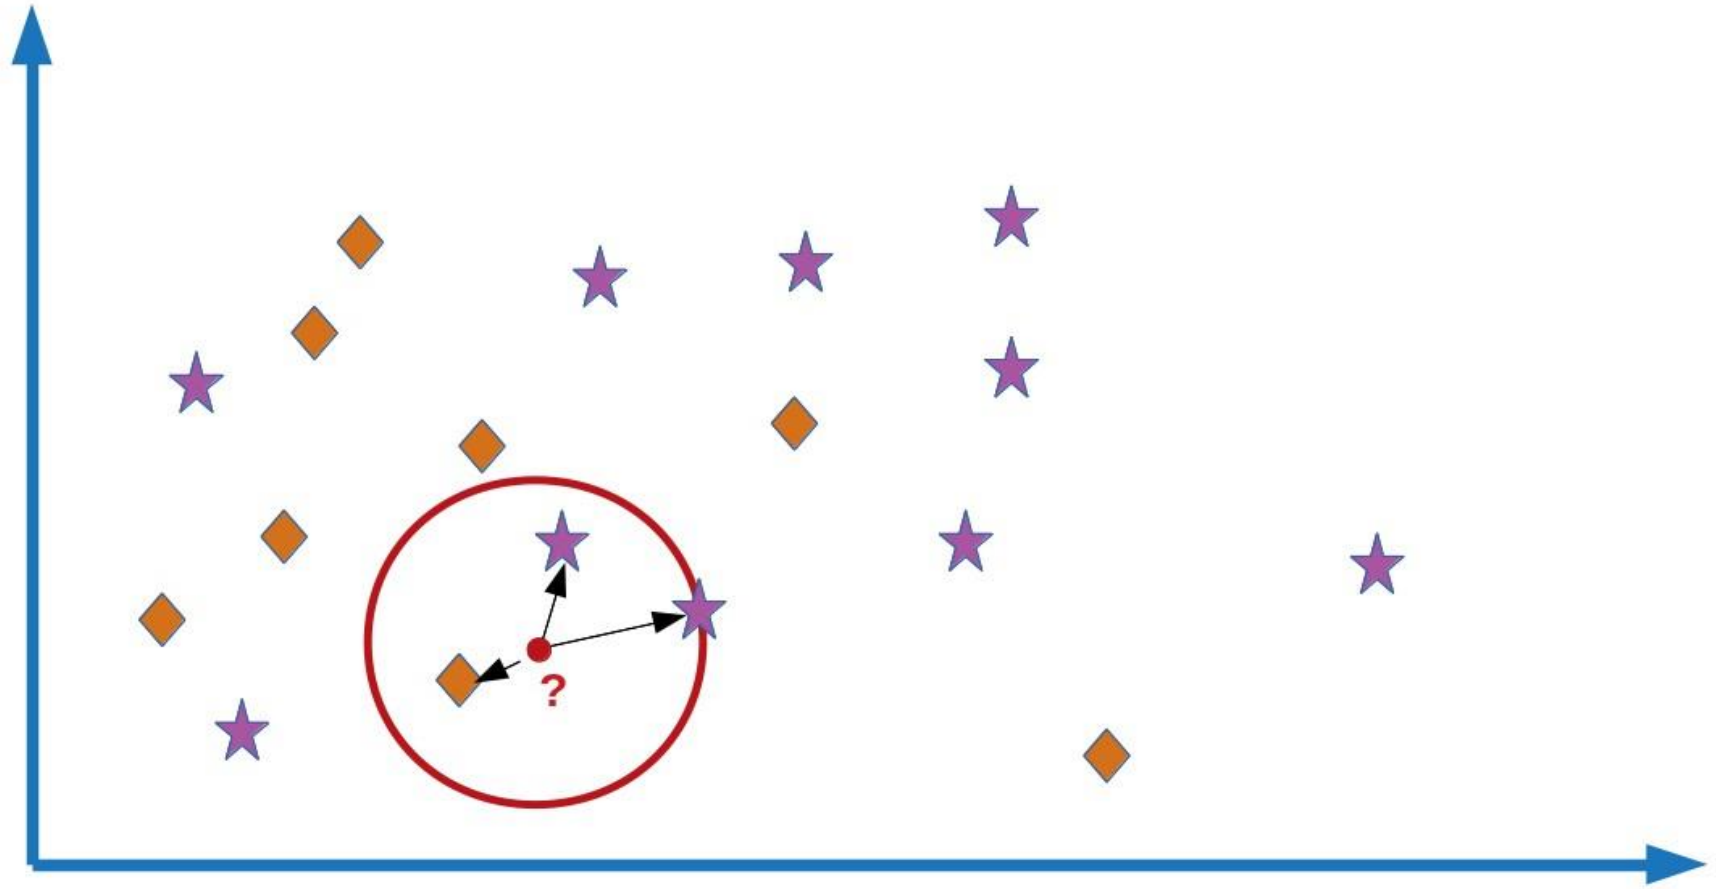

TU Delft

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Model-based learning

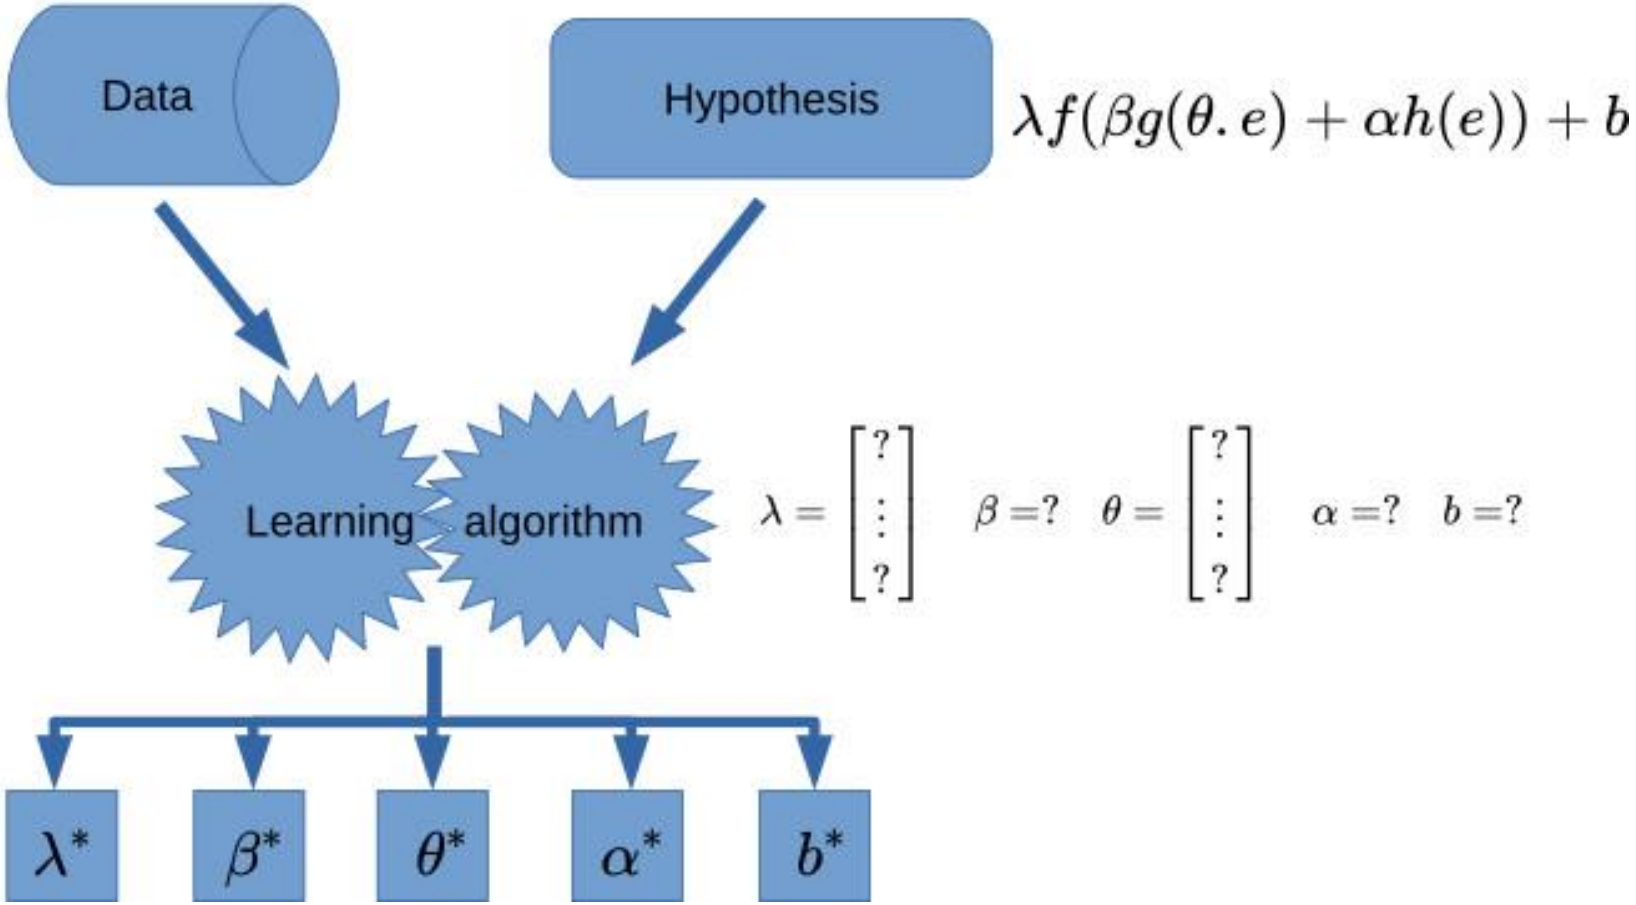

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Classification by Training Behavior

