

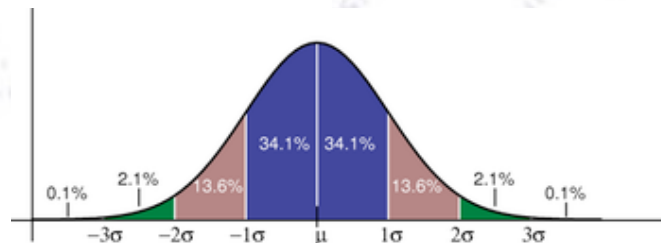


# Big Data Analysis

## Generative Adversarial Networks



Troels C. Petersen (NBI)



*"Statistics is merely a quantisation of common sense - Machine Learning is a sharpening of it!"*

# Generative Adversarial Networks

Invented (partly) by Ian Goodfellow in 2014, Generative Adversarial Networks (GANs) is a method for learning how to produce new (simulated) datasets from existing data.

The basic idea is, that **two networks “compete” against each other:**

- **Generative Network:** Produces new data trying to make it match the original.
- **Adversarial (Discriminatory) Network:** Tries to classify original and new data.

Typically, the generator is a de-convolutional NN, while the discriminating (adversarial) is convolutional NN.


“The coolest idea in machine learning in the last twenty years”

[Yann LeCun, French computer scientist]

# GAN drawing

Imagine that you want to write numbers that looks like hand writing.

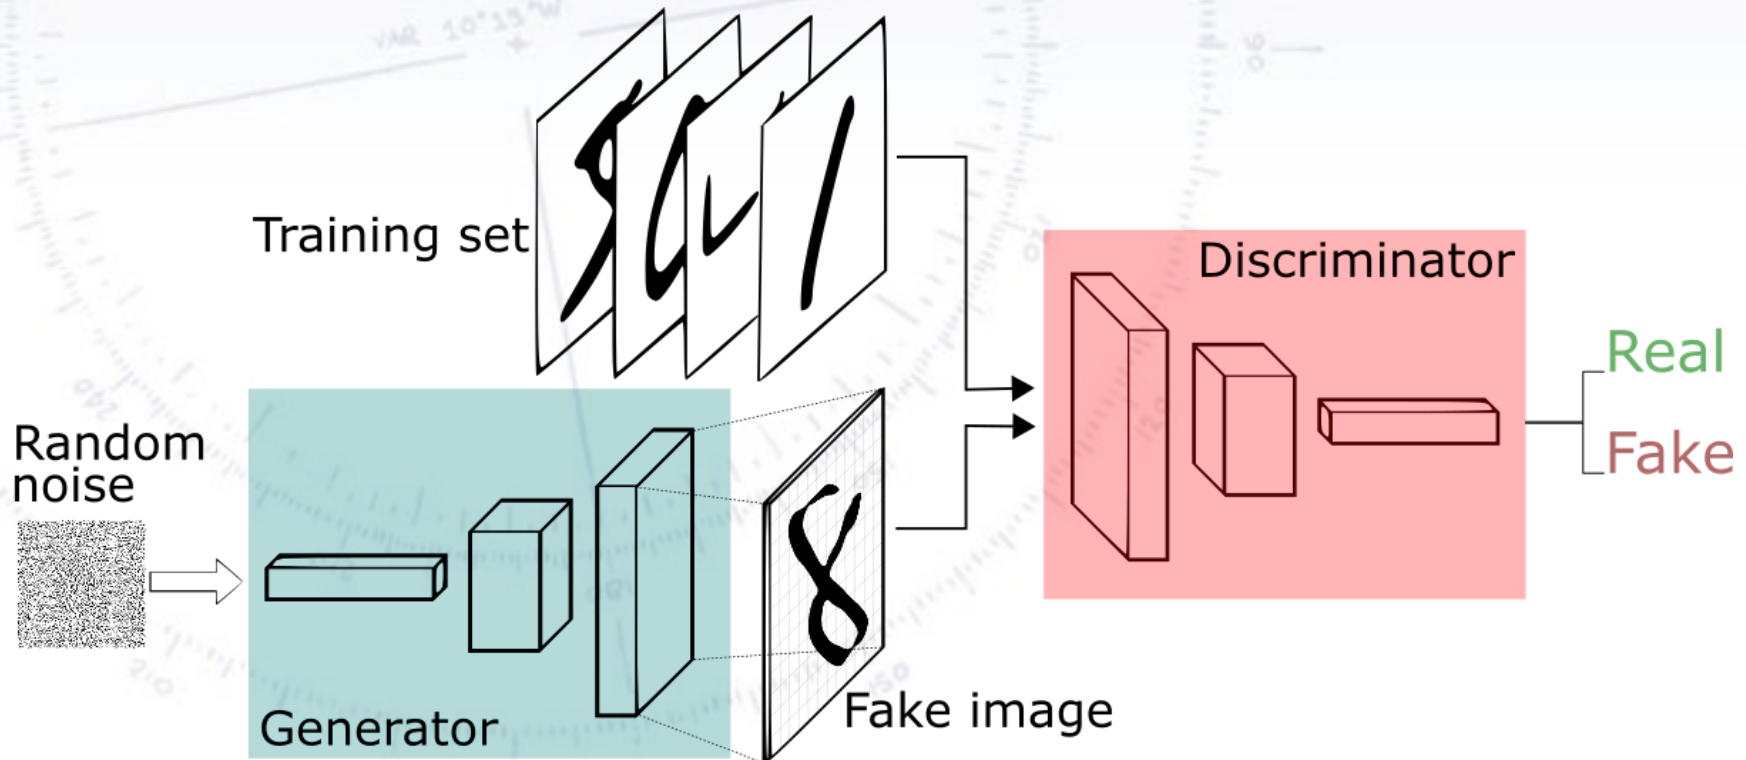
Given a large training set, you can ask you GAN to produce numbers. At first it will do poorly, but as it is “punished” by the discriminator, it improves, and at the end it might be able to produce numbers of **equal quality to real data**:



