Quand l'essentiel tient en si peu de choses :-)

Ou, la magie des espaces latents...

Un exemple d'apprentissage non supervisé.

FIDLE

Journée GRICAD - 3 novembre 2020

Soraya ARIAS – INRIA
Eric MALDONADO – INRAE
Jean-Luc PAROUTY – SIMaP

GRIC

Thanks to :

Remenber that **Bigfoot** and **Jean-Zay** are good for you !

Cooking Ingredients :

[*-learning]

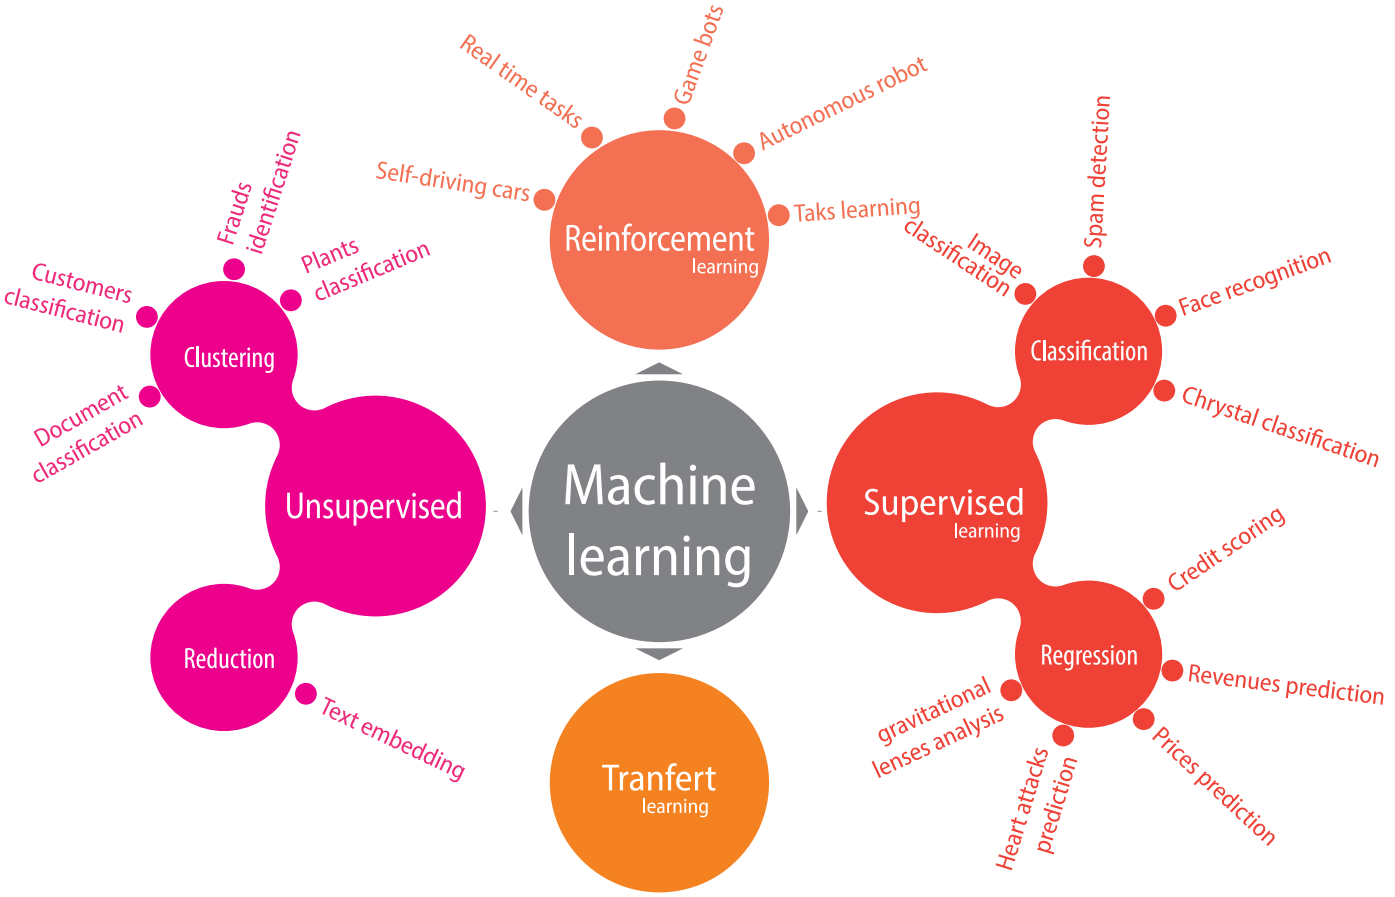
Machine learning
Supervised learning
Classification
Face recognition
Chrystal classification
Spam detection
Image classification
Regression
Credit scoring
Revenues prediction
Prices prediction
Heart attacks prediction
gravitational lenses analysis
Unsupervised
Clustering
Frauds identification
Plants classification
Customers classification
Document classification
Reduction
Text embedding
Reinforcement learning
Real time tasks
Game bots
Autonomous robot
Taks learning
Self-driving cars
Tranfert learning