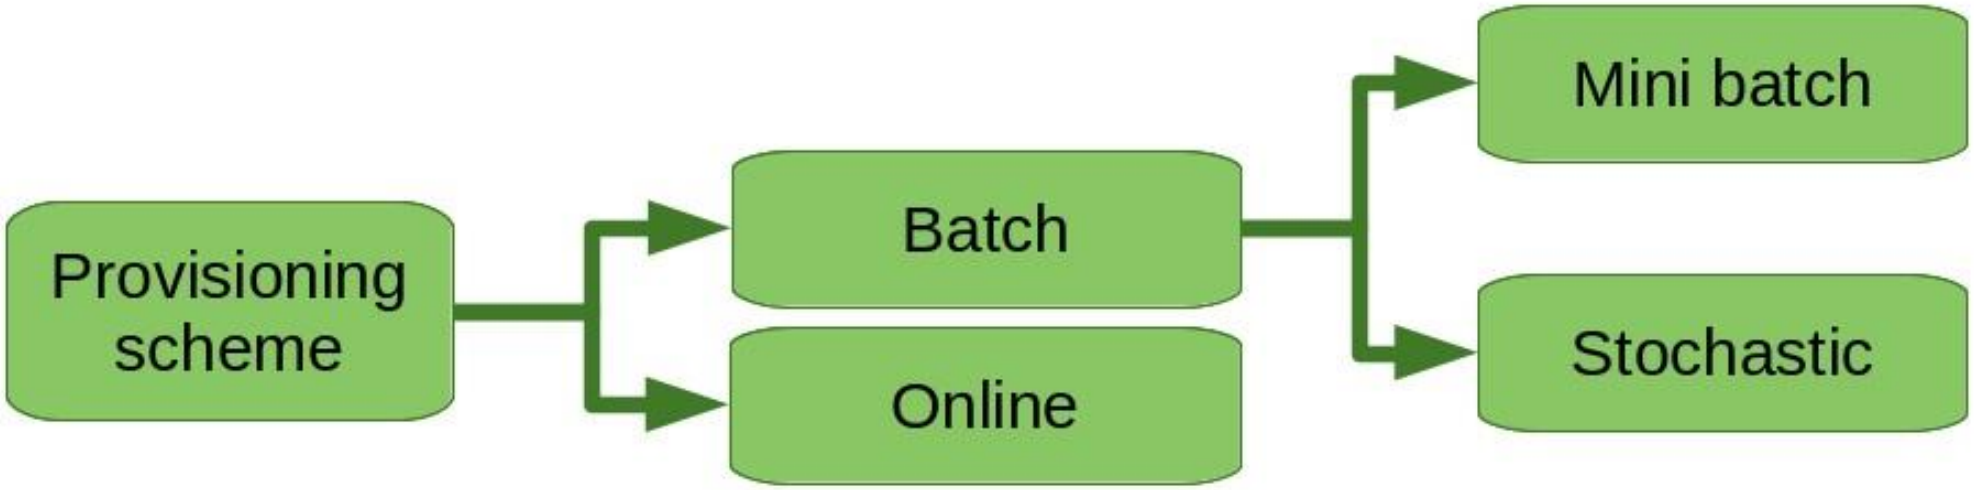

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Batch learning

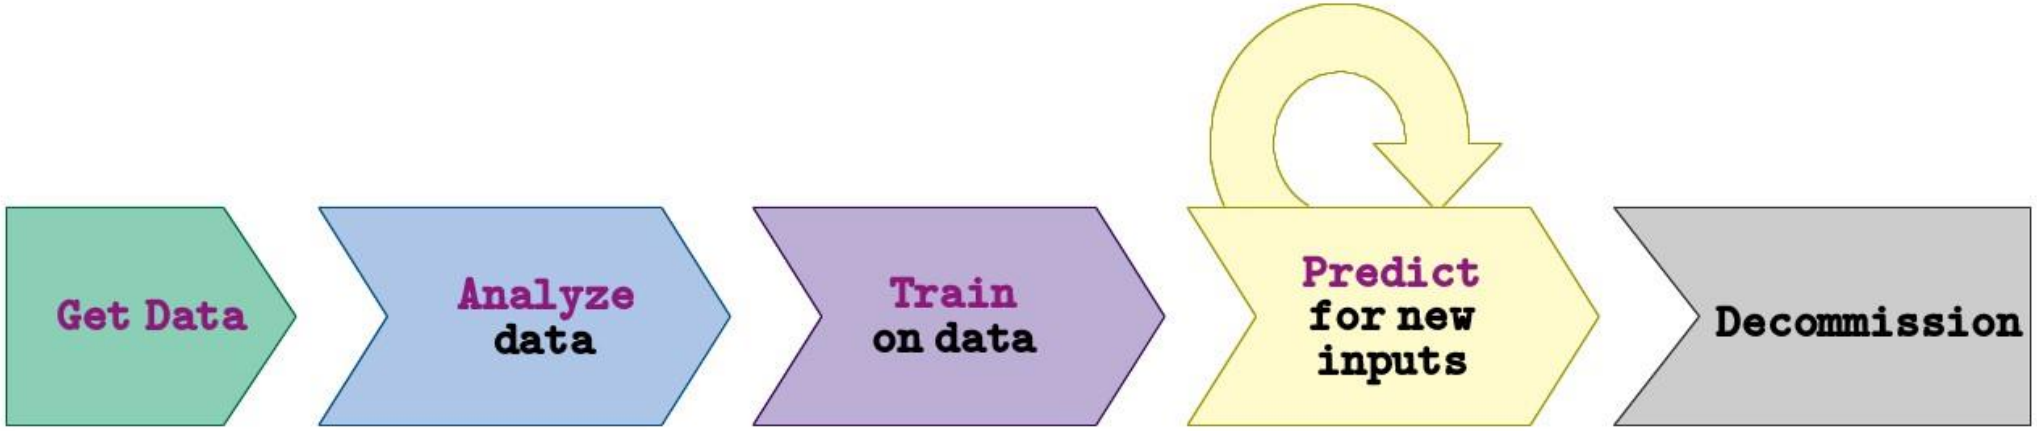

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Online learning

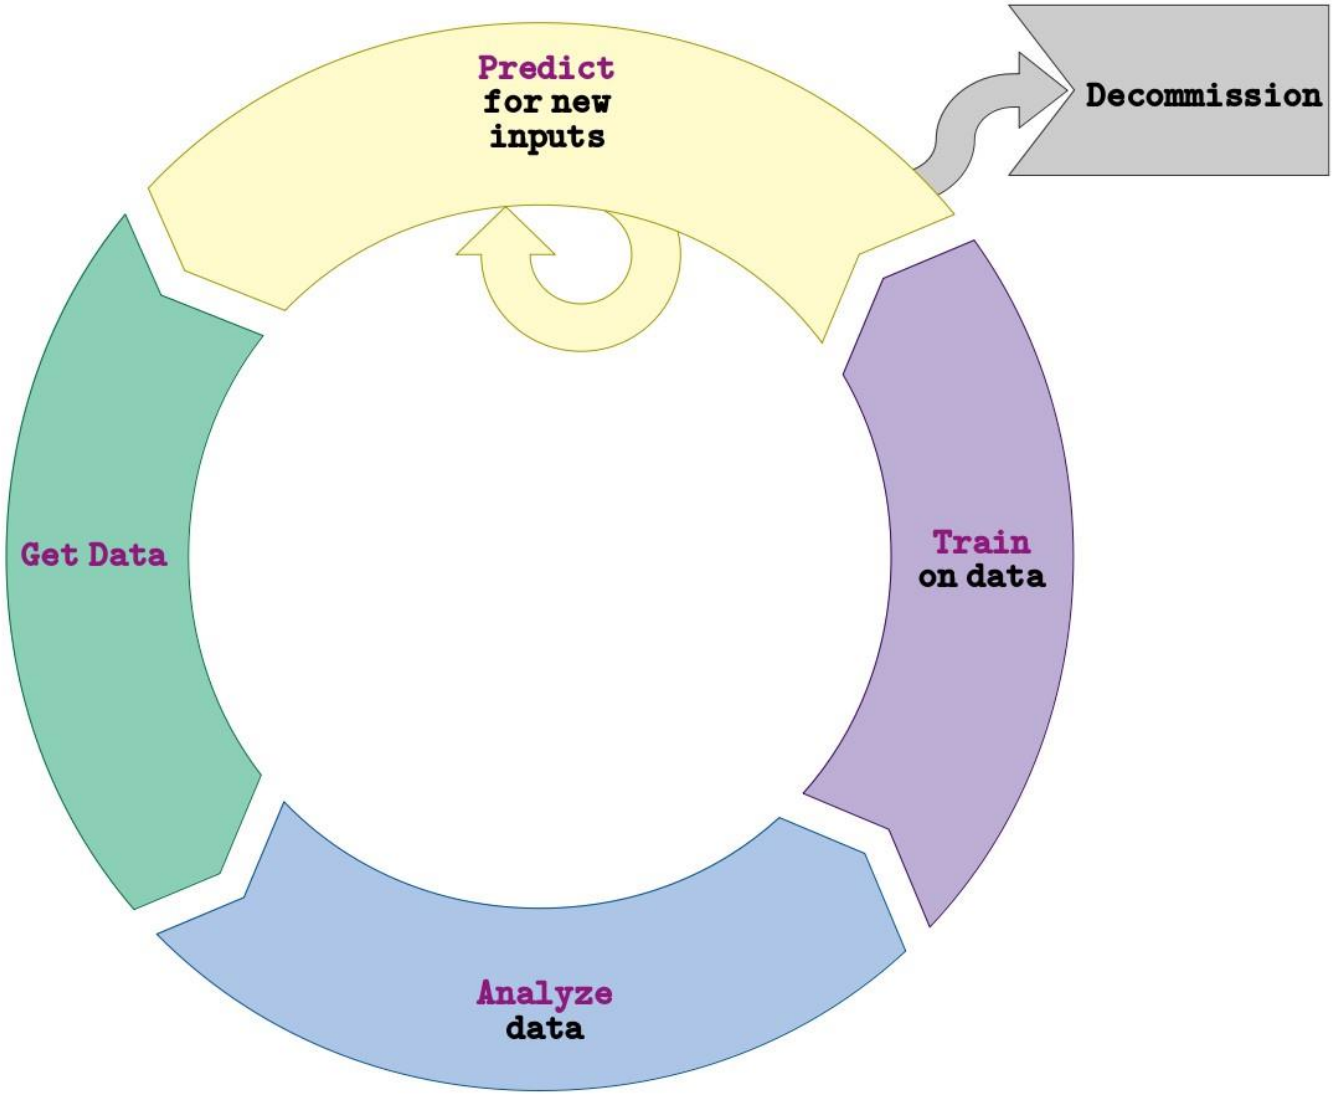

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Classification by Task Type

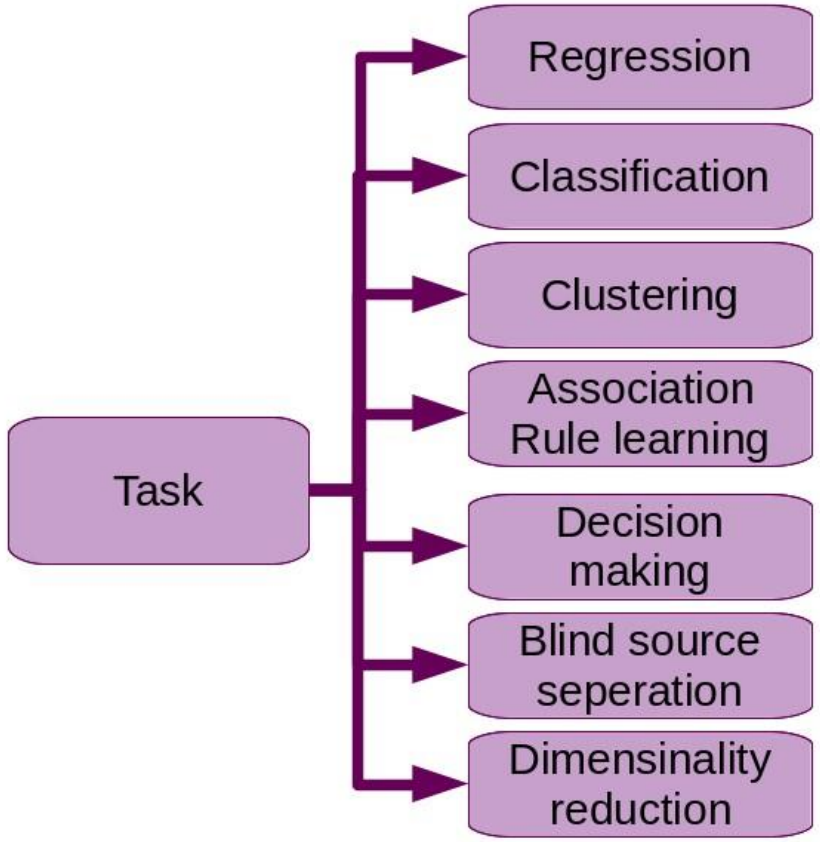

Source: https://hermit-notebook.site/en/notebook/computer-sciences/artificial-intelligence/machine-learning/taxonomy-of-machine-learning/

Classification

- $Y = f(X)$
- The values of Y are determined by a human
- $Y \in \{C_1, \ldots, C_k\}$ is a discrete variable
  - C_1 = Triangle
  - C_2 = Circle
- f is learned from the data through ML

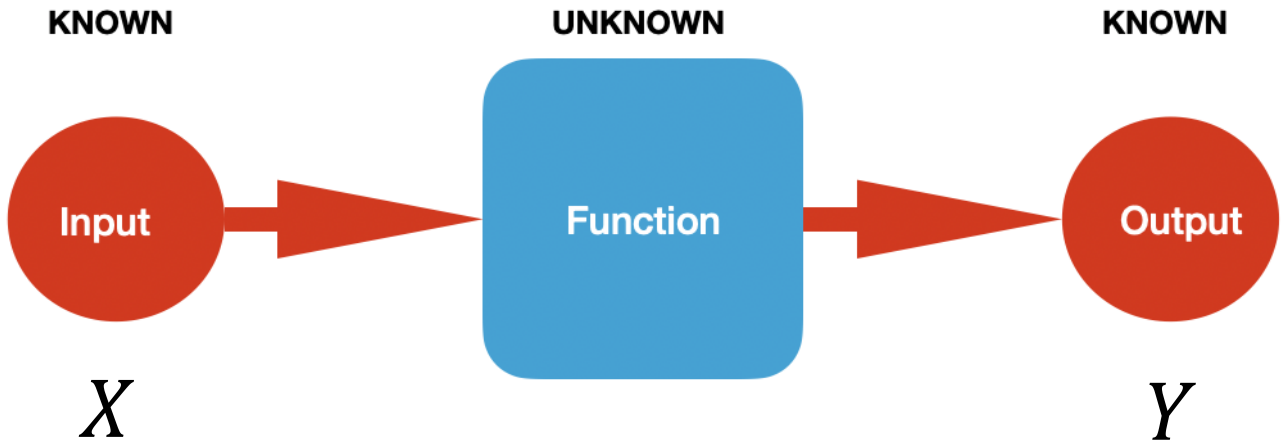

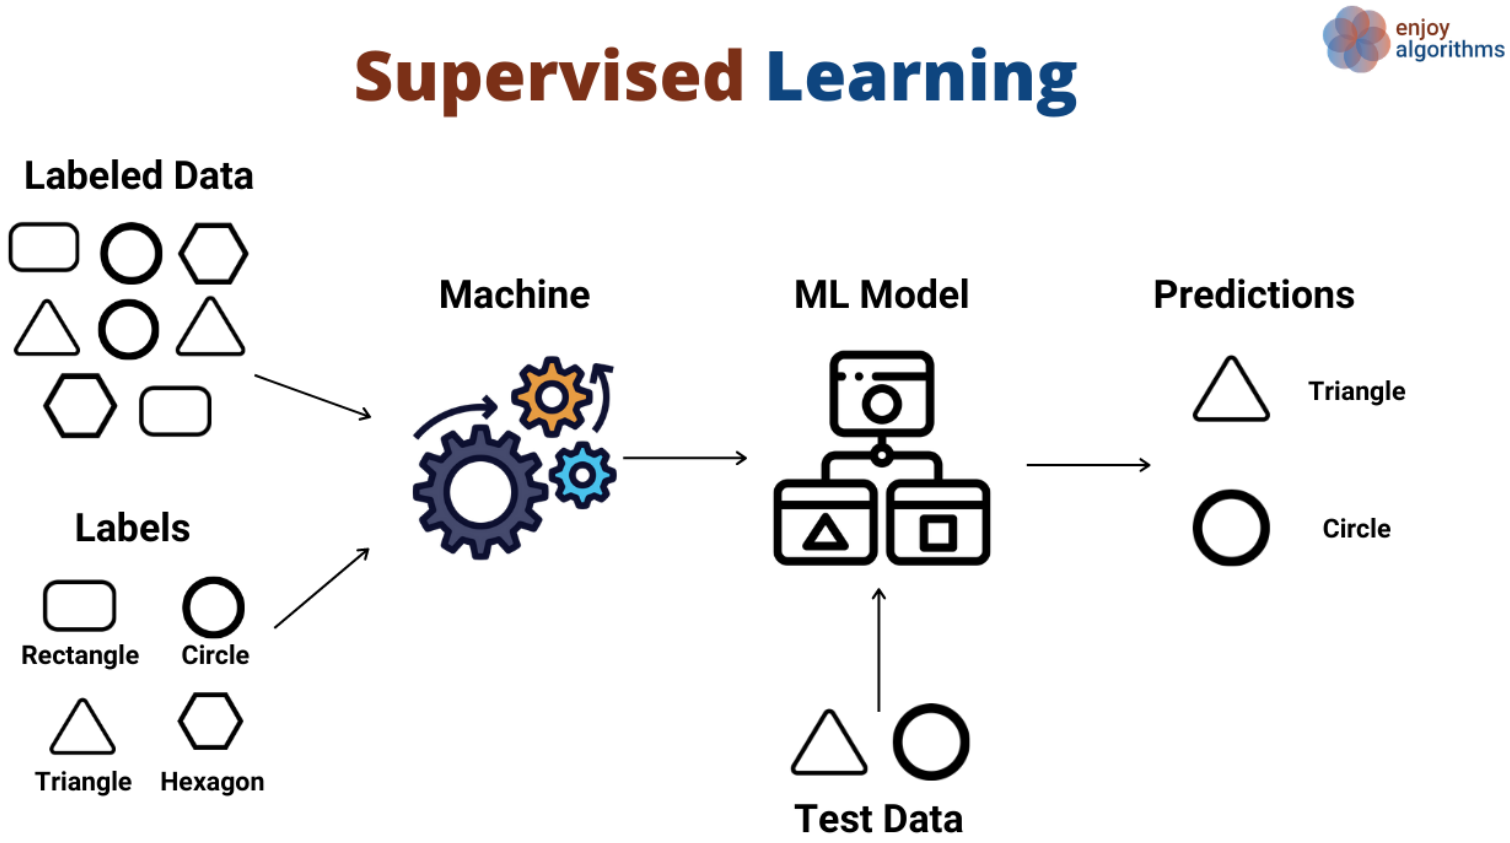

Source: https://www.enjoyalgorithms.com/blogs/supervised-unsupervised-and-semisupervised-learning

TU Delft

Regression

- $Y = f(X)$
- The values of Y are determined by a human
- $Y \in \mathbb{R}$ is a continuous variable
- f is learned from the data through ML

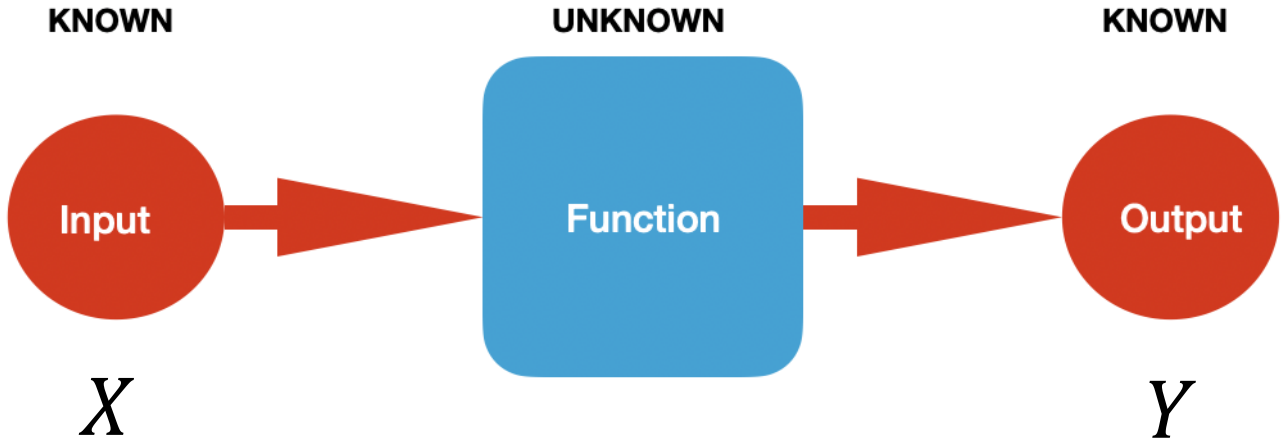

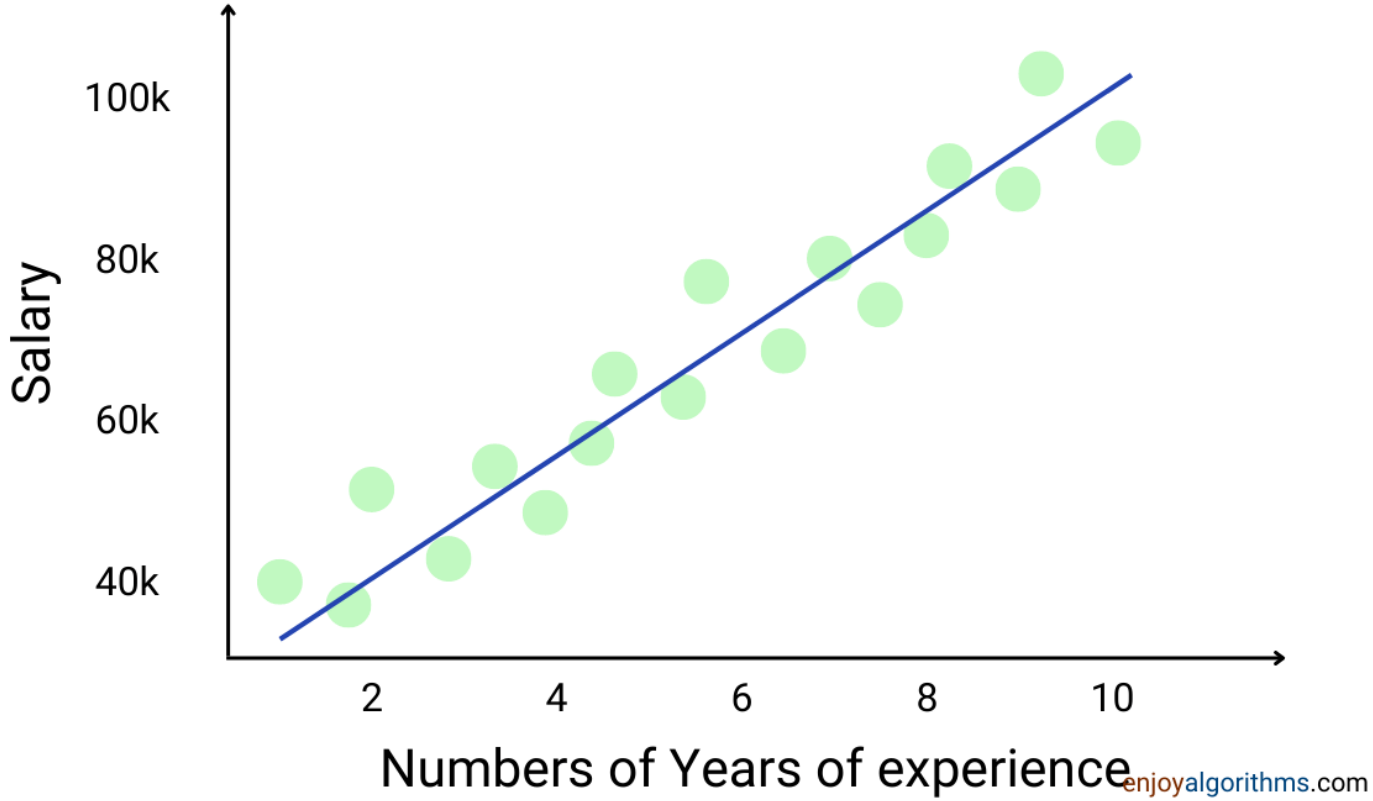

Source: https://www.enjoyalgorithms.com/blogs/supervised-unsupervised-and-semisupervised-learning

TU Delft

Clustering

- The real values of Y are unknown
- The ML algorithm tries to identify existing patterns in the data (without prior supervision)