# GAN drawing

The discriminator/adversarial can also be seen as an addition to loss function, penalising (with  $\lambda$ ) an ability to see differences between real and fake:

$$\text{Loss} = \text{Loss} + \lambda \cdot L_{\text{Adversarial}}$$

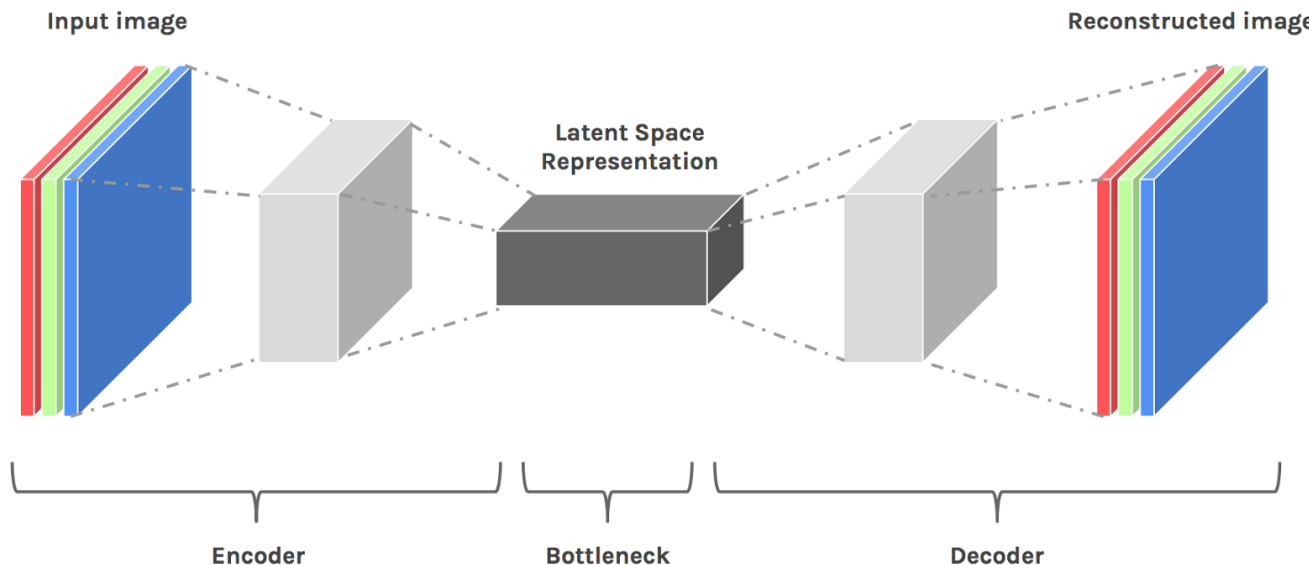


# Latent space

**Latent variables** are variables that are **inferred** instead of directly **observed**. They may correspond to some physical reality, e.g. temperature, (then also called hidden variables) but can also correspond to abstract concepts, e.g. mental state.