Supervised learning

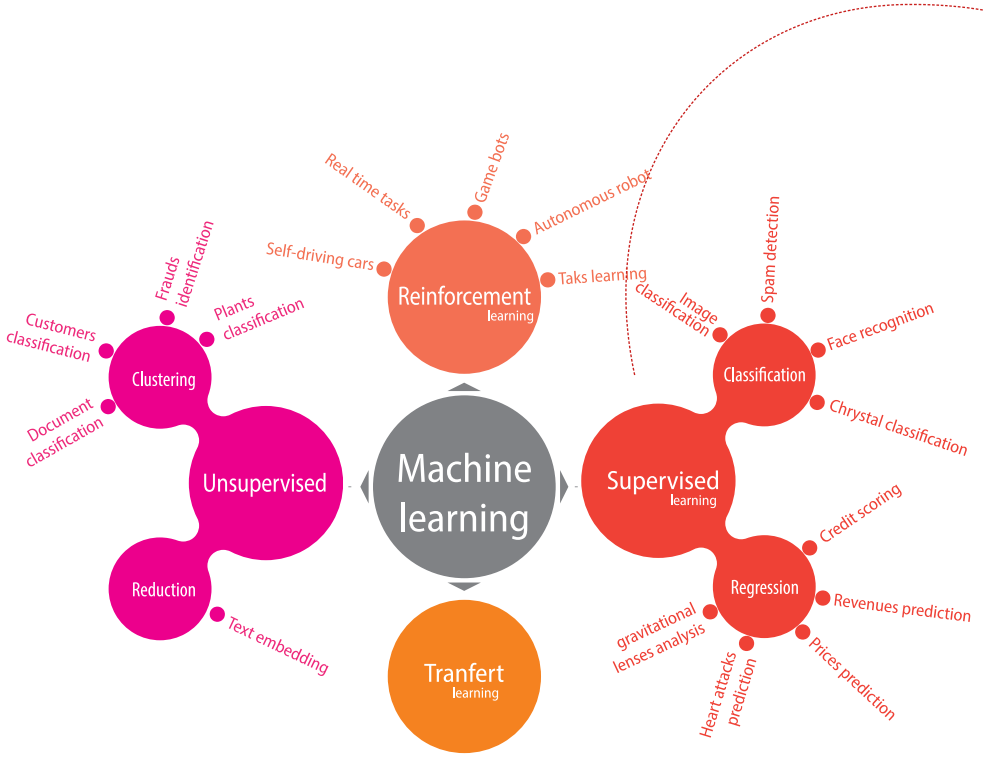

Learning
from **examples**

Supervised learning

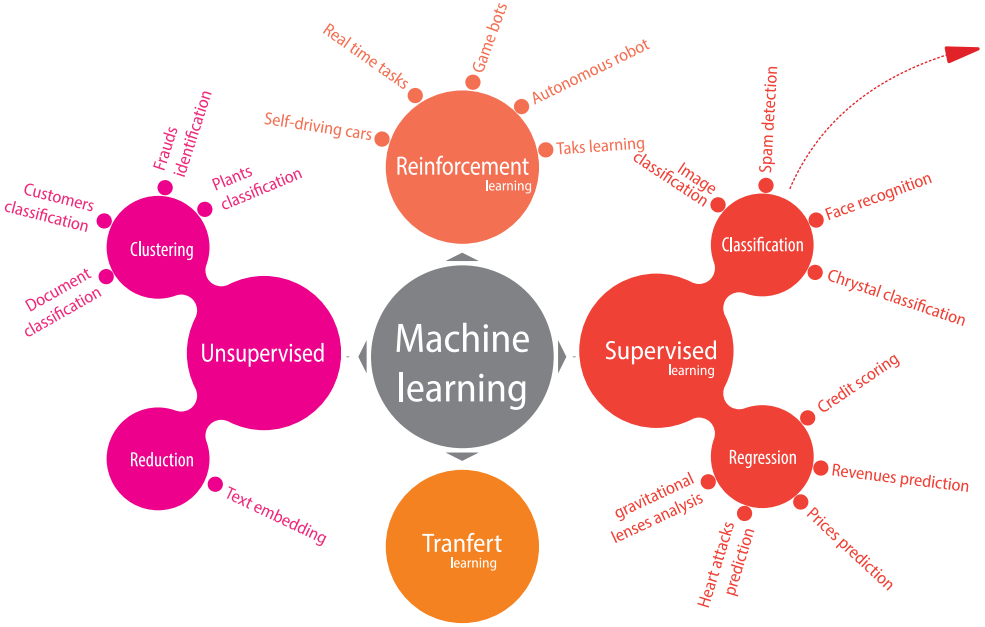

Classification :

Predict qualitative informations

This is a cat