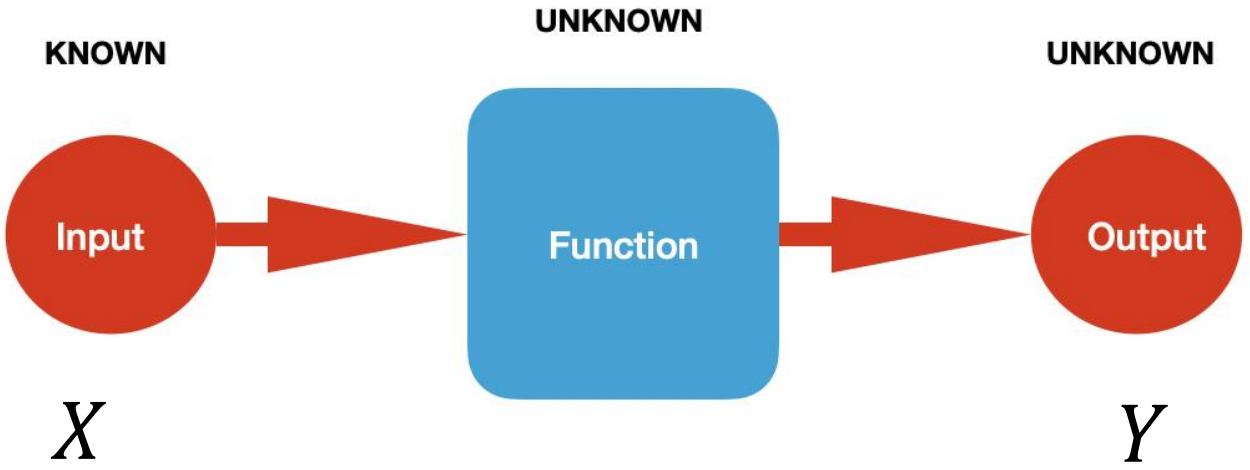

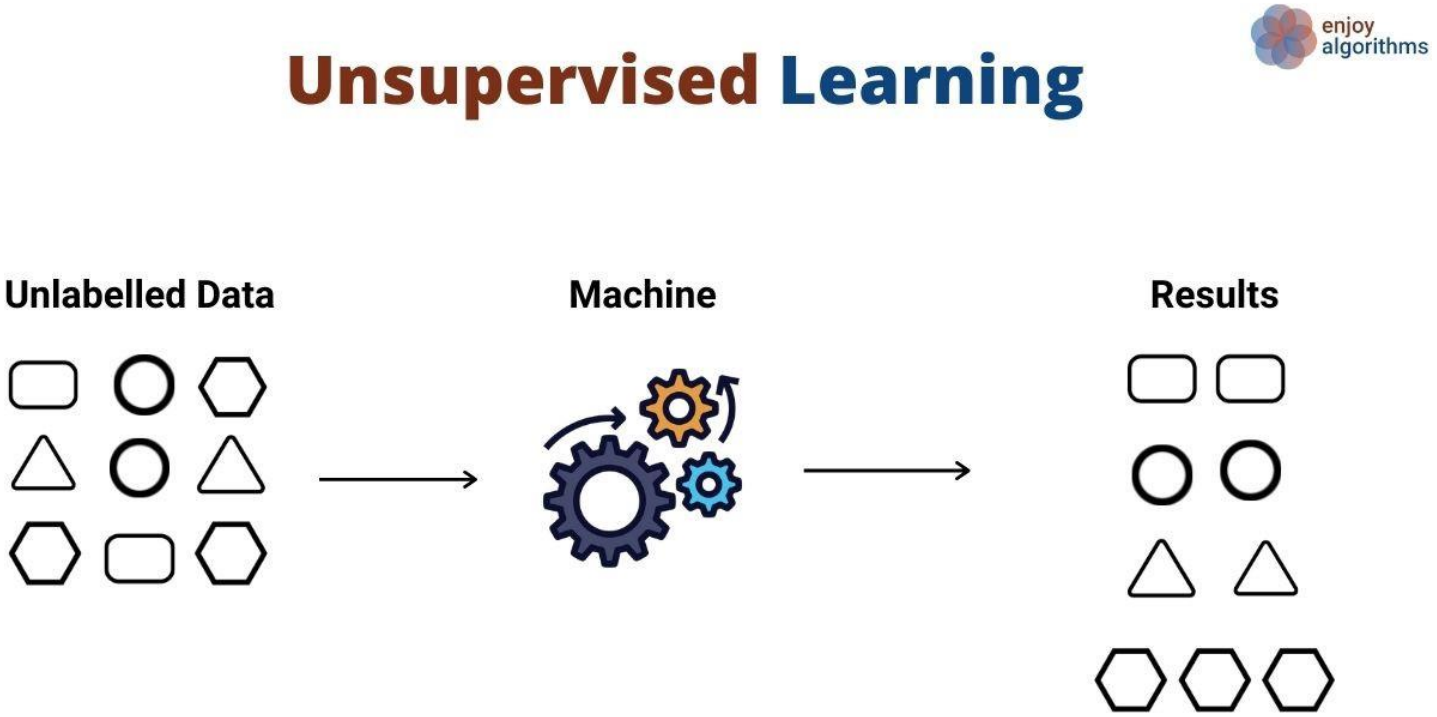

Source: https://www.enjoyalgorithms.com/blogs/supervised-unsupervised-and-semisupervised-learning

TU Delft

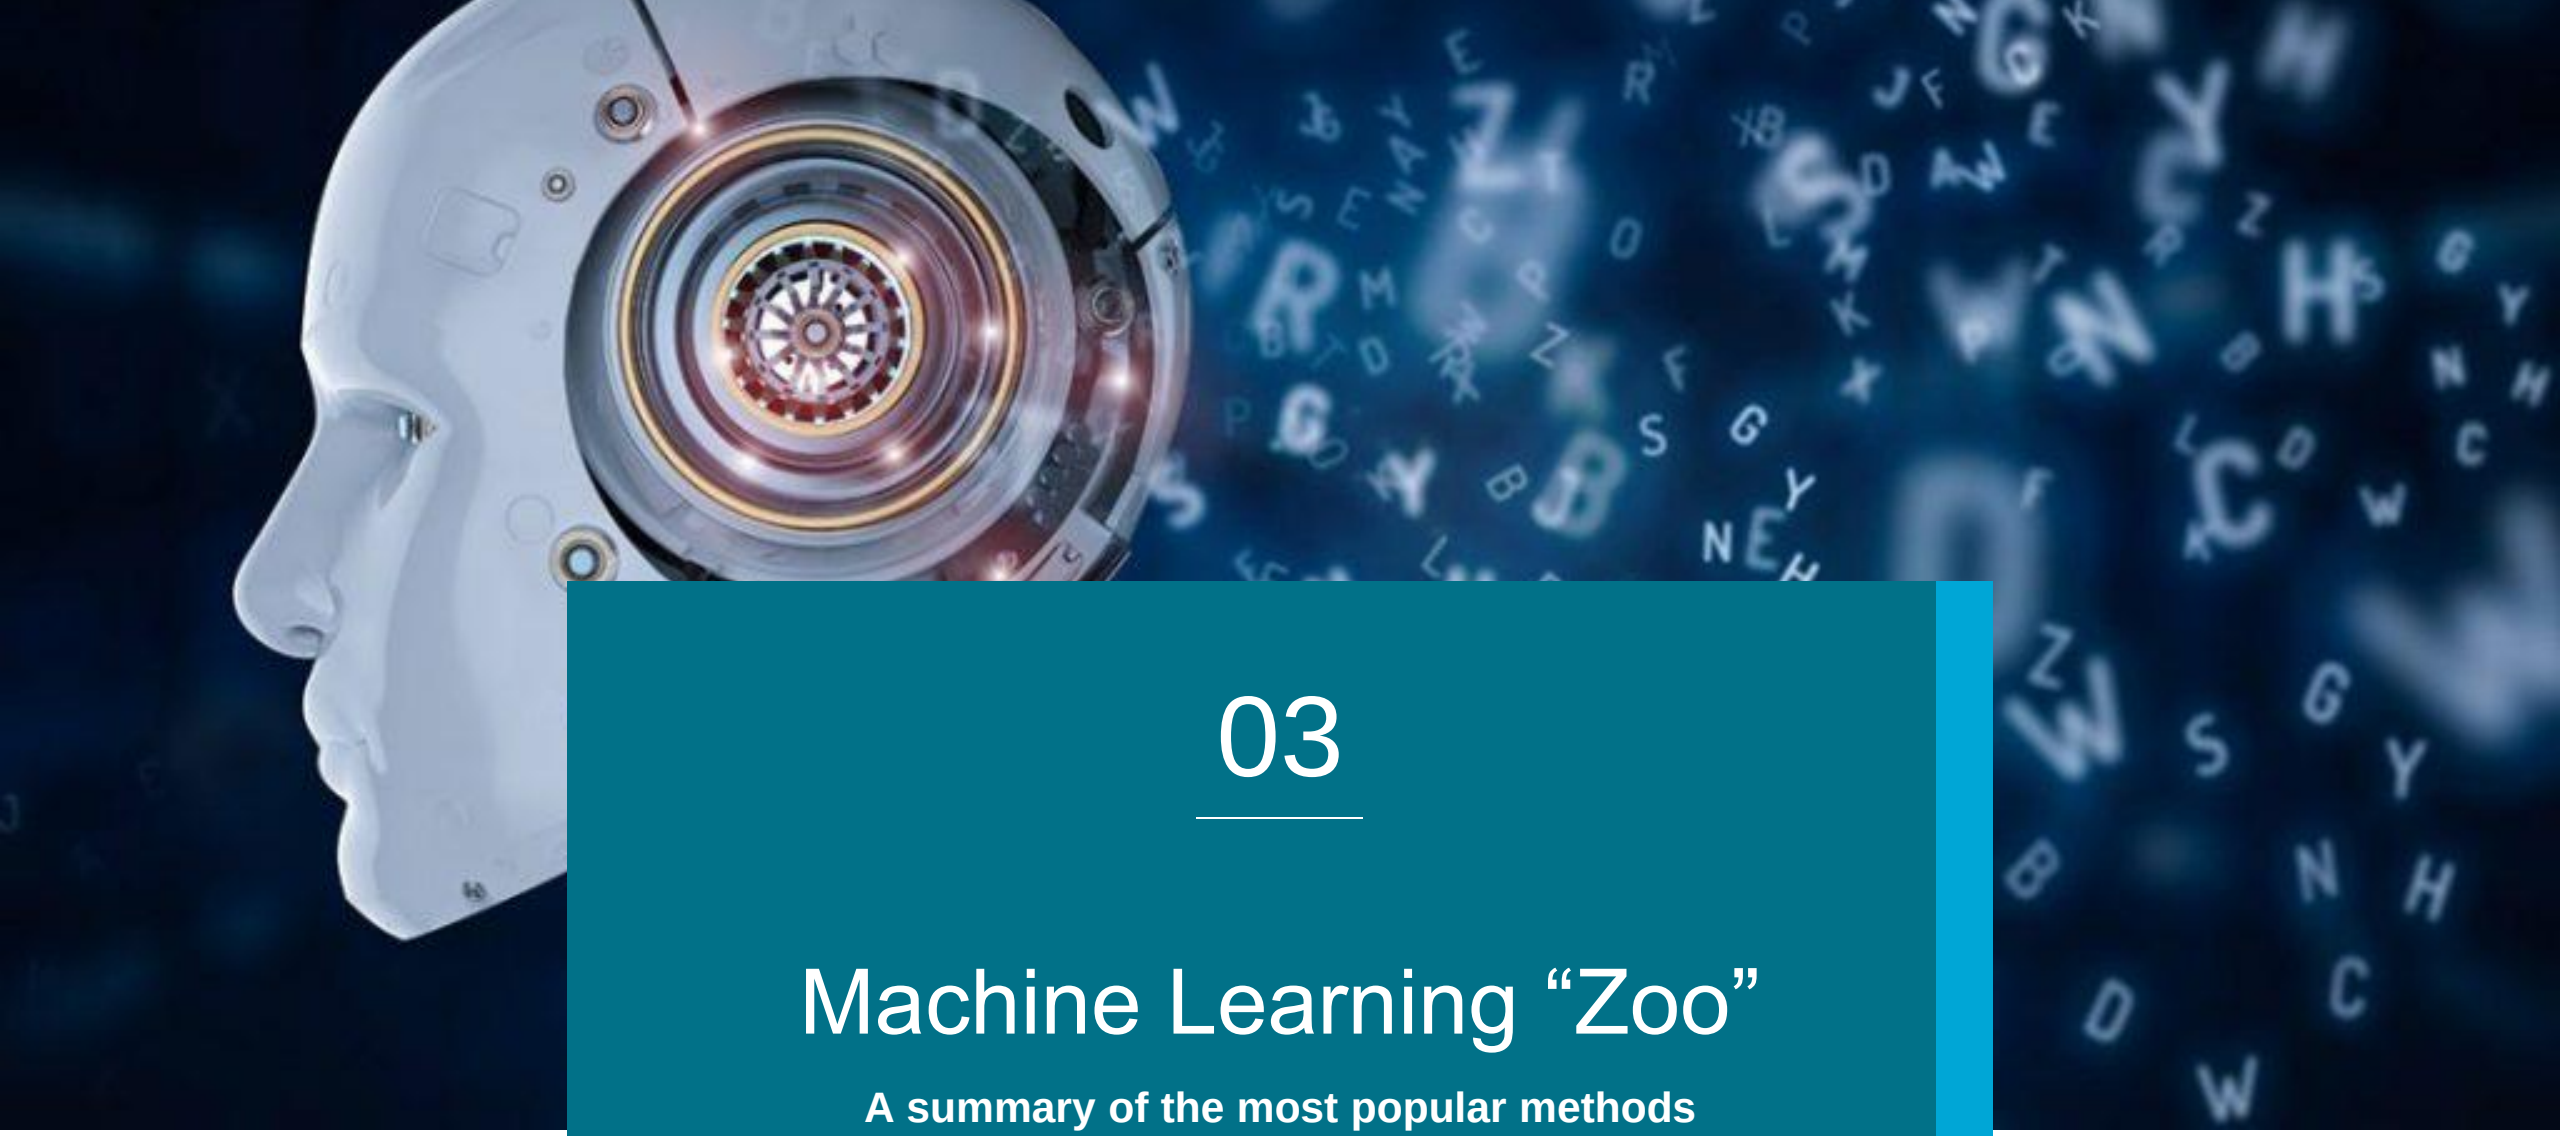

03

Machine Learning “Zoo”

A summary of the most popular methods

TU Delft

What ML models/techniques do you know?

TU Delft

Linear Regression

- **Dataset requirement :** *Supervised*
- **Data provisioning:** *Batch*
- ***Model representation:*** *Model-based :* $Y = \beta X + \varepsilon$
- ***Task:*** *Regression*
  - For *Classification, the equivalent model is Logistic Regression*

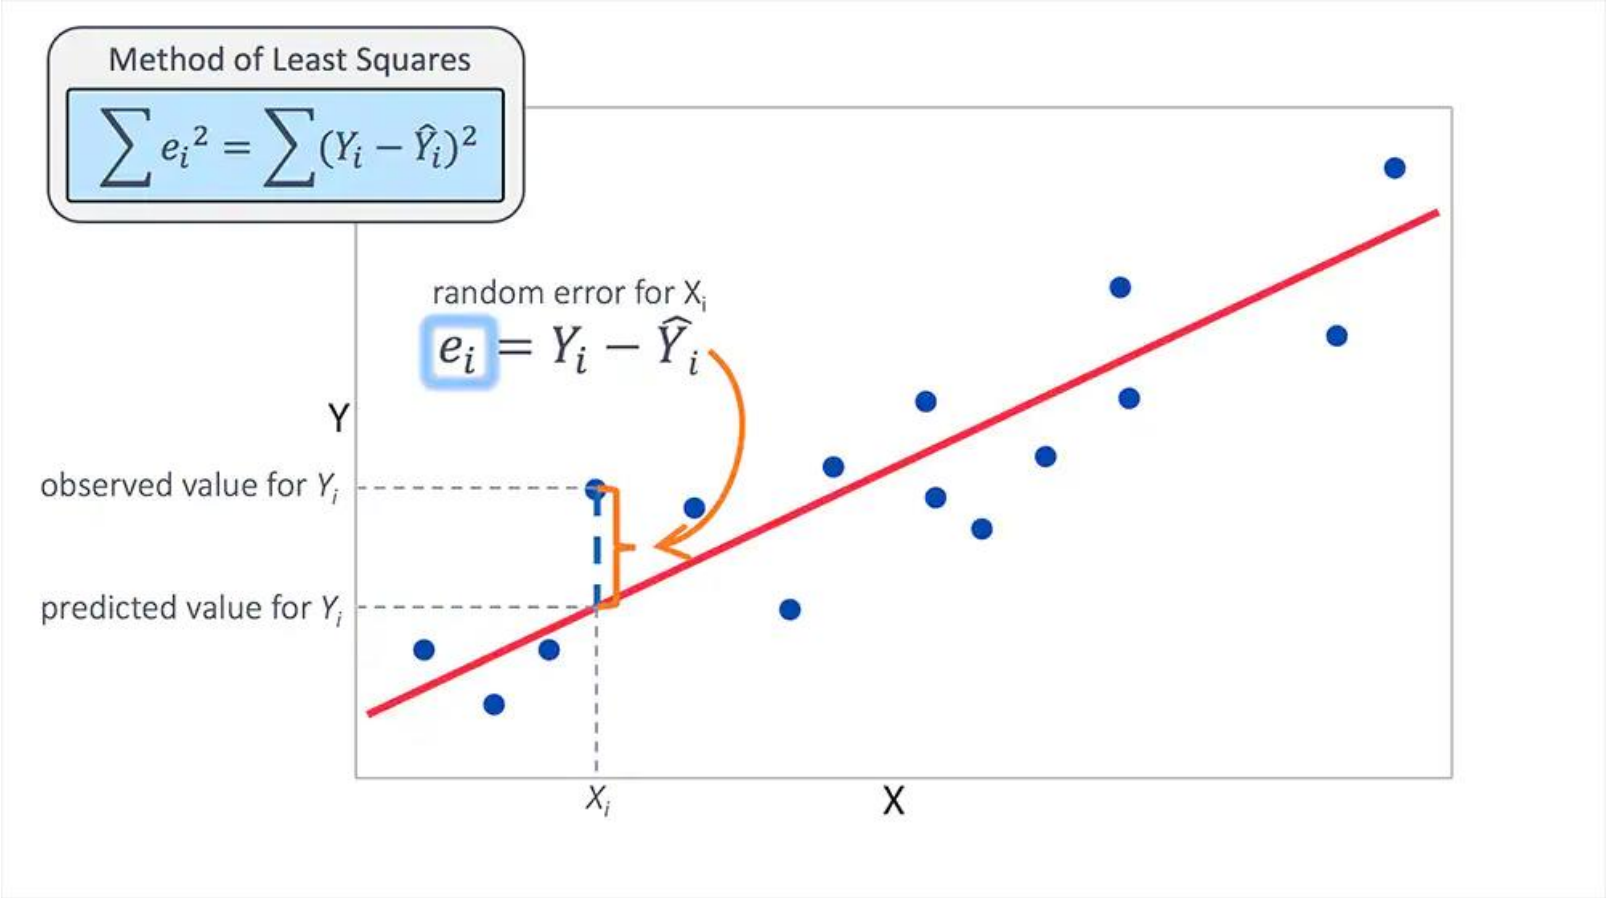

Source: https://heartbeat.comet.ml/understanding-the-mathematics-behind-decision-trees-22d86d55906