One advantage of using latent variables is that they can serve to reduce the dimensionality of data. Also, latent variables link observable data in the real world to symbolic data in the modelled world.

A **latent space** is one spanned by latent variables, thus containing the main features.



# Example: Latent space for PCA

Consider a 3 dimensional space on which we apply a PCA analysis.


Then the principle component will fall in some direction spanned by the three dimensions.

If we choose only to use this component, then this 1D direction forms the latent space:

- All 3D points can be boiled down to this line, and
- this line can give an approximation to all 3D points.

**This is a linear example in low dimensionality.**

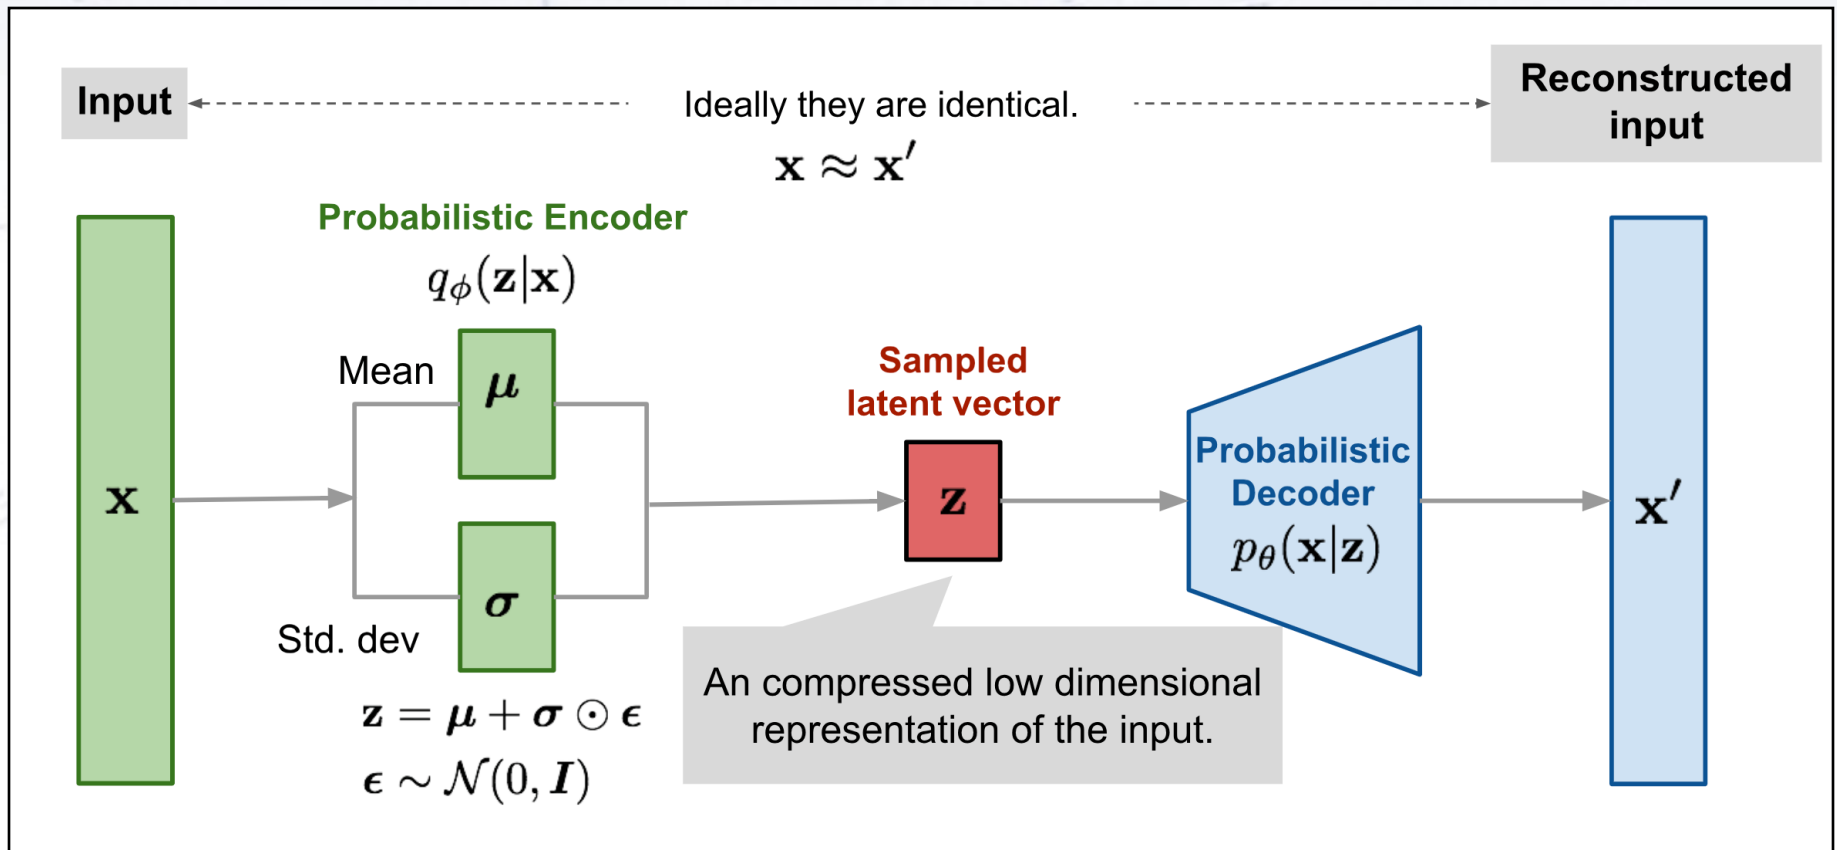
Typically, ML-problems are non-linear and in high dimensionality. Therefore, the latent spaces can also have significant dimensionality, though it should of course always have a (much) lower dimensionality than the problem itself.