This is a rabbit

Tell me,
what is it ?

Supervised learning

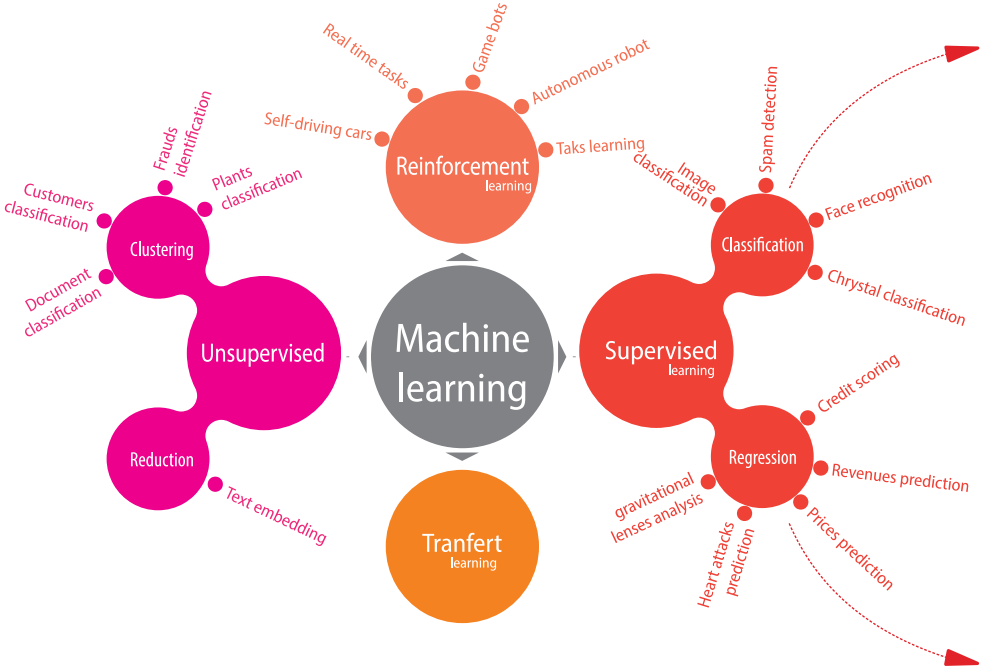

Classification :

Predict qualitative informations

This is a cat

This is a rabbit

Tell me, what is it ?