Support Vector Machines (SVM)

- **Dataset requirement :** *Supervised*
- **Data provisioning:** *Batch*
- ***Model representation:*** *Model-based :* $Y = K(\beta X) + \varepsilon$
- ***Task:*** *Classification*
  - For *Regression, the equivalent model is Support Vector Regression*

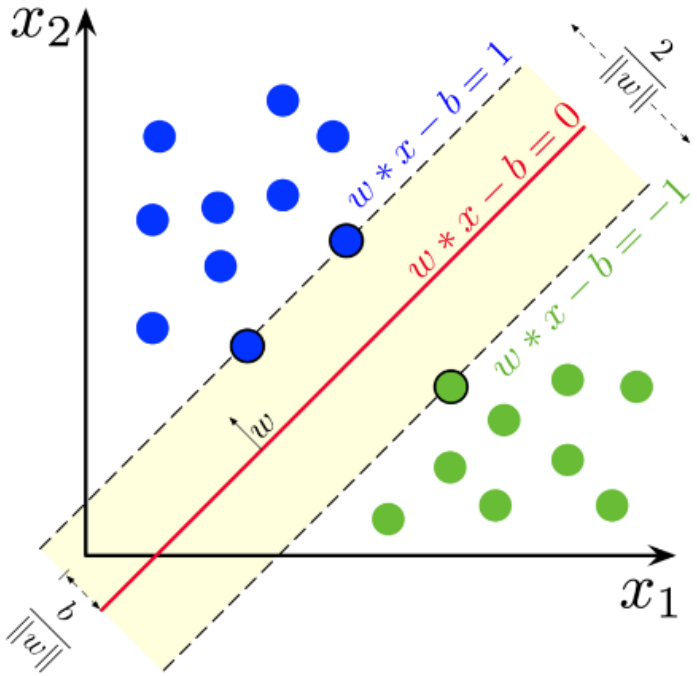

Source: https://en.wikipedia.org/wiki/Support_vector_machine#/media/File:SVM_margin.png

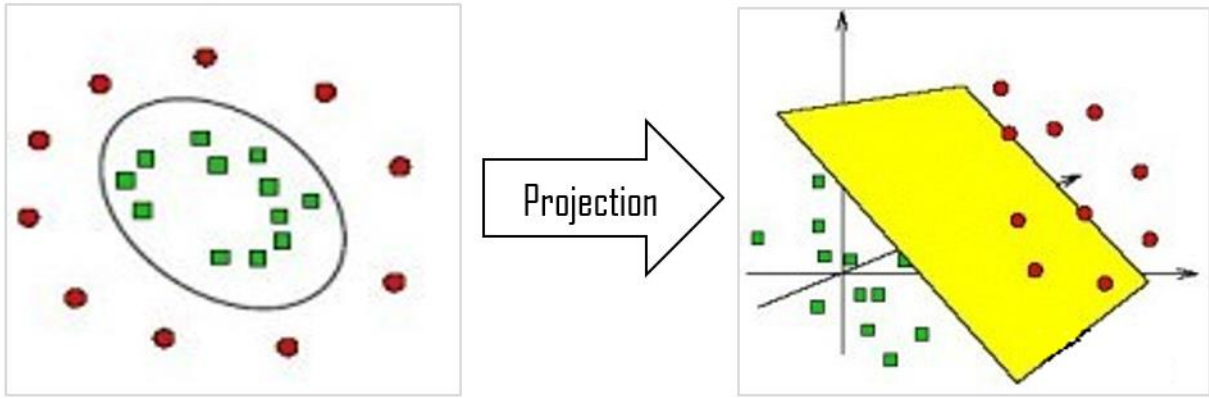

Complex segmentation in low-dimensional space

Easy segmentation in high-dimensional space

Source: https://forum.huawei.com/enterprise/en/machine-learning-algorithms-support-vector-machine-svm/thread/722247-895

Decision Tree

- **Dataset requirement :** *Supervised*
- **Data provisioning:** *Batch*
- ***Model representation:*** *Instance-based*
- ***Task:*** *Regression/Classification*

X

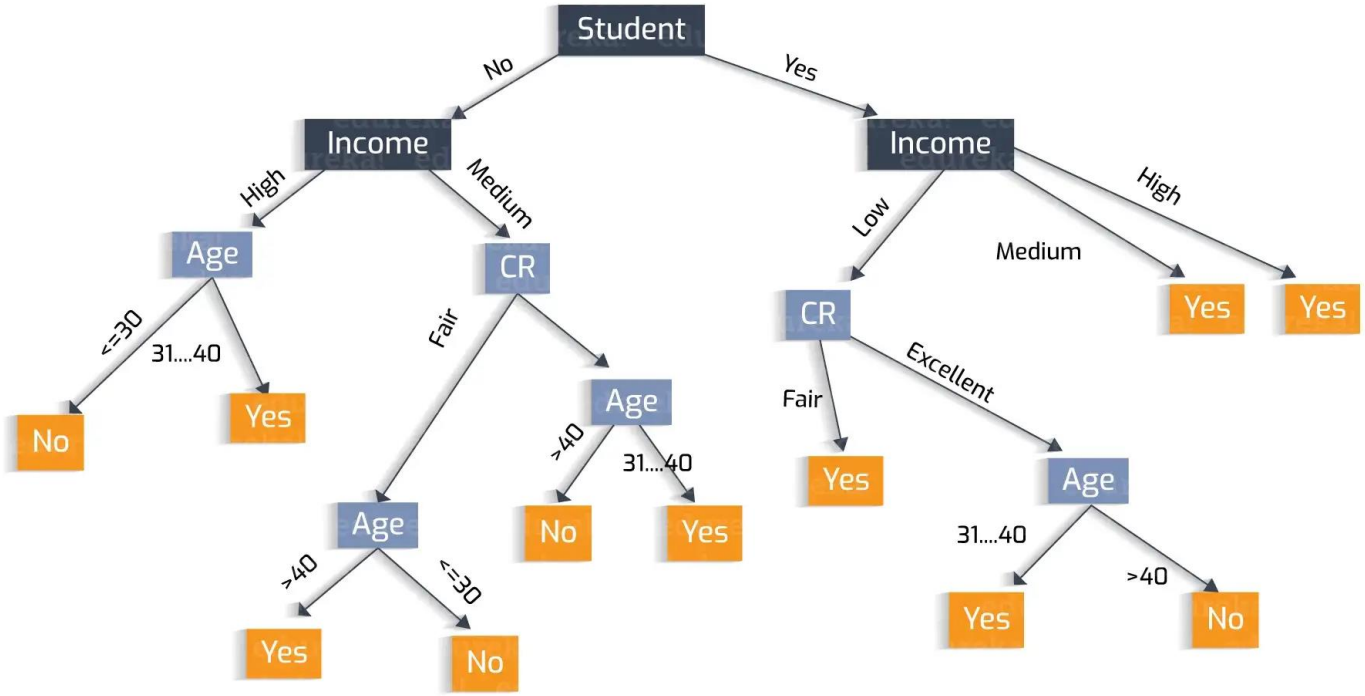

Y

Source: https://heartbeat.comet.ml/understanding-the-mathematics-behind-decision-trees-22d86d55906

Random forest

- **Dataset requirement :** *Supervised*
- **Data provisioning:** *Batch*
- ***Model representation:*** *Instance-based*
- ***Task:*** *Regression/Classification*
- ***Ensemble model***

X

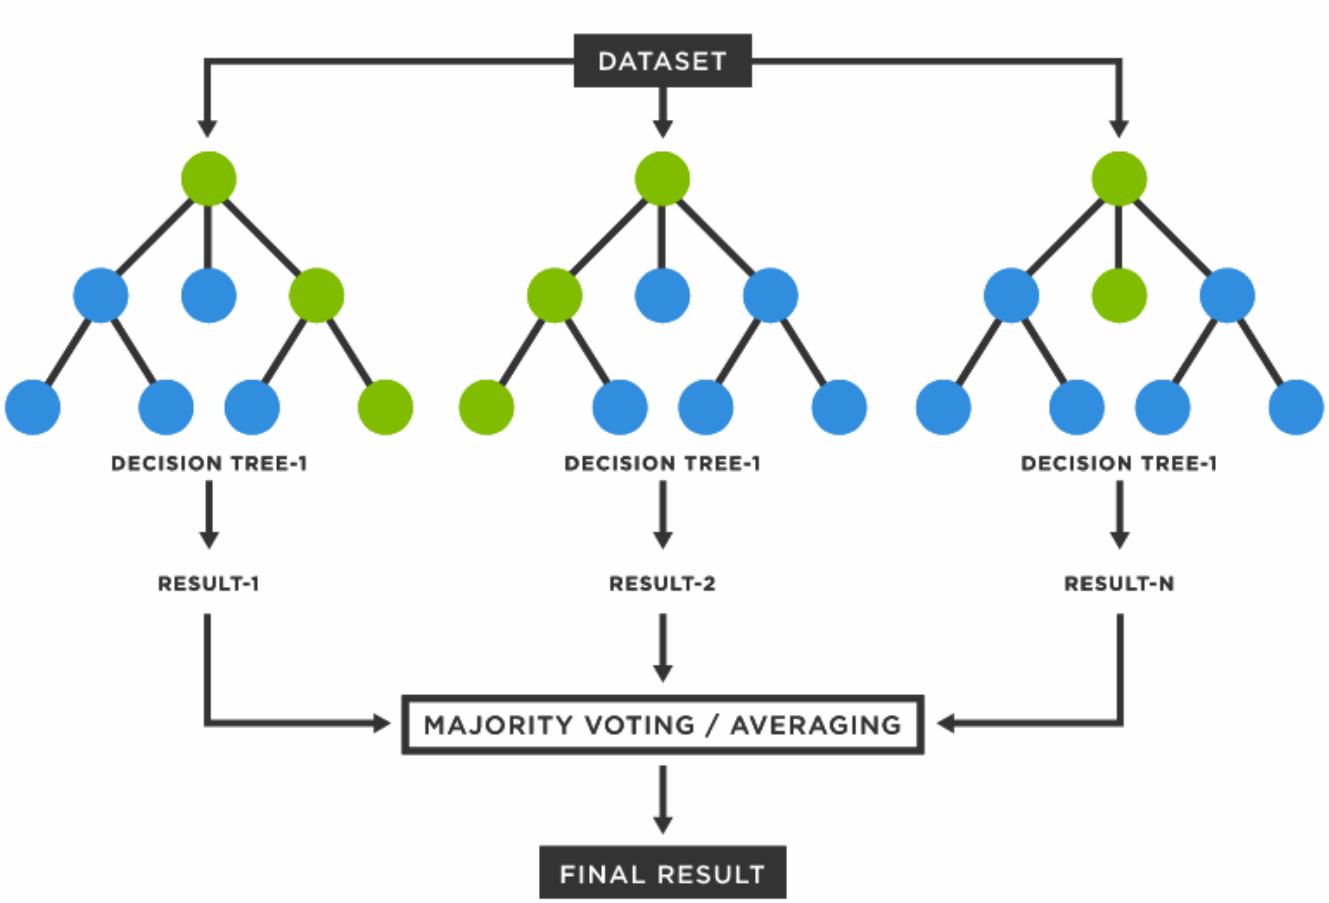

Y

Source: https://www.tibco.com/reference-center/what-is-a-random-forest

Artificial Neural Networks

- **Dataset requirement :**
  - *Supervised (ANN, RNN, CNN, GAN)*
  - *Unsupervised (Autoencoders)*
- **Data provisioning:** *Batch/Online*
- ***Model representation:*** *Model-based*
- ***Task:*** *Regression/Classification*
- ***Ensemble model***

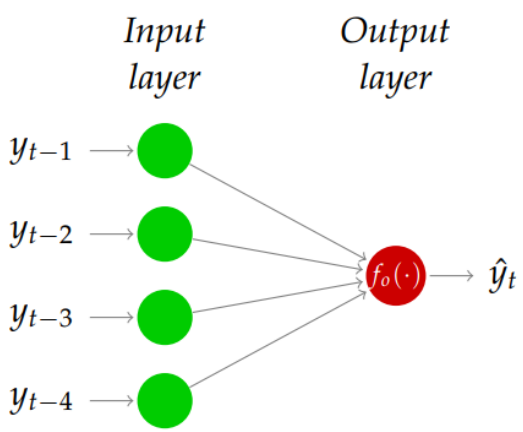

$$\hat{y}_t = f_o\left(b_o + \sum_{i=1}^{|I|} w_{io} y_{t-i}\right) \tag{2.47}$$

(a)
Perceptron
(equivalent to linear regression)

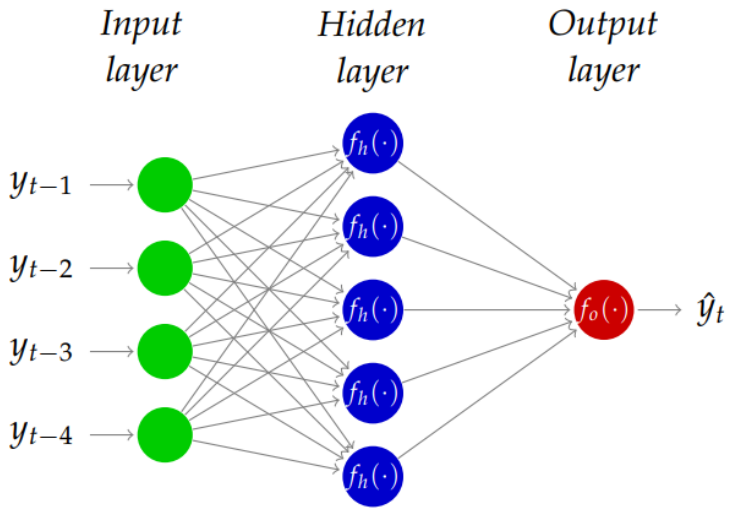

$$\hat{y}_t = f_o\left(b_o + \sum_{j=1}^{|H|} w_{jo} \cdot f_h\left(\sum_{i=1}^{|I|} w_{ij} y_{t-i} + b_j\right)\right) \tag{2.48}$$

(b)
Multi-layer perceptron

Source: [My Own PhD Thesis ☺]

Artificial Neural Networks

- **Dataset requirement :**
  - *Supervised (ANN, RNN, CNN, GAN)*
  - *Unsupervised (Autoencoders)*
- **Data provisioning:** *Batch/Online*
- ***Model representation:*** *Model-based*
- ***Task:*** *Regression/Classification*

TU Delft

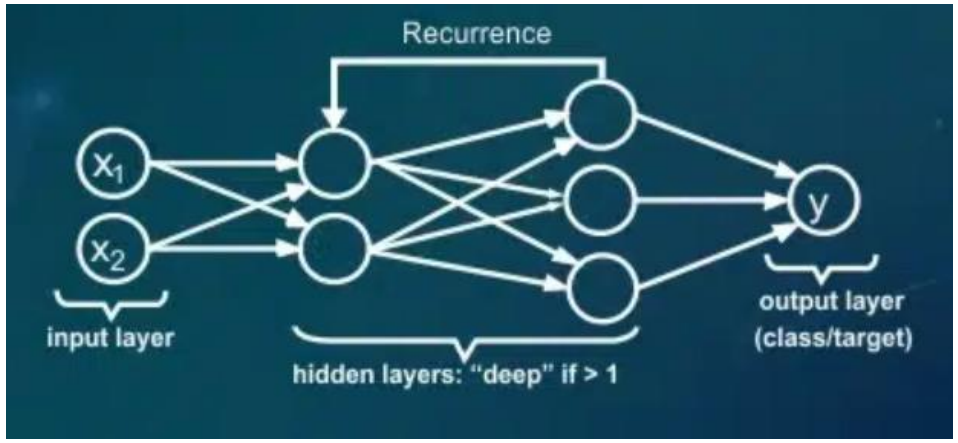

(a)
Recurrent (Deep) Neural Network

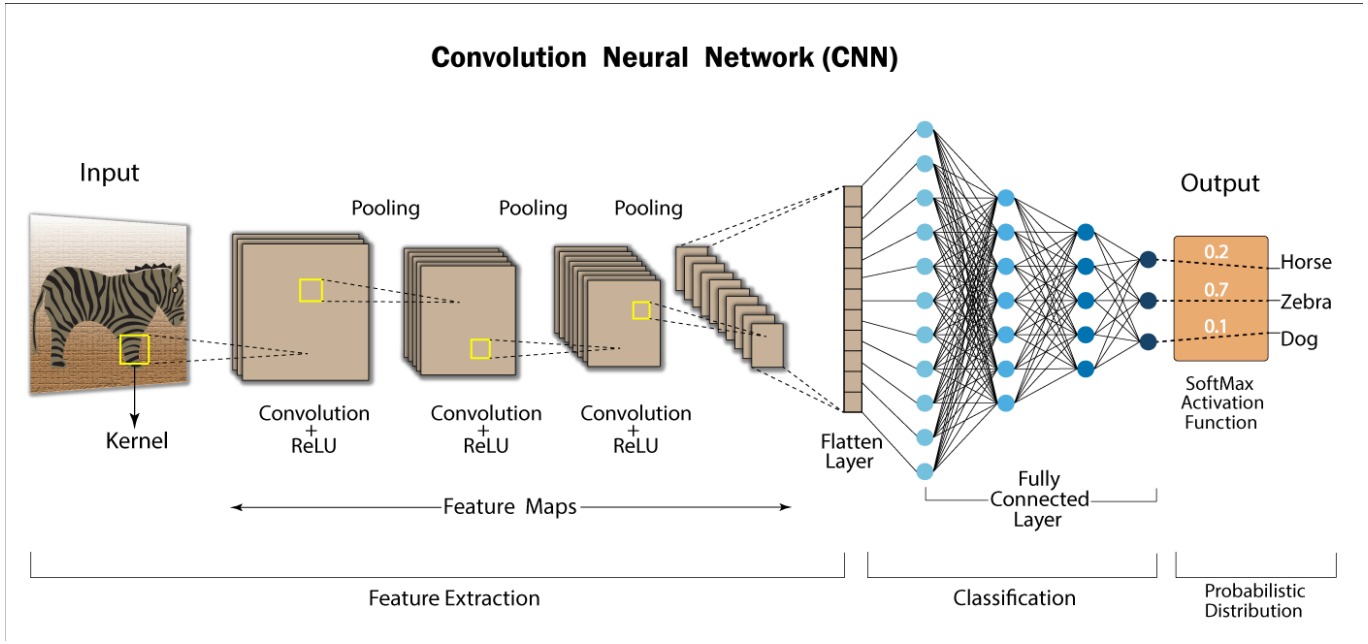

(b)
Convolutional Neural Networks

K-Means Clustering

- **Dataset requirement :** *Unsupervised*
- **Data provisioning:** *Batch*
- ***Model representation:*** *Instance-based*
- ***Task:*** *Clustering/pattern recognition*
- ***N.B.*** *: As clustering is unsupervised, multiple solutions can be found*