# Variable AutoEncoders

An auto-encoder is a method (typically neural network) to learn efficient data codings in an unsupervised manner (hence the “auto”).

This dimensionality reduction is schematically shown below, and closely related to Generative Models.




# Latent space illustration

The below animation shows how latent spaces are a simplified representation of the more complex objects, containing the main features of these.

For this reason, one can do arithmetics (typically interpolate) between the inputs:

## Arithmetic in Latent Space





# Latent space illustration

The below animation shows how latent spaces are a simplified representation of the more complex objects, containing the main features of these.

For this reason, one can do arithmetics (typically interpolate) between the inputs:

## Interpolation in Latent Space



# GANs producing face images

In 2017, Nvidia published the result of their “AI” GANs for producing celebrity faces. There is of course a lot of training data...



# Evolution in facial GANs

There is quiet a fast evolution in GANs, and their ability to produce realistic results....



2014



2015



2016

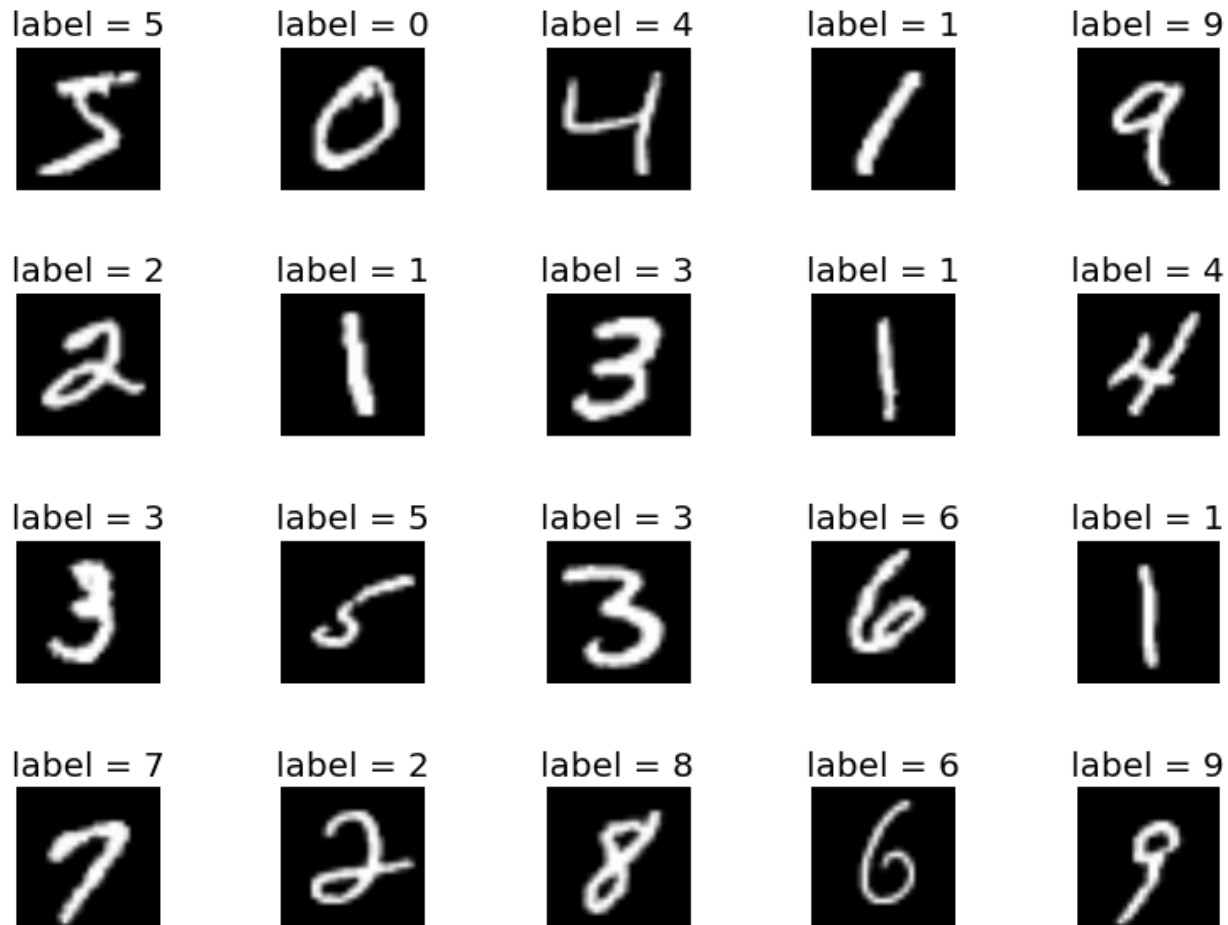


2019

**FAKE!**

# MNist data: Handwritten numbers

A “famous case” has been hand written numbers. The data consists of 28x28 gray scale images of numbers. While that spans a large space, the latent space is probably (surely!) much smaller, as far from all combinations of pixels and intensities are present.



# MNist data: Handwritten numbers

A “famous case” has been hand written numbers. The data consists of 28x28 gray scale images of numbers. While that spans a large space, it is probably (surely!) much smaller, as far from the origin as the intensities are present.