Régression :

Predict quantitative informations

150 K€

400 K€

120 K€

100 K€

Tell me, what's the price ?

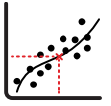

Unsupervised learning

Learning from data **alone**

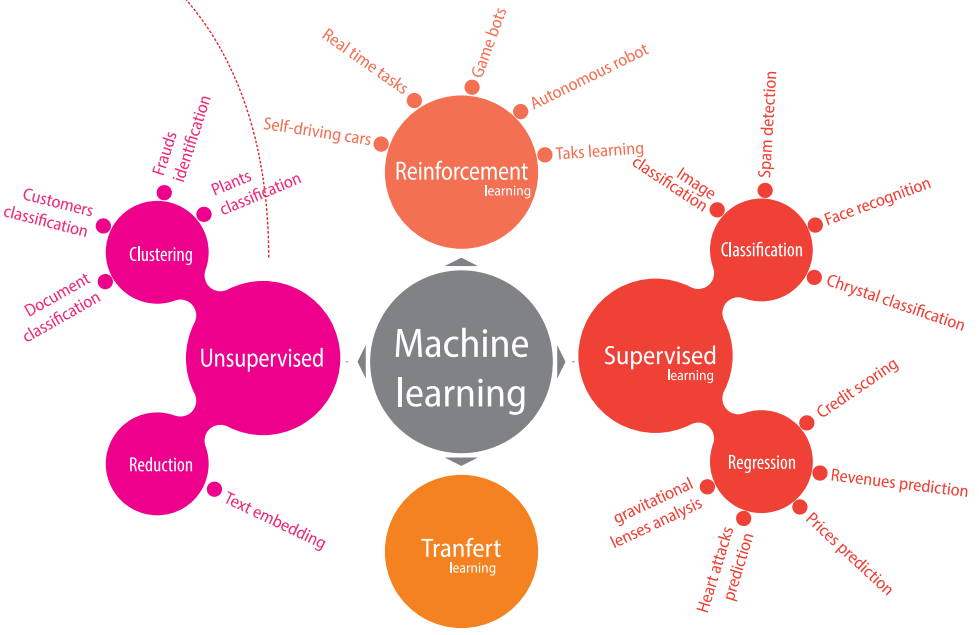

Unsupervised learning

Learning from data alone

Clustering :
Finding Common Relationships

What is the relationship between these data ?