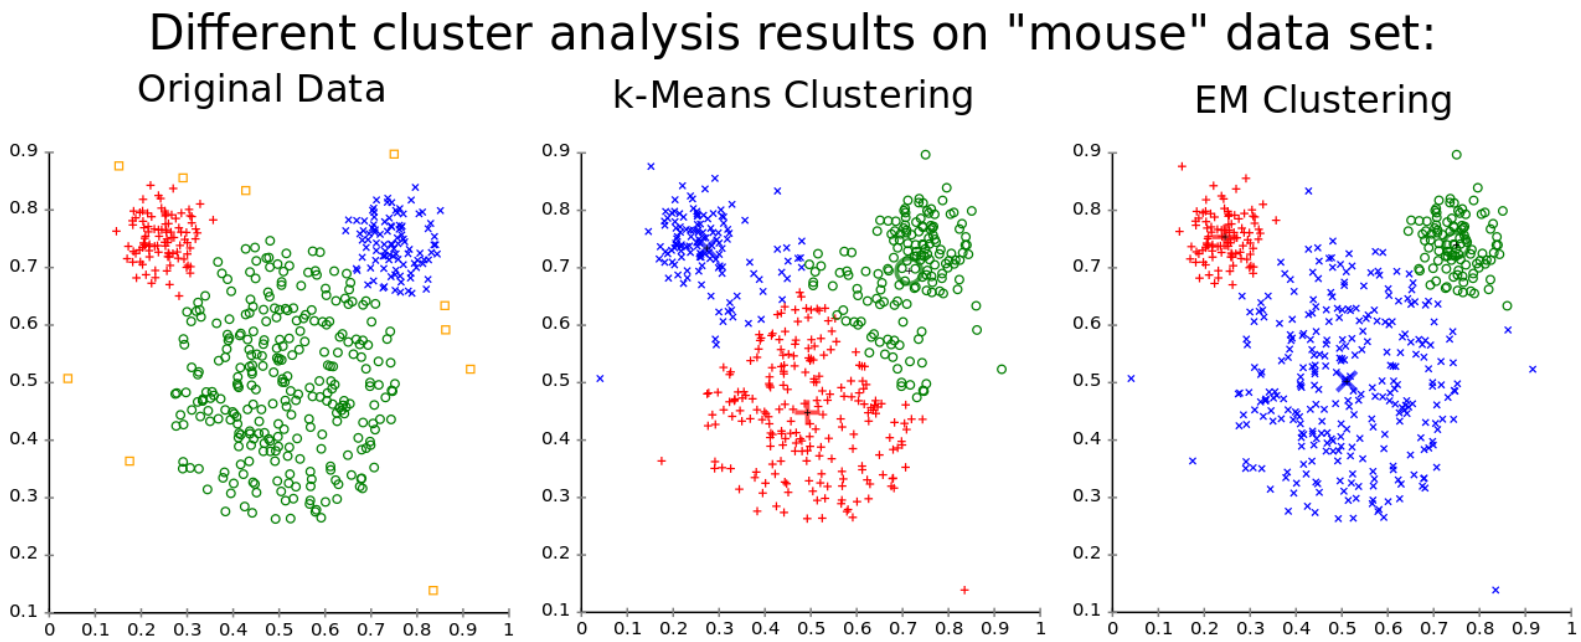

TU Delft

Hierarchical clustering

- **Dataset requirement :** *Unsupervised*
- **Data provisioning:** *Batch*
- ***Model representation:*** *Instance-based*
- ***Task:*** *Clustering/pattern recognition*

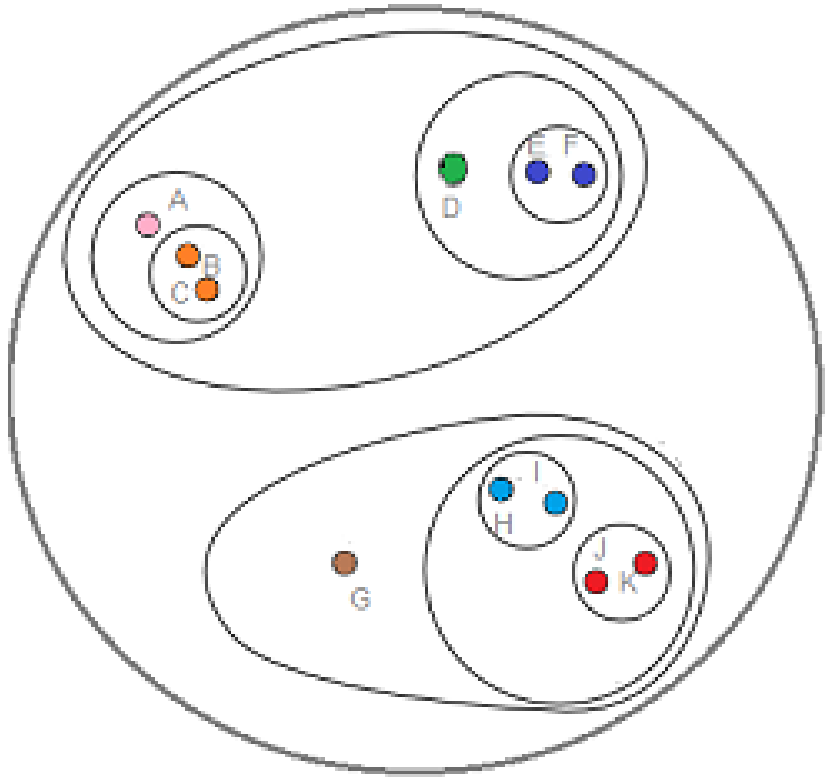

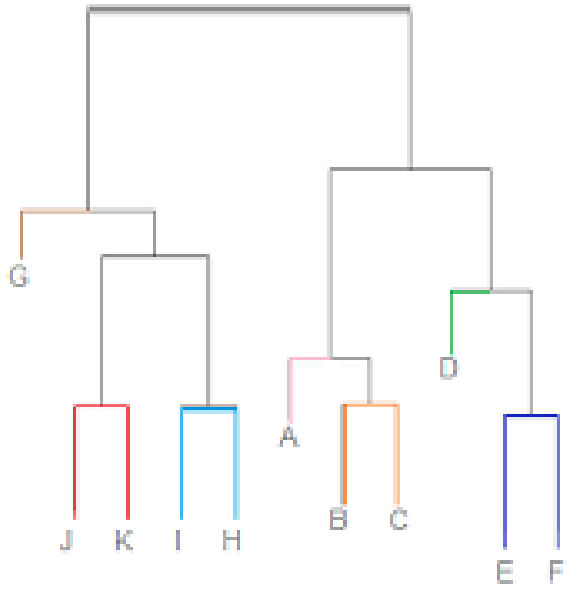

TU Delft

And many more…

- For the sake of time and/or mathematical complexity, I will not talk in detail about the following methods…
- …but feel free to ask questions at the end of the talk!

TU Delft

Dimensionality Reduction
- PCA
- t-SNE
- Autoencoders

Clustering
- DBSCAN
- Self-organizing maps

Regression
- Gradient Boosting
- K-NN

Classification
- Gradient Boosting
- K-NN

Reinforcement Learning
- Q-Learning
- Deep Q-Learning

…

TU Delft

03

AI as a socio-technical technology

From an academic perspective

TU Delft

What we have learned so far?

- AI is concerned with “intelligent-like” behaviour
- ML is a subfield of AI focusing on learning from data to improve over time performances on a task
- Several different ML models/learning algorithms exist to fulfil different purposes
- A learning algorithm is still inherently an algorithm (i.e. a sequence of instructions) with no pre-conceived notion of good or bad…
- … but lives in a socio-technical domain

TU Delft

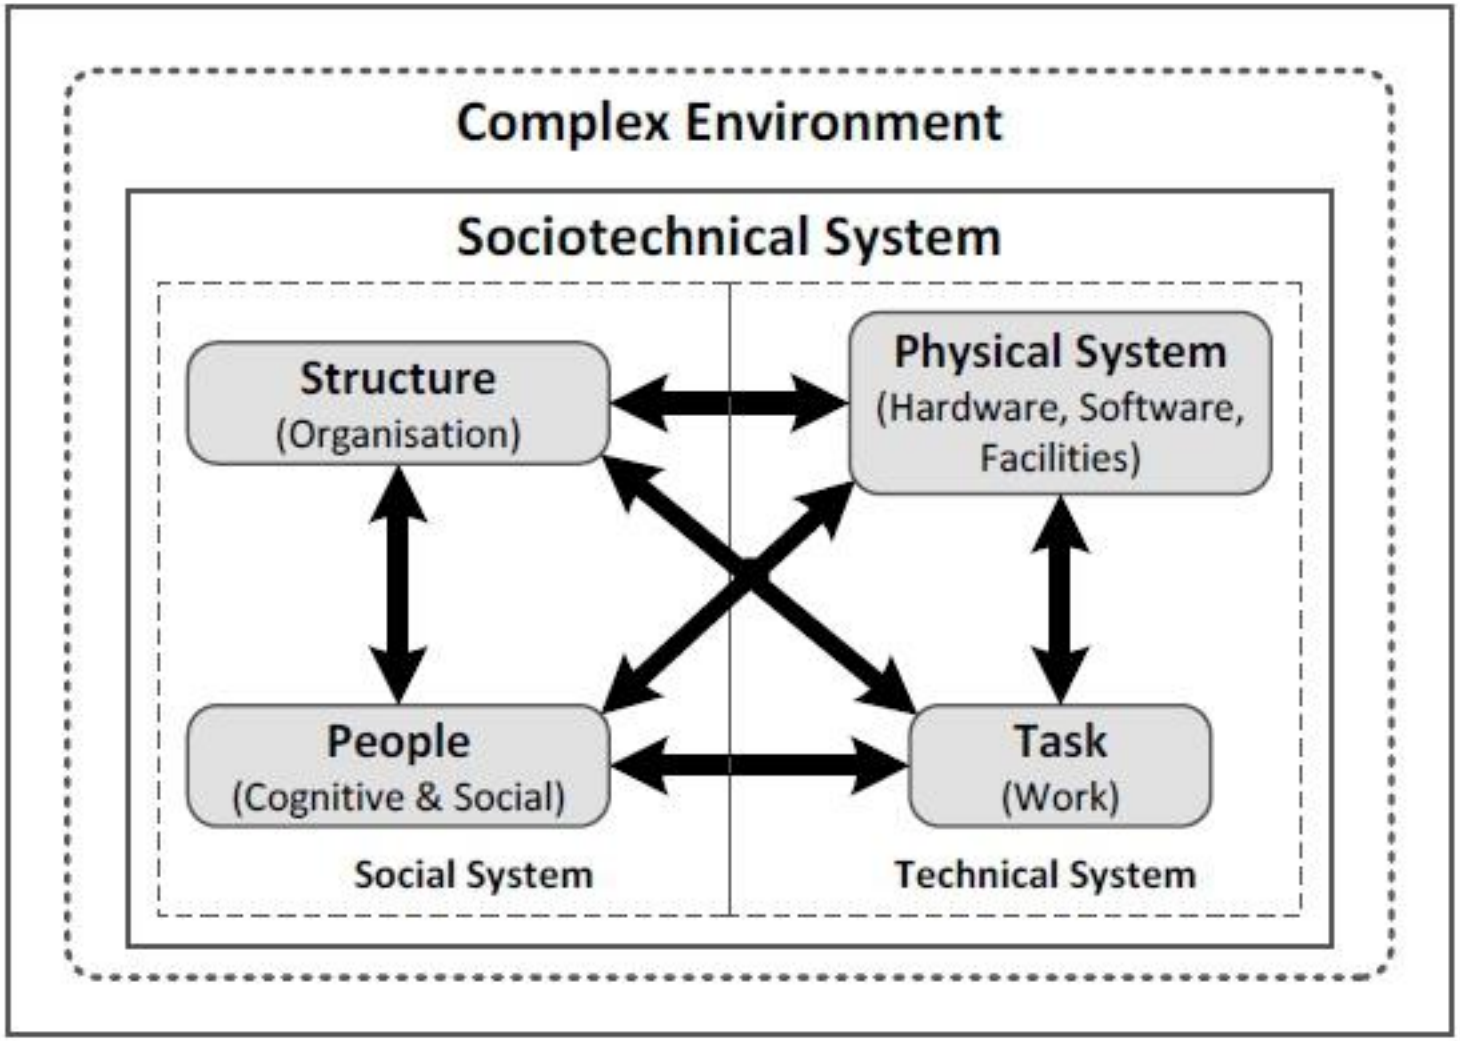

Figure 2: Sociotechnical system (STS) [4]

Source: Oosthuizen, R., & Pretorius, L.. (2016). Assessing the impact of new technology on complex sociotechnical systems. *South African Journal of Industrial Engineering*, *27*(2), 15-29. https://dx.doi.org/10.7166/27-2-1144

How should be AI applied in practice?

From the perspective of the data:

1. Descriptive Analysis
2. Diagnostic Analysis
3. Predictive Analysis
4. Prescriptive Analysis

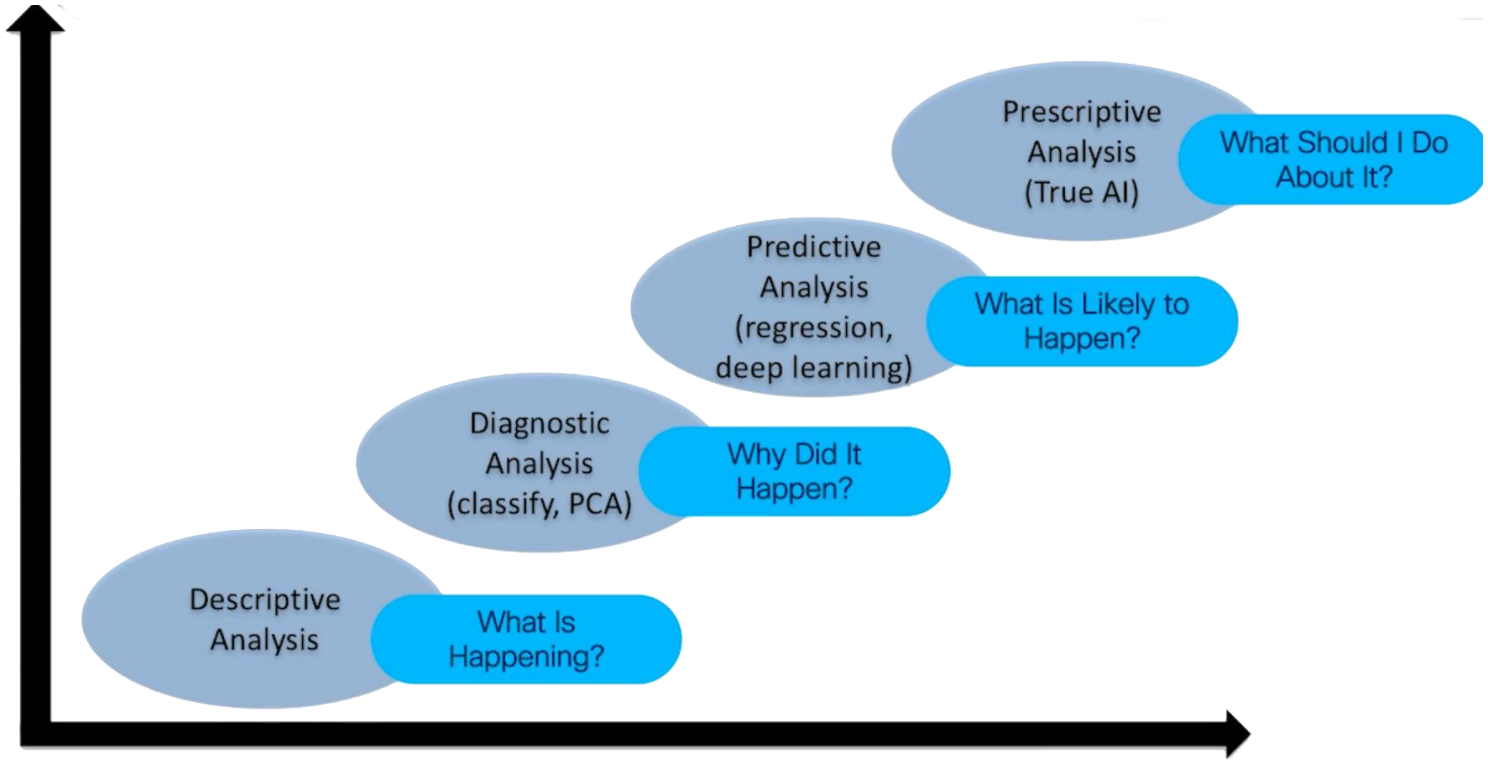

Source picture : Screenshot from Data Analytics and Machine Learning Fundamentals LiveLessons Video Training by Jerome Henry

TU Delft

How should be AI applied in practice?