*With GANs, you can produce handwritten letters again - sort of!*

label = 5



label = 7



label = 2



label = 8



label = 6



label = 1



label = 6



label = 9



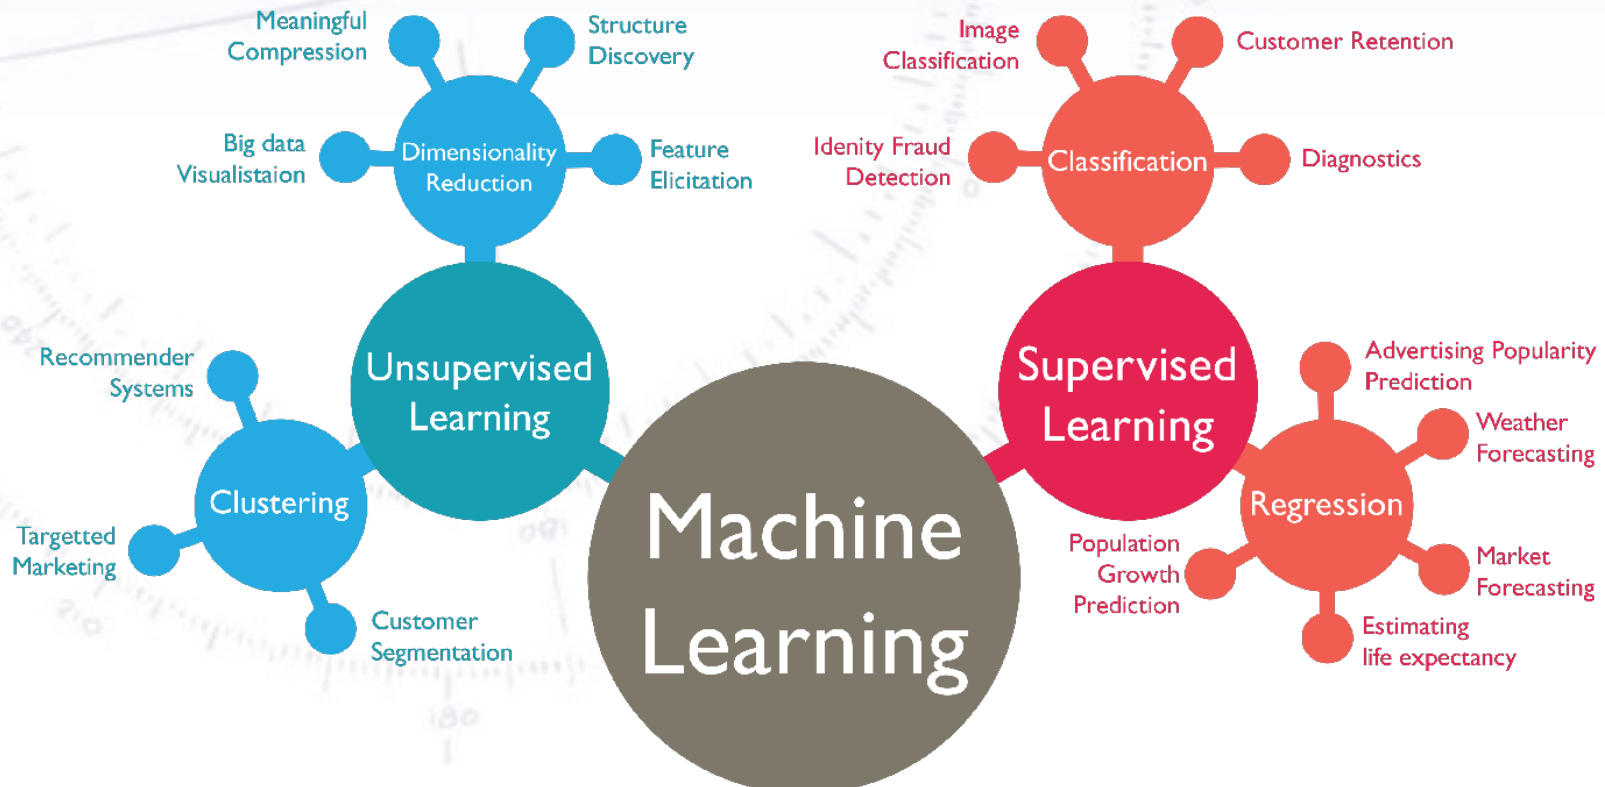