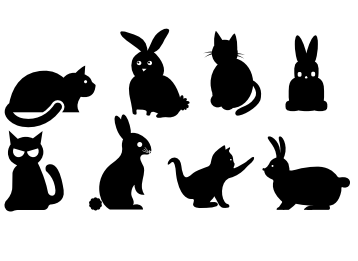

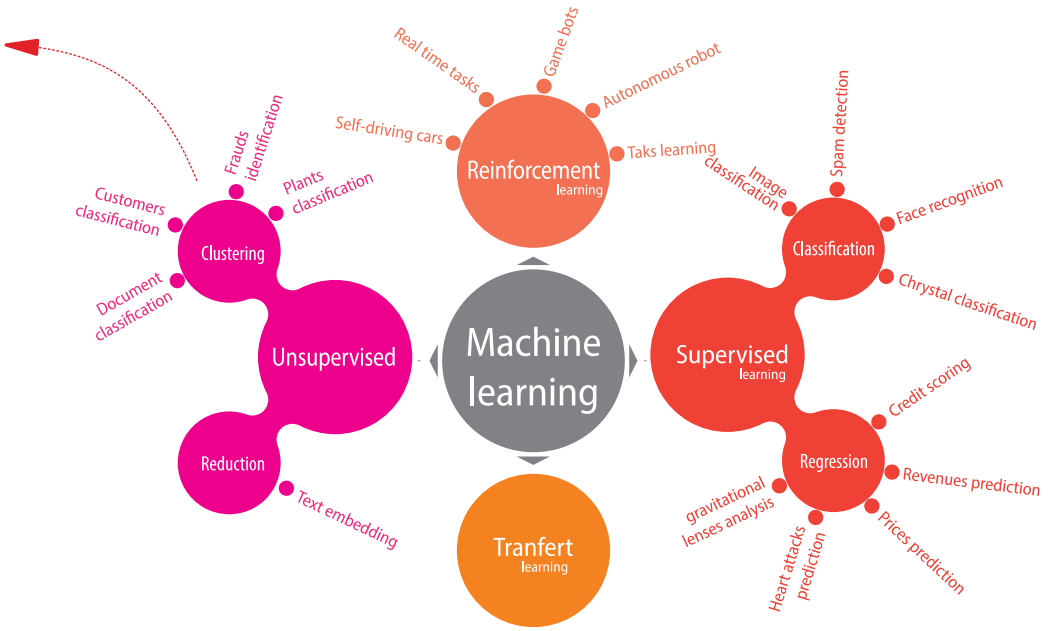

Unsupervised learning

Learning from data alone

Clustering :
Finding Common Relationships

What is the relationship between these data ?

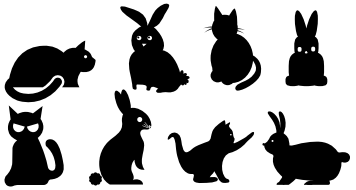

Reduction :
Reduce the number of dimensions

Simplify while keeping meaning

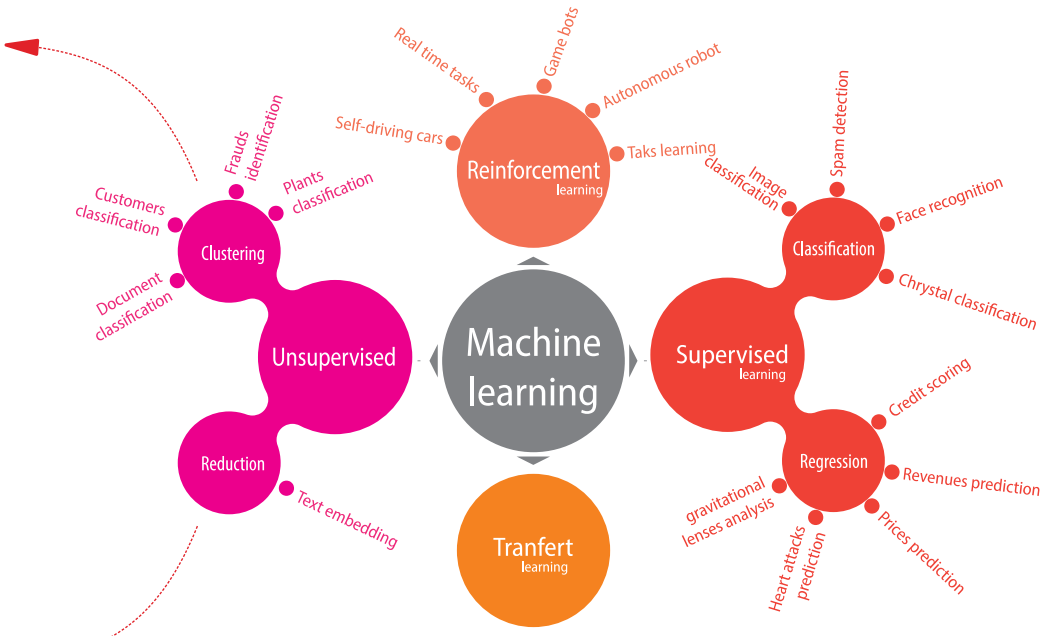

Variational Antoencoder (VAE)