From the perspective of the model:

- Purely technical aspects
- Socio-technical aspects
- Overarching guiding principles

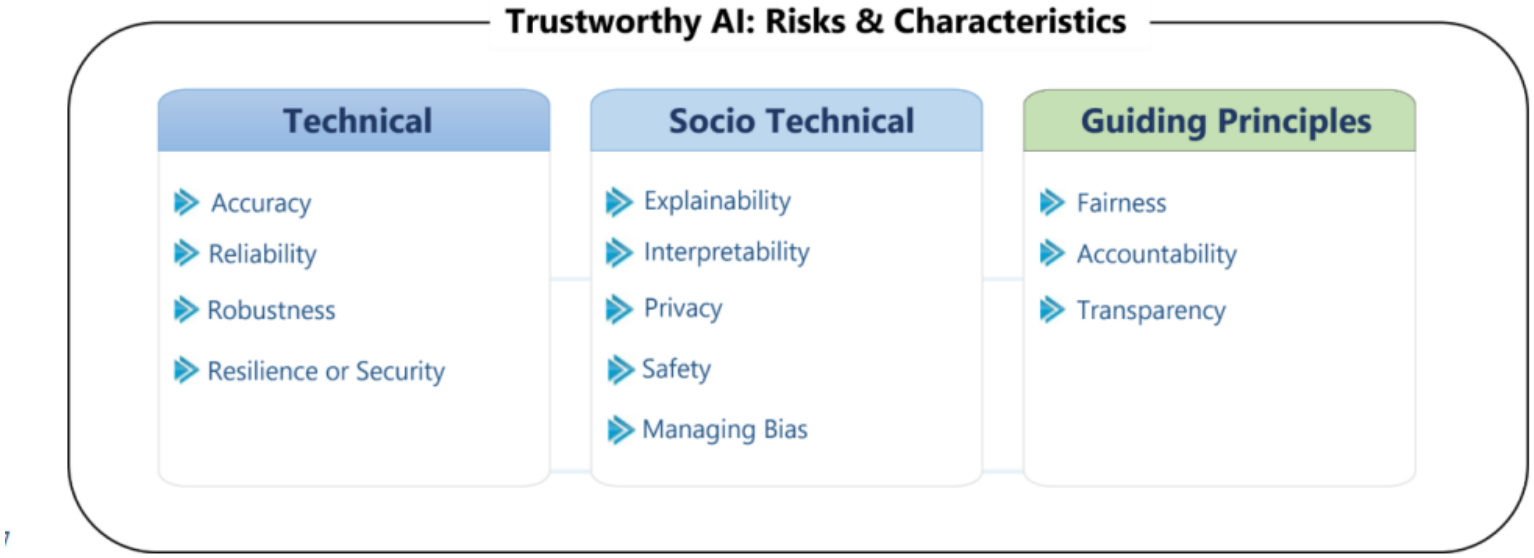

Source: https://www.equalai.org/blog/2022/08/10/nist-will-cultivate-trust-in-ai-by-developing-a-framework-for-ai-risk-management/

TU Delft

How should be AI applied in practice?

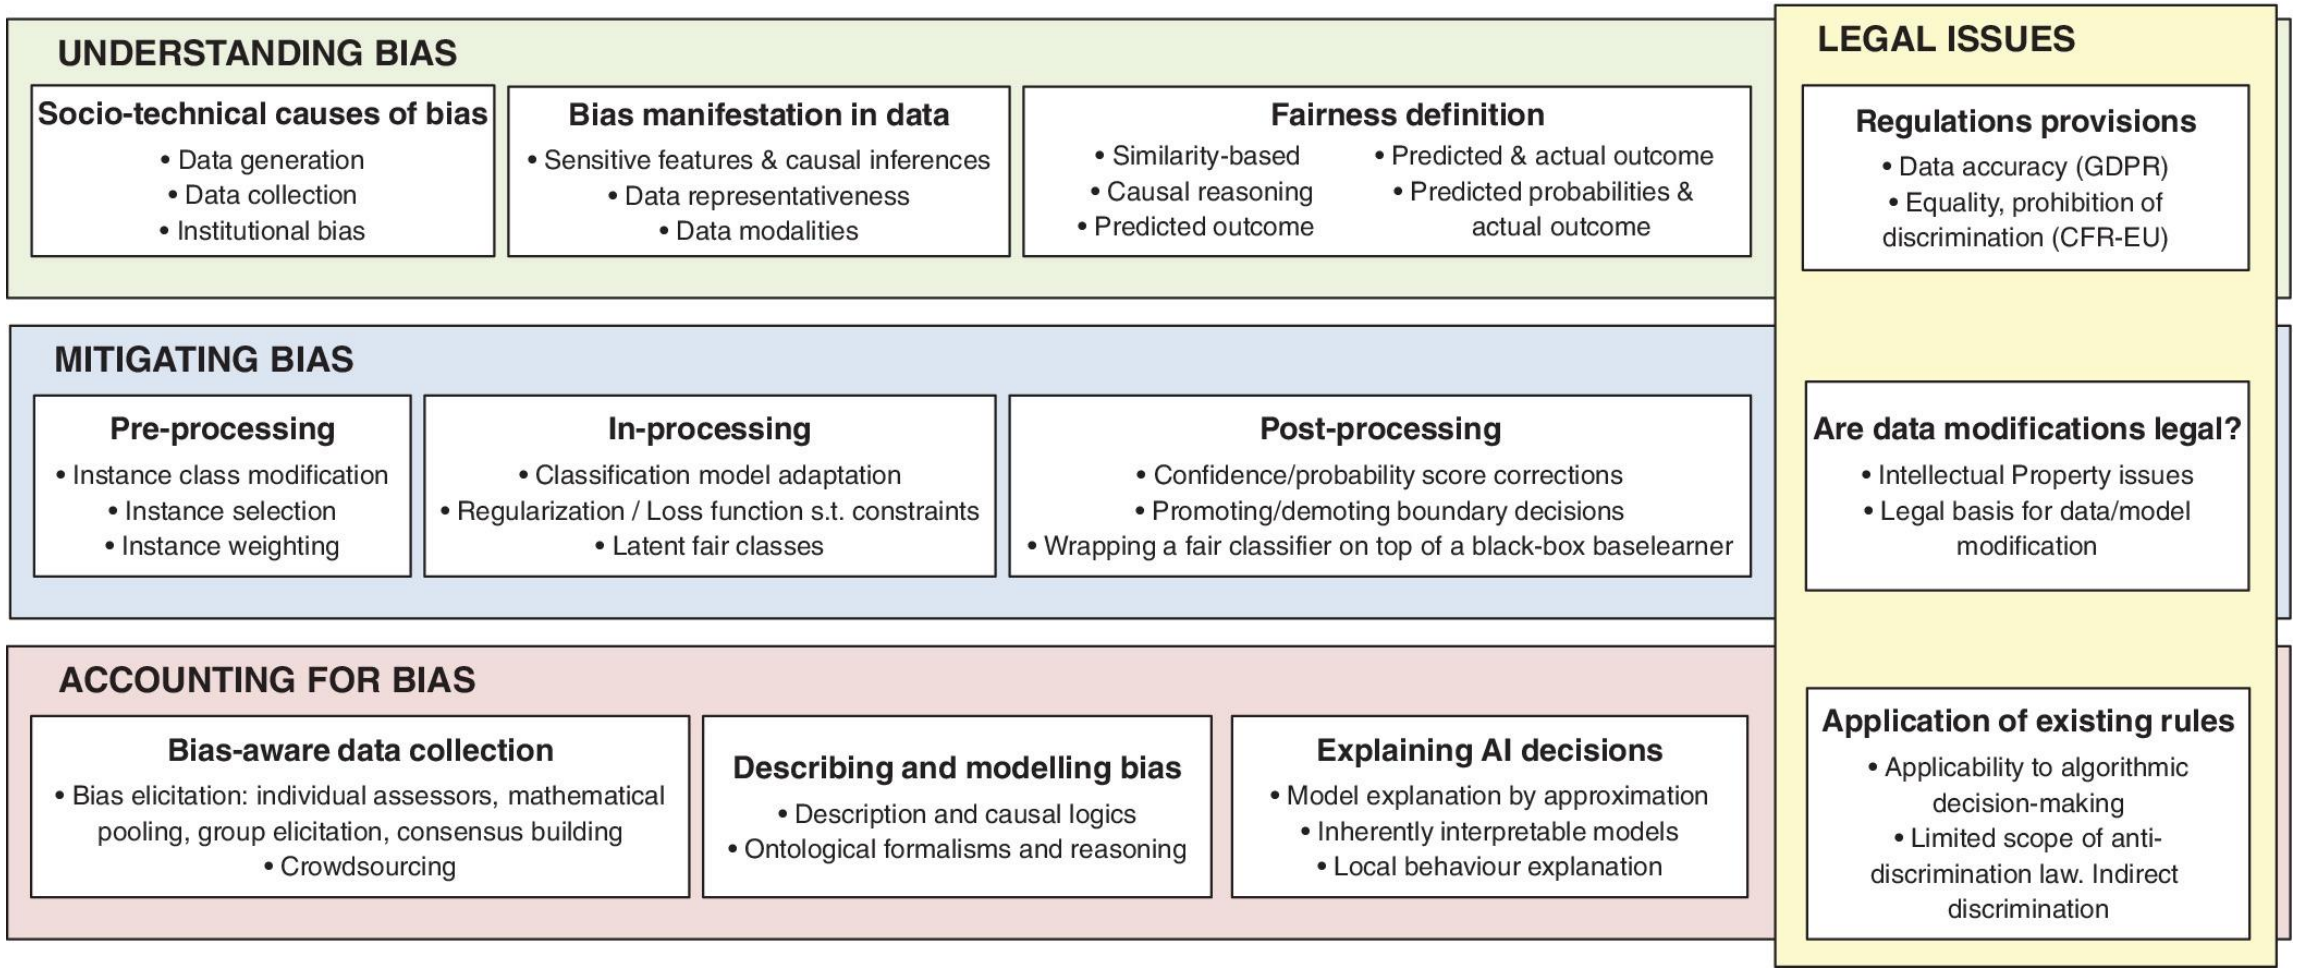

TU Delft

Source: Ntoutsi, E, Fafalios, P, Gadiraju, U, et al. Bias in data-driven artificial intelligence systems—An introductory survey. *WIREs Data Mining Knowl Discov*. 2020; 10:e1356. https://doi.org/10.1002/widm.1356

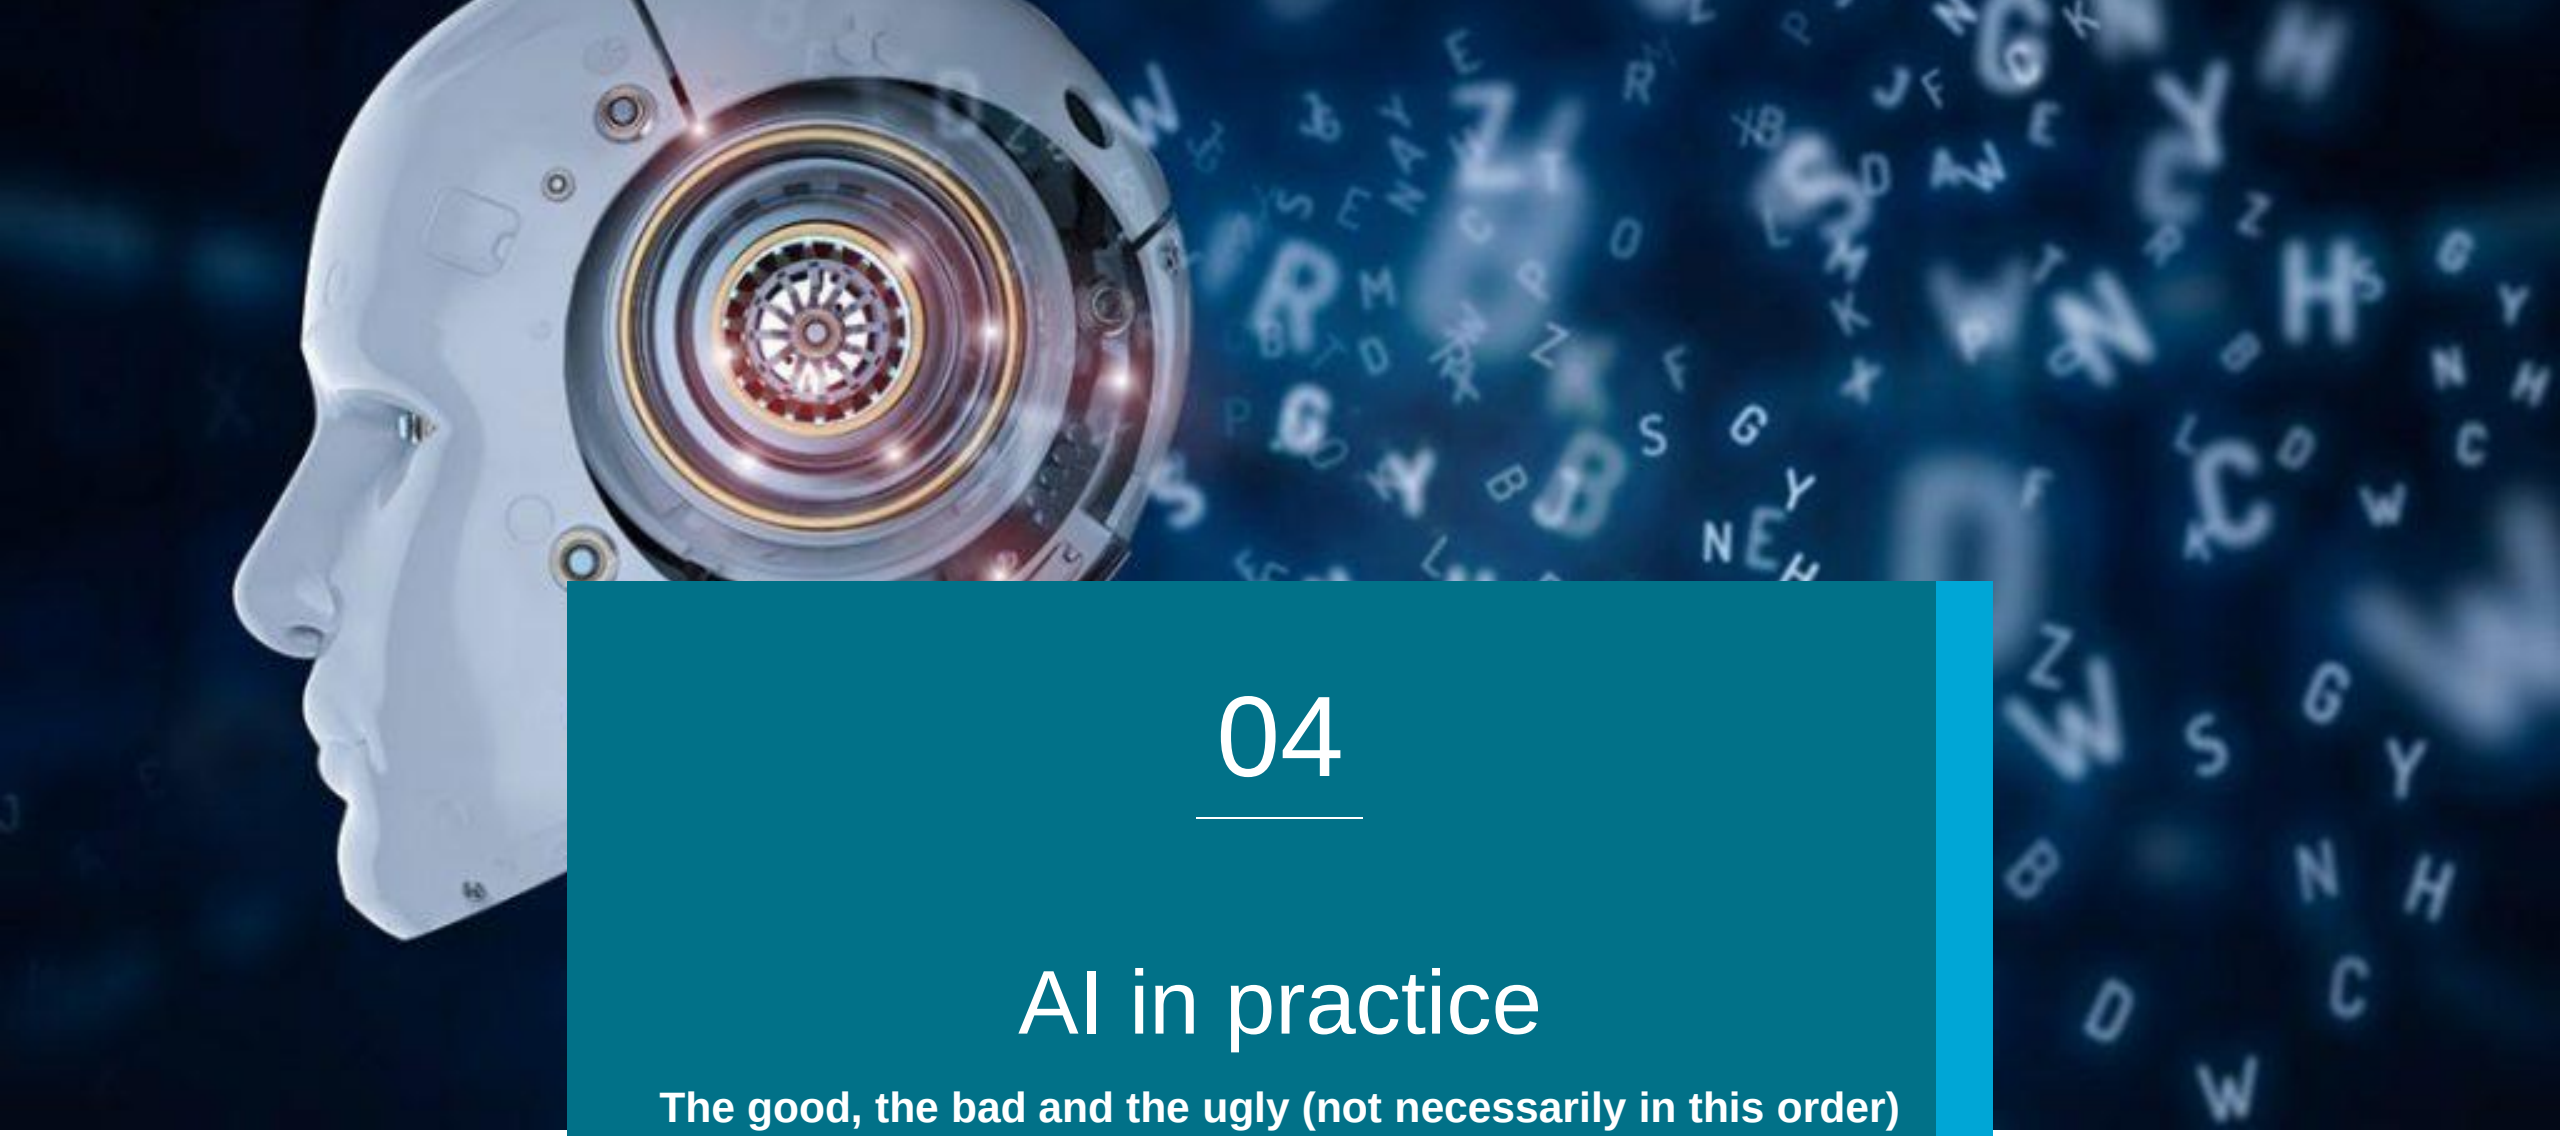

04

AI in practice

The good, the bad and the ugly (not necessarily in this order)

TU Delft

09-10-2020

How AI is actually applied in practice?

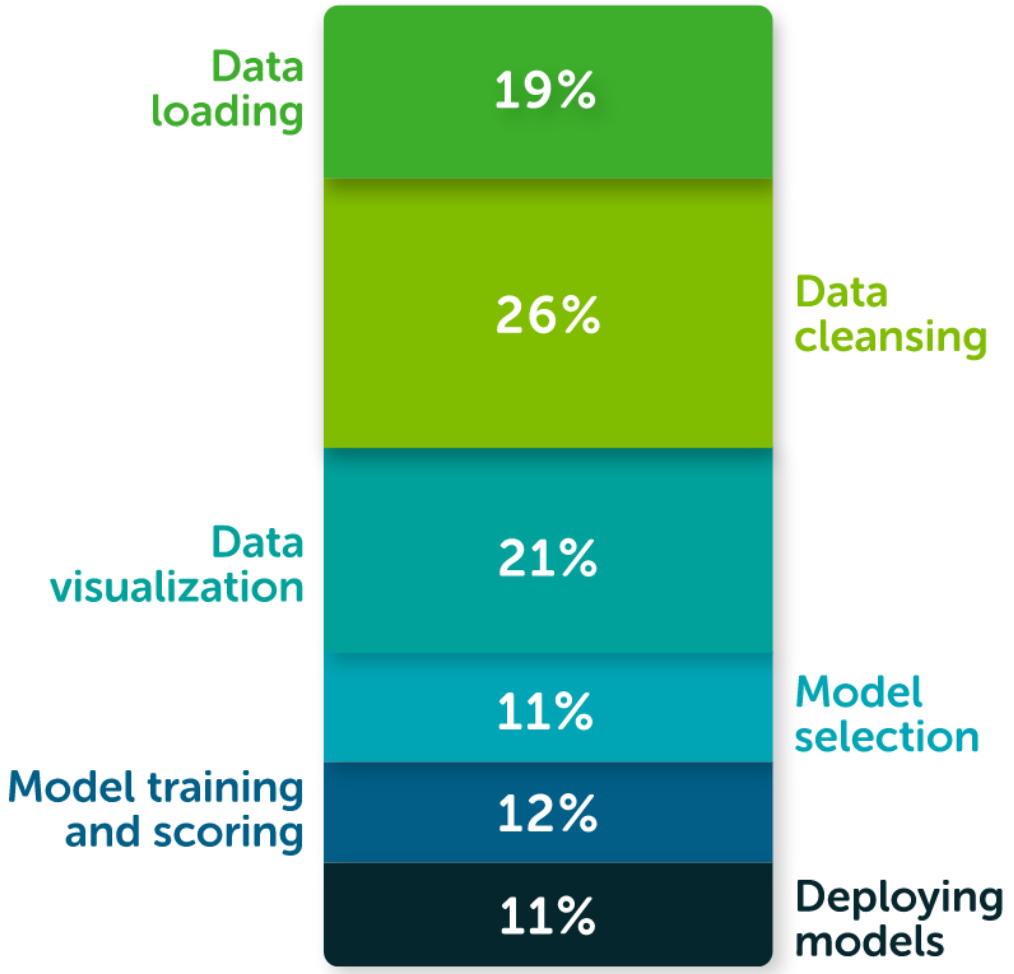

Source: https://www.anaconda.com/state-of-data-science-2020

TU Delft

How AI is actually applied in practice?

Source: Algortihmia report
https://info.algorithmia.com/hubfs/2020/Reports/2021-Trends-in-ML/Algorithmia_2021_enterprise_ML_trends.pdf?hsLang=en-us

TU Delft

Customer experience and process automation represent the top AI/ML use cases

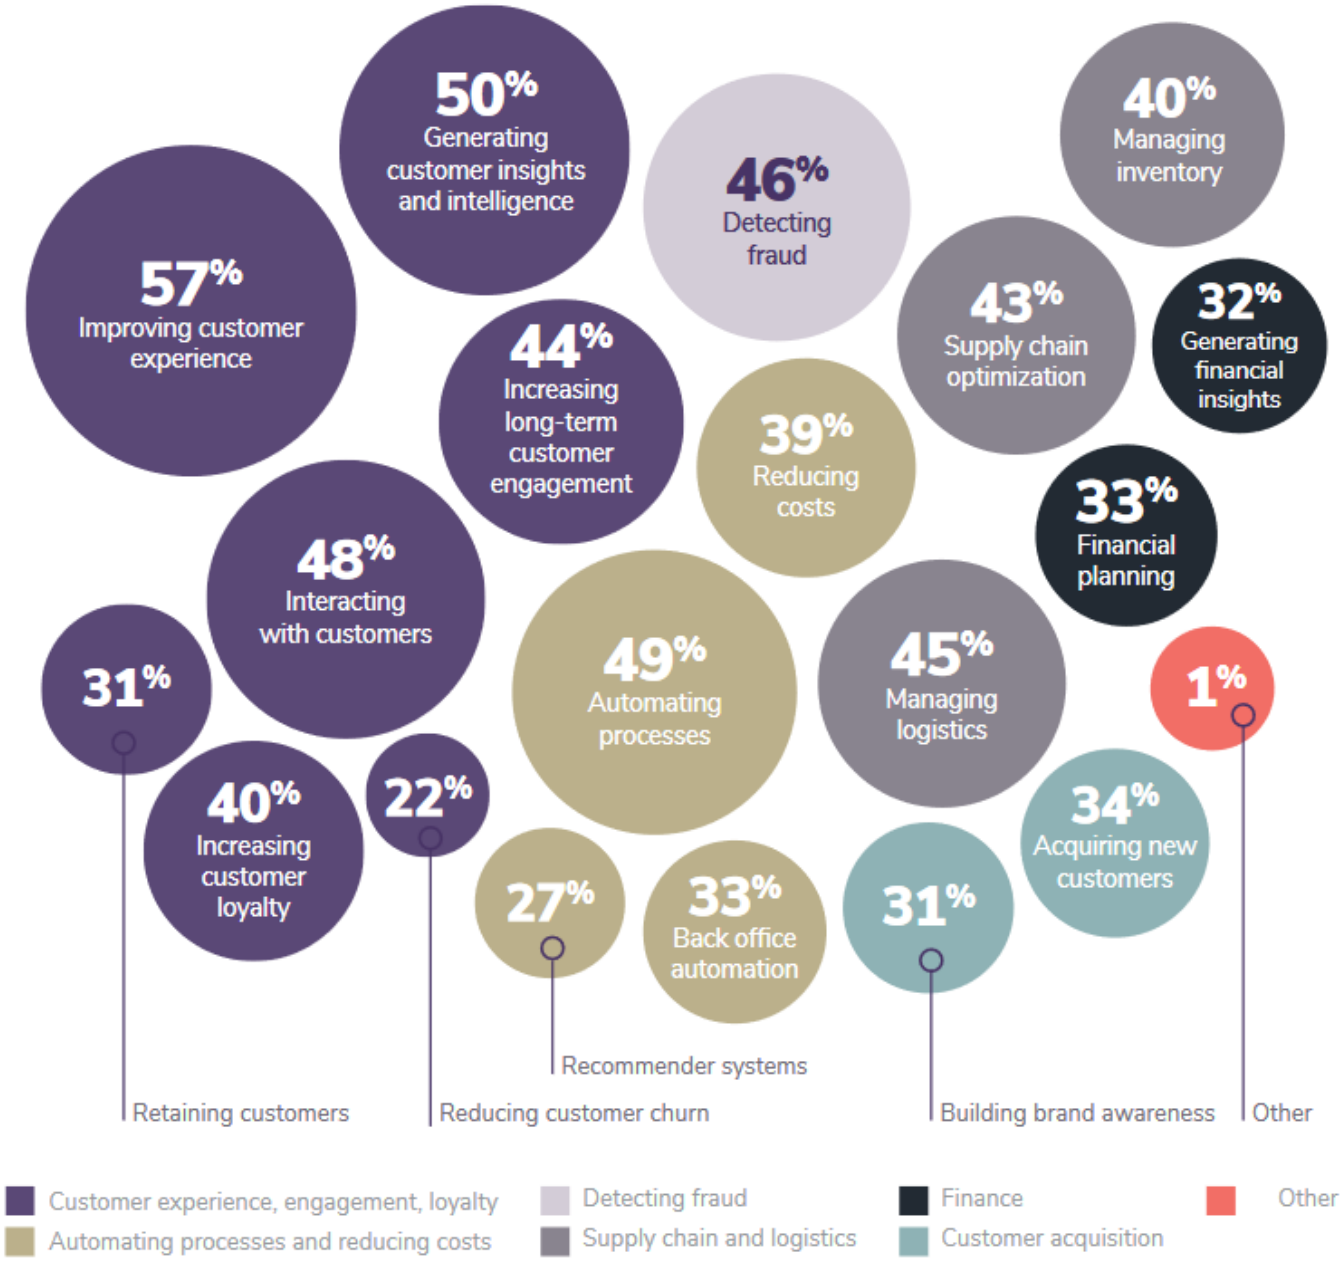

How AI is actually applied in practice?

Source: Algortihmia report https://info.algorithmia.com/hubfs/2020/Reports/2021-Trends-in-ML/Algorithmia_2021_enterprise_ML_trends.pdf?hsLang=en-us

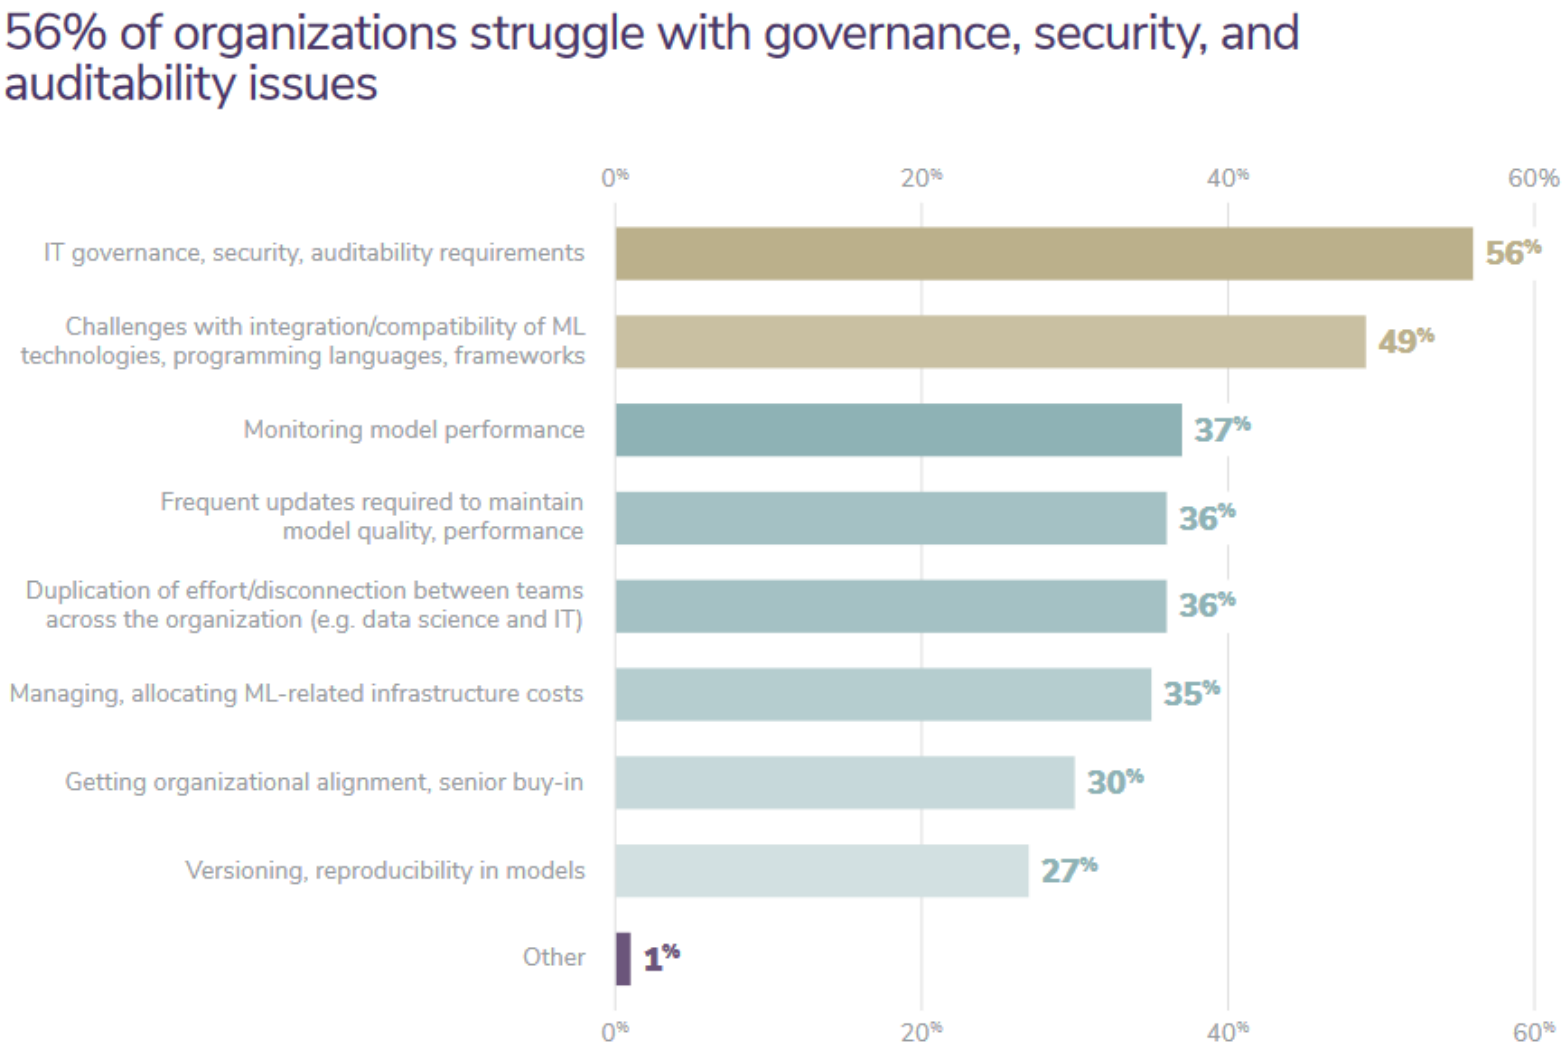

TU Delft

How AI is actually applied in practice?

Source: Algortihmia report https://info.algorithmia.com/hubfs/2020/Reports/2021-Trends-in-ML/Algorithmia_2021_enterprise_ML_trends.pdf?hsLang=en-us

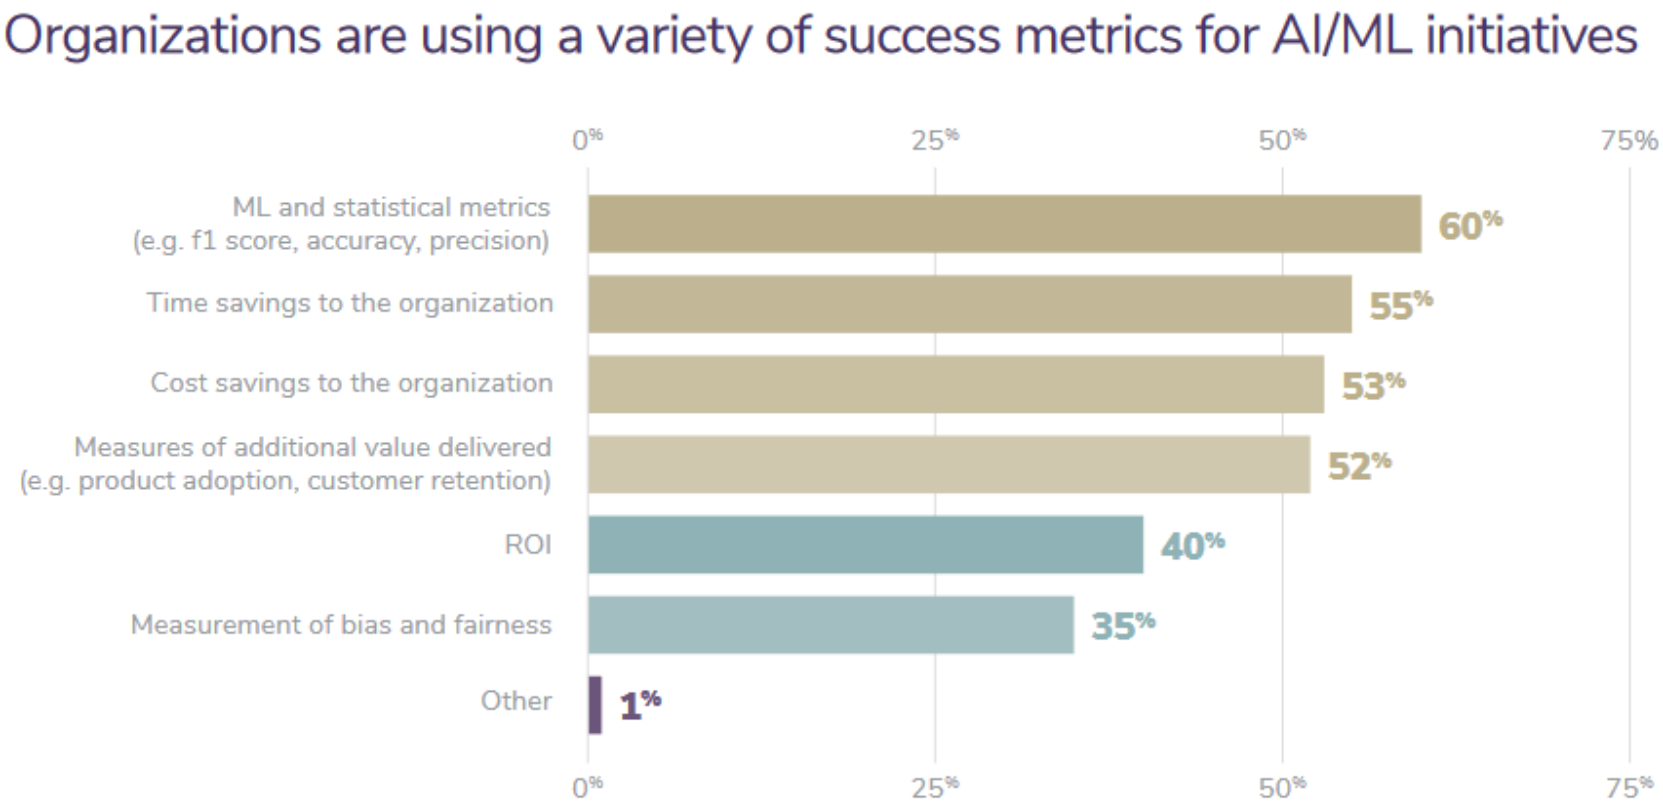

TU Delft

What happened in practice?

- Microsoft case (2016)
- Chat-bot AI trained on available online data

Microsoft 'deeply sorry' for racist and sexist tweets by AI chatbot

Company finally apologises after 'Tay' quickly learned to produce offensive posts, forcing the tech giant to shut it down after just 16 hours

Microsoft's artificial intelligence chatbot Tay didn't last long on Twitter. Photograph: Twitter

Microsoft has said it is "deeply sorry" for the racist and sexist Twitter messages generated by the so-called chatbot it launched this week.

The company released an official apology after the artificial intelligence program went on an embarrassing tirade, likening feminism to cancer and suggesting the Holocaust did not happen.

What happened in practice?

- Amazon case (2018)
- CV-Screening AI trained on available historical data

That is because Amazon's computer models were trained to vet applicants by observing patterns in résumés submitted to the company over a 10-year period. Most came from men, a reflection of male dominance across the tech industry.

In effect, Amazon's system taught itself that male candidates were preferable. It penalized résumés that included the word "women's", as in "women's chess club captain". And it downgraded graduates of two all-women's colleges, according to people familiar with the matter.

TU Delft

Amazon ditched AI recruiting tool that favored men for technical jobs

Specialists had been building computer programs since 2014 to review résumés in an effort to automate the search process

Amazon's automated hiring tool was found to be inadequate after penalizing the résumés of female candidates. Photograph: Brian Snyder/Reuters

Amazon's machine-learning specialists uncovered a big problem: their new recruiting engine did not like women.

The team had been building computer programs since 2014 to review job applicants' résumés, with the aim of mechanizing the search for top talent, five people familiar with the effort told Reuters.

What happened in practice?

- Amazon case (2018)
- CV-Screening AI trained on available historical data

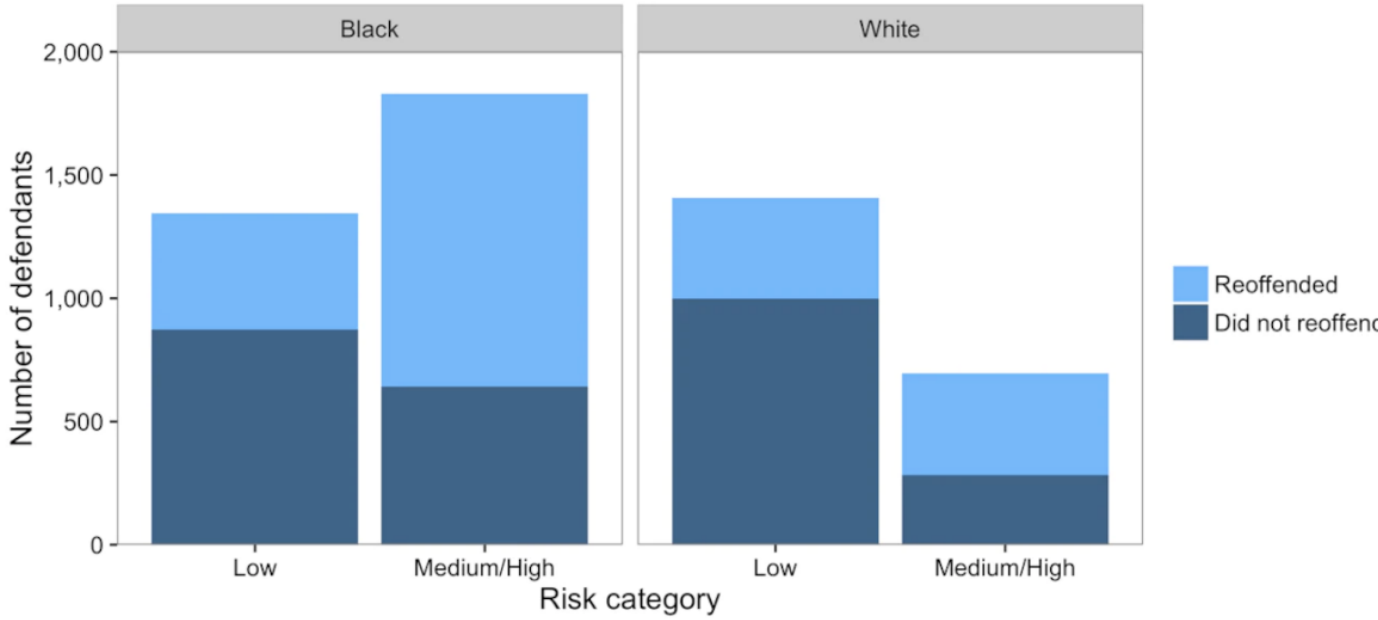

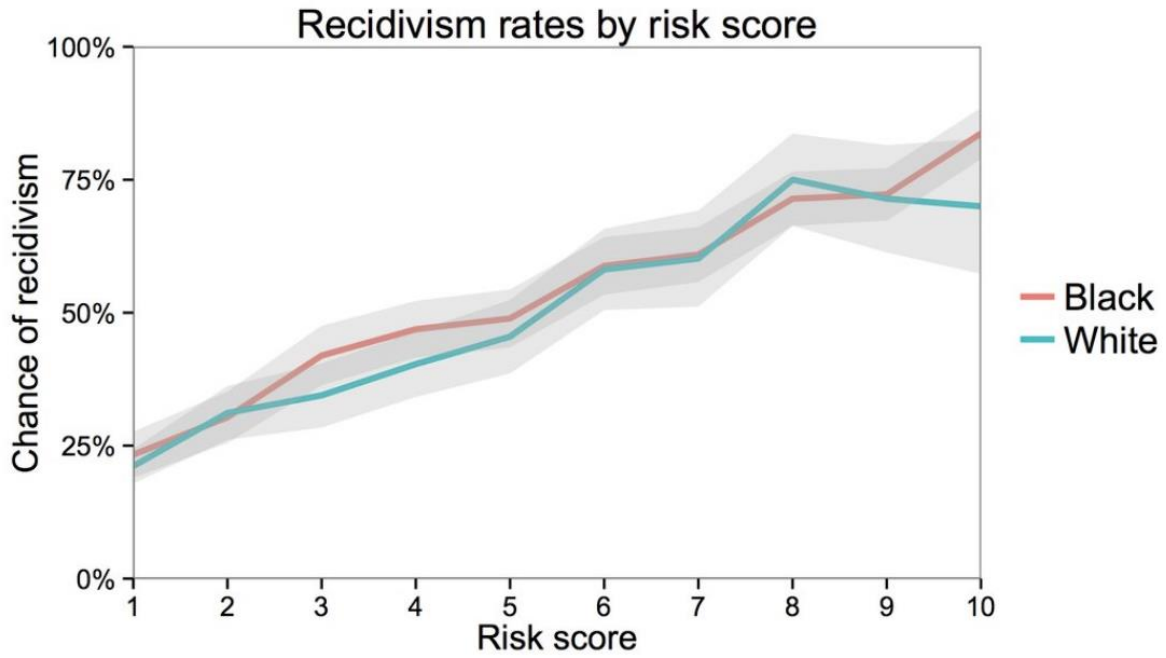

TU Delft

Do you have other examples of ML gone wrong? If so, why?

TU Delft

AI applications are not always huge failures, though!

- Politecnico di Milano, Musical Acoustics Lab (2021)
- AI approach to predict the sound of the violin through non-invasive measurements

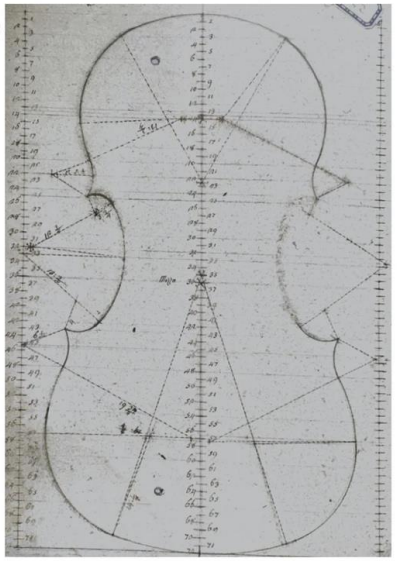

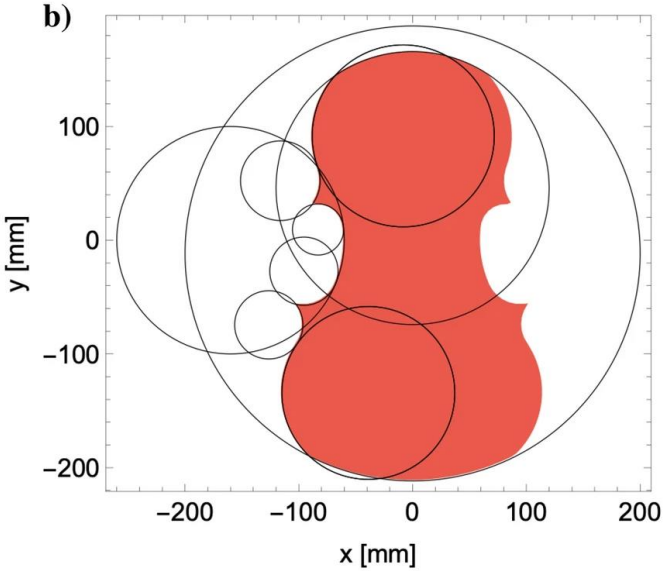

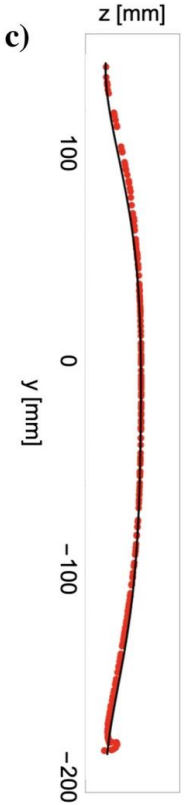

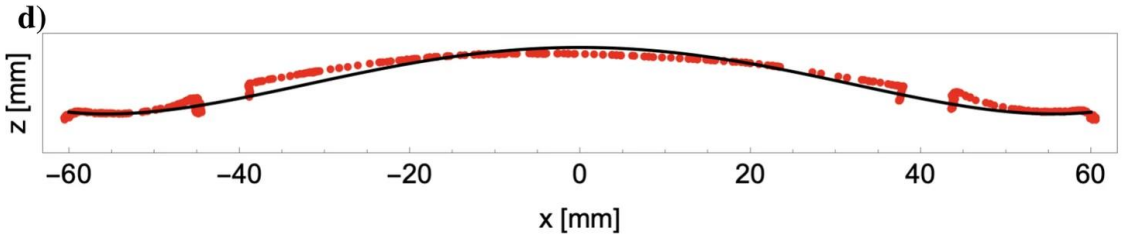

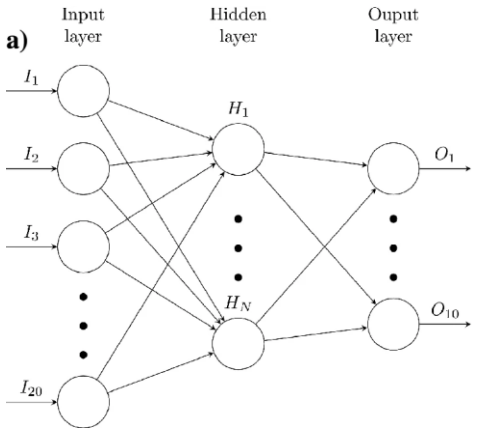

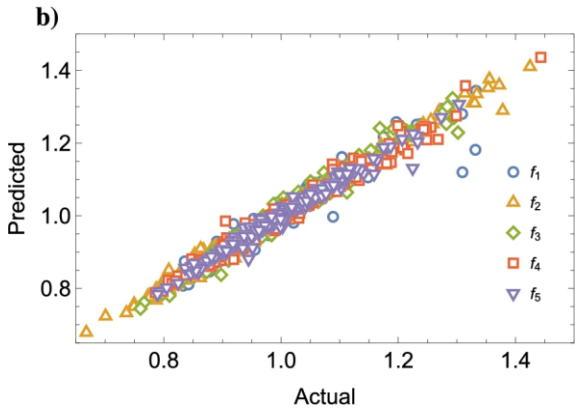

Source: Gonzalez, S., Salvi, D., Baeza, D. *et al.* A data-driven approach to violin making. *Sci Rep* **11**, 9455 (2021). https://doi-org.tudelft.idm.oclc.org/10.1038/s41598-021-88931-z

TU Delft

AI applications are not always huge failures, though!

- Universite' Libre de Bruxelles, IB2 Lab (2021)
- AI approach to predict the pathogenicity of rare diseases based on genetic information.

Source: Papadimitriou, S. *et al.* Predicting disease-causing variant combinations, *Proceedings of the National Academy of Sciences*. May 2019, DOI: https://doi.org/10.1073/pnas.1815601116.

TU Delft

AI algorithm analyses gene pair variants to improve rare disease diagnosis

A team of Belgian researchers are using artificial intelligence (AI) to identify the causes of rare disease and improve their diagnosis.

Chloe Kent

VarCoPP makes it possible to simultaneously test the combinations of different variants in gene pairs to predict their potential pathogenicity. Credit: Shutterstock

AI applications are not always huge failures, though!

- From a professor now at TUDelft, ESS-ICT (2020)
- AI approach to improve orthopaedics surgery processes.

Source: Oosterhoff, J. H., & Doornberg, J. N. (2020). Artificial intelligence in orthopaedics: false hope or not? A narrative review along the line of Gartner's hype cycle. *EFORT Open Reviews*, *5*(10), 593-603.

TU Delft

Jacobien H.F. Oosterhoff[1,2]
Job N. Doornberg[1,3]
Machine Learning Consortium[4]

- Artificial Intelligence (AI) in general, and Machine Learning (ML)-based applications in particular, have the potential to change the scope of healthcare, including orthopaedic surgery.
- The greatest benefit of ML is in its ability to learn from real-world clinical use and experience, and thereby its capability to improve its own performance.
- Many successful applications are known in orthopaedics, but have yet to be adopted and evaluated for accuracy and efficacy in patients' care and doctors' workflows.
- The recent hype around AI triggered hope for development of better risk stratification tools to personalize orthopaedics in all subsequent steps of care, from diagnosis to treatment.
- Computer vision applications for fracture recognition show promising results to support decision-making, overcome bias, process high-volume workloads without fatigue, and hold the promise of even outperforming doctors in certain tasks.
- In the near future, AI-derived applications are very likely to assist orthopaedic surgeons rather than replace us. 'If the computer takes over the simple stuff, doctors will have more time again to practice the art of medicine'.[76]

Do you have other examples of ML gone good?

TU Delft

Take-home messages

- AI is a socio-technical technology: technical aspects + social aspects

- AI is a data-driven technology: the quality of the predictions depends of the quality of the input data: Garbage In Garbage Out

- As most of the technologies, the technology itself is not inherently good or bad, but depends on the use you make of it!

- General tendency in the media is to:

1. Cherry-pick examples of AI to give an extreme representation of good and bad features of AI **(as I purposely did)**
2. Put the blame on AI for solely technical aspects, while neglecting the human component in the AI life-cycle: from data annotation, over to training and operationalization

- Be critical and analyse the whole context surrounding new and hyped technologies (e.g. AI, Metaverse, Blockchain).

TU Delft

Thank you for your attention!
Any questions?

Dr. Ir. Jacopo De Stefani – J.deStefani@tudelft.nl

TU Delft