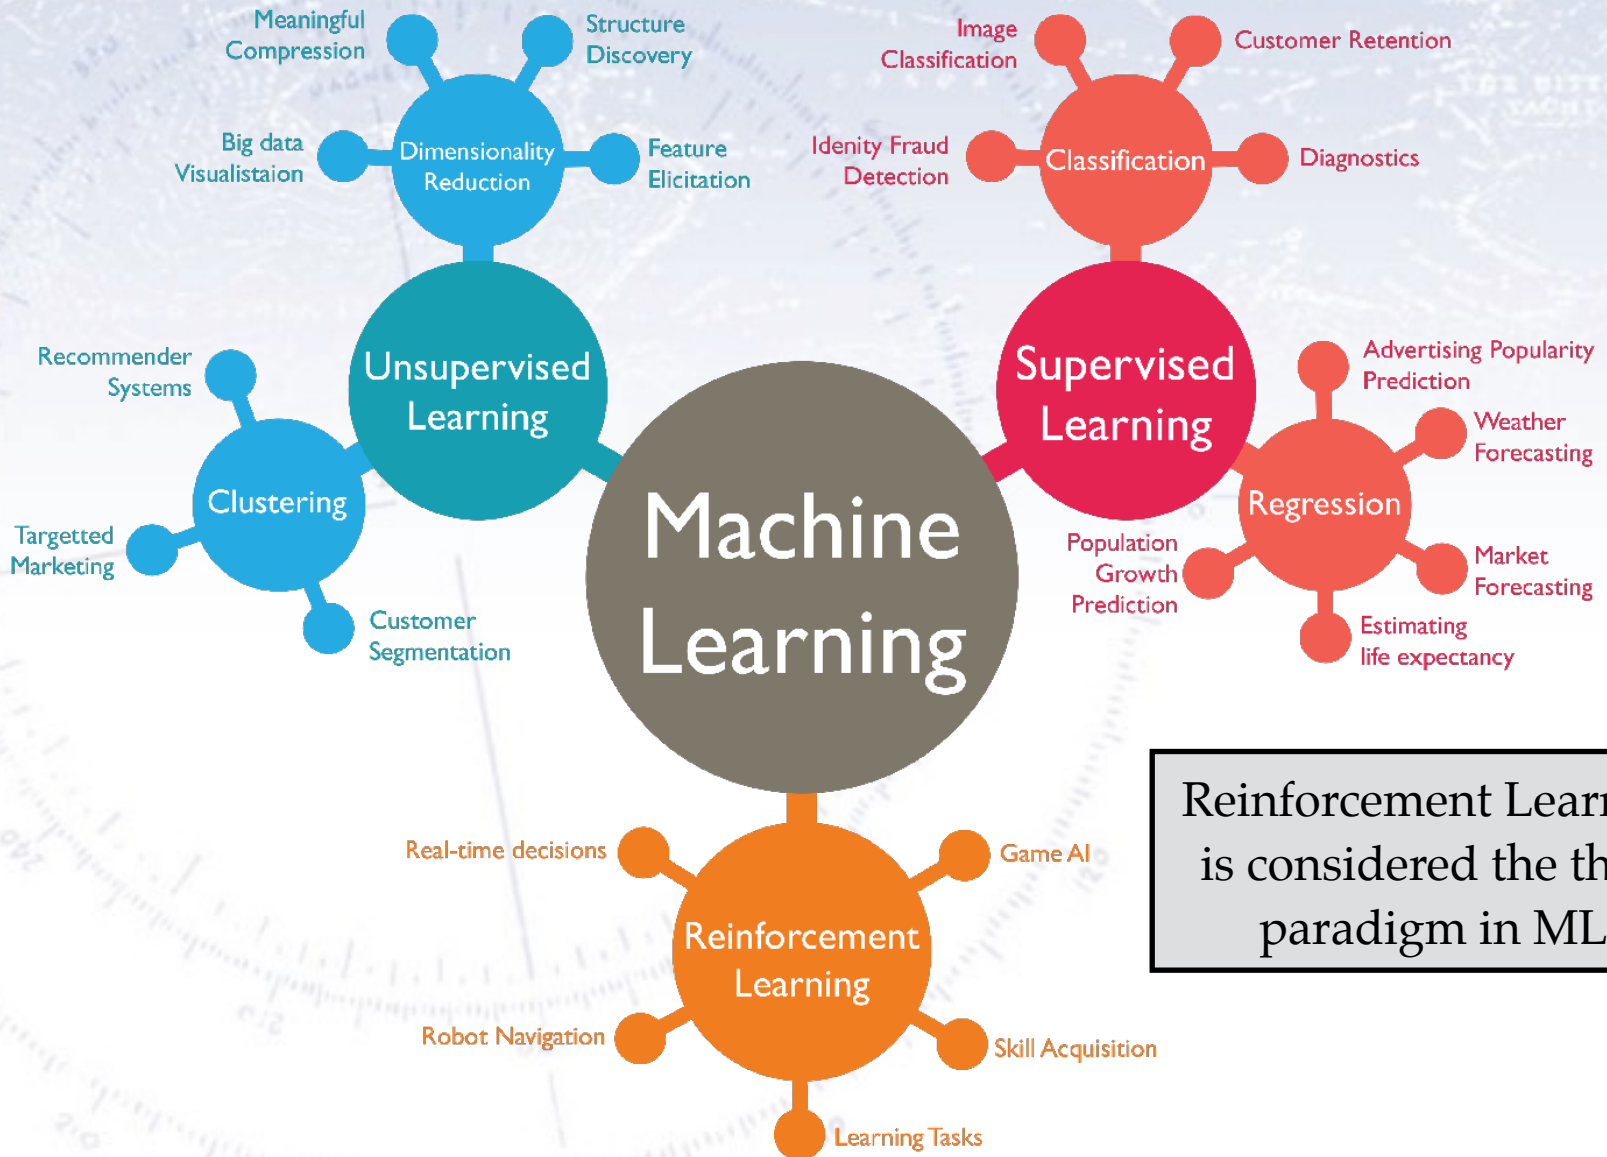
# Reinforcement Learning

# Classification vs. Regression

## Unsupervised learning vs. supervised

Machine Learning can be supervised (you have correctly labelled examples) or unsupervised (you don't)... [or reinforced]. Following this, one can be using ML to either classify (is it A or B?) or for regression (estimate of X).





Reinforcement Learning is considered the third paradigm in ML.