FIDLE

Variational Autoencoder

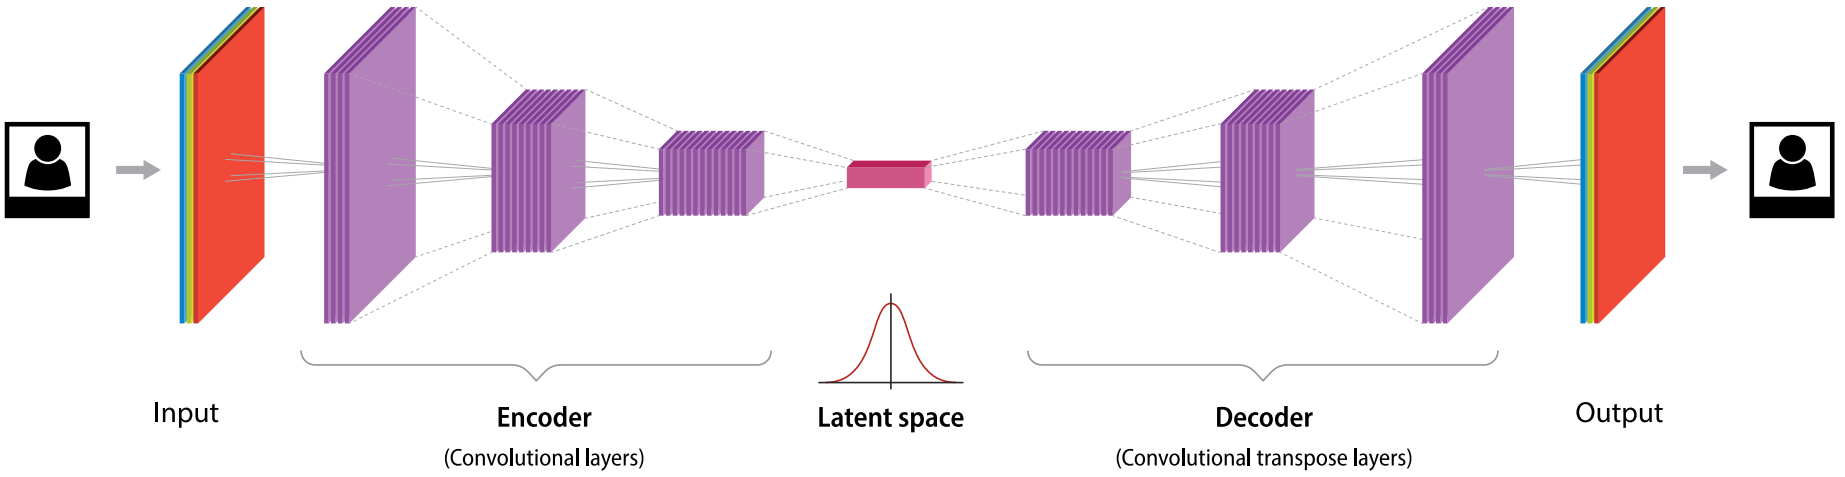

Encoder

By Jan Kroon, from Pexels.com

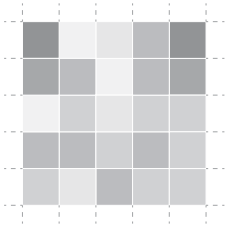

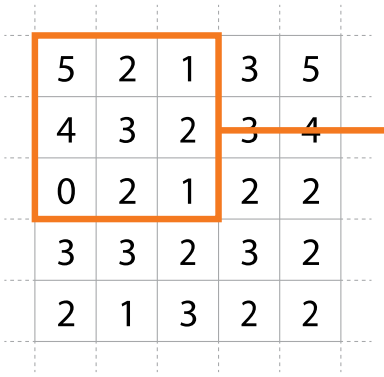

Image piece

5	2	1
4	3	2
0	2	1

X

⊗

Kernel 3x3

1	0	1
0	1	0
1	0	1

ω

=

10

y

$$y = 5x1 + 2x0 + 1x1 + 4x0 + 3x1 + 2x0 + 0x1 + 2x0 + 1x1 = 10$$

$$y = \sum_{i=1}^{n} \sum_{j=1}^{m} x_{i,j} \cdot \omega_{i,j} \quad \text{with} \begin{cases} n & \text{kernel width} \\ m & \text{kernel height} \end{cases}$$

2D convolution

⊗ is Hadamard product

Encoder

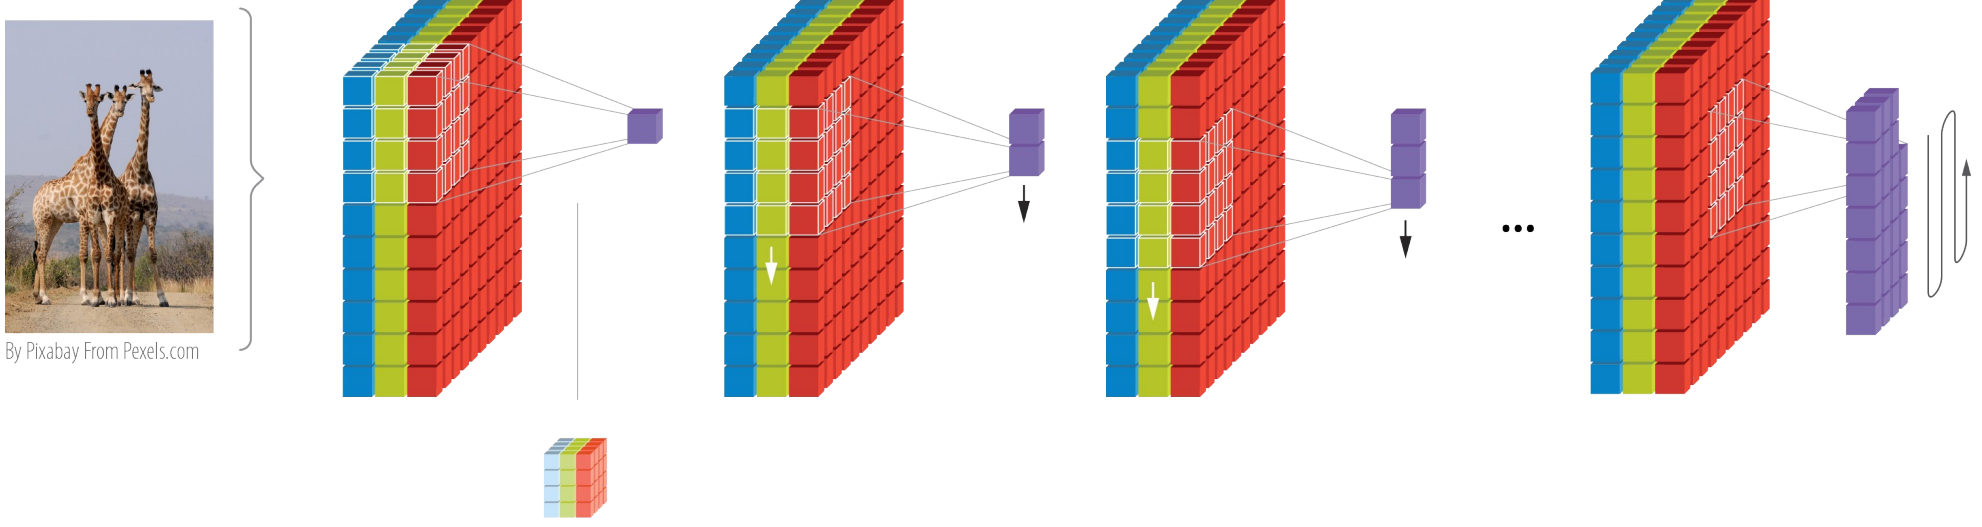

3D convolution

Convolutions

tf.keras.layers.Conv2D

Convolution layer