# Reinforcement Learning

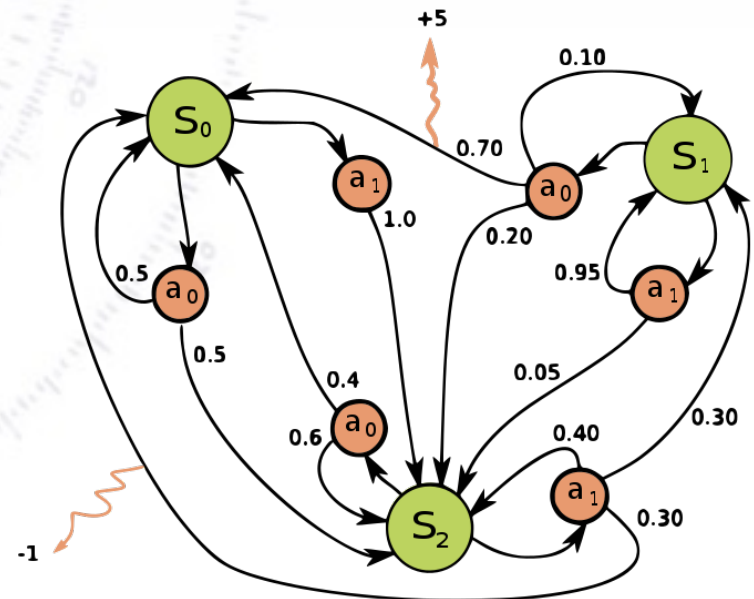
Reinforcement Learning (RL) does not need data per se, but rather an environment / set of rules in which it needs to optimise its actions / behaviour.

In doing so, the RL needs to find a balance between exploration (of uncharted territory) and exploitation (of current knowledge).

The environment can be formulated as a Markov Decision Process (MDP), as shown below.

Reinforcement Learning does not assume knowledge of the MDP (i.e. it doesn't know what environment it is in - all it needs is a score).

And typically RL has great success in (potentially very) large environments, such as "real life".



# One program to rule them all

Last December, AlphaZero was introduced to play three classic strategy board games...

A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play

David Silver<sup>1,2,\*†</sup>, Thomas Hubert<sup>1,\*</sup>, Julian Schrittwieser<sup>1,\*</sup>, Ioannis Antonoglou<sup>1</sup>, Matthew Lai<sup>1</sup>, Arthur Guez<sup>1</sup>, Marc Lanctot<sup>1</sup>, Laurent Sifre<sup>1</sup>, Dharshan Kumaran<sup>1</sup>, Thore Graepel<sup>1</sup>, Timothy Lillicrap<sup>1</sup>, Karen Simonyan<sup>1</sup>, Demis Hassabis<sup>1,†</sup>

<sup>1</sup>DeepMind, 6 Pancras Square, London N1C 4AG, UK.

<sup>2</sup>University College London, Gower Street, London WC1E 6BT, UK.

↔<sup>†</sup>Corresponding author. Email: [dauidsilver@google.com](mailto:dauidsilver@google.com) (D.S.); [dhcontact@google.com](mailto:dhcontact@google.com) (D.H.)

↔\* These authors contributed equally to this work.

- Hide authors and affiliations

Science 07 Dec 2018:

Vol. 362, Issue 6419, pp. 1140-1144

DOI: 10.1126/science.aar6404

# One program to rule them all

Last December, AlphaZero was introduced to play three classic strategy board games...

A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play

David Silver<sup>1,2,\*†</sup>, Thomas Hubert<sup>1,\*</sup>, Julian Schrittwieser<sup>1,\*</sup>, Ioannis Antonoglou<sup>1</sup>, Matthew Lai<sup>1</sup>, Arthur Guez<sup>1</sup>, Marc Lanctot<sup>1</sup>, Laurent Sifre<sup>1</sup>, Dhharshan Kumaran<sup>1</sup>, Thore Graepel<sup>1</sup>, Timothy Lillicrap<sup>1</sup>, Karen Simonyan<sup>1</sup>, Demis Hassabis<sup>1,†</sup>

<sup>1</sup>DeepMind, 6 Pancras Square, London N1C 4AG, UK.

<sup>2</sup>University College London, Gower Street, London WC1E 6BT, UK.

↔<sup>†</sup>Corresponding author. Email: [dauidsilver@google.com](mailto:dauidsilver@google.com) (D.S.); [dhcontact@google.com](mailto:dhcontact@google.com) (D.H.)

↔\* These authors contributed equally to this work.

- Hide authors and affiliations

Science 07 Dec 2018:

Vol. 362, Issue 6419, pp. 1140-1144

DOI: 10.1126/science.aar6404

After four hours of training it beat the best chess program (28 wins, 0 losses, and 72 draws).

**Within 24 hours AlphaZero achieved a superhuman level of play in ALL three games by defeating world-champion programs.... using only the rules!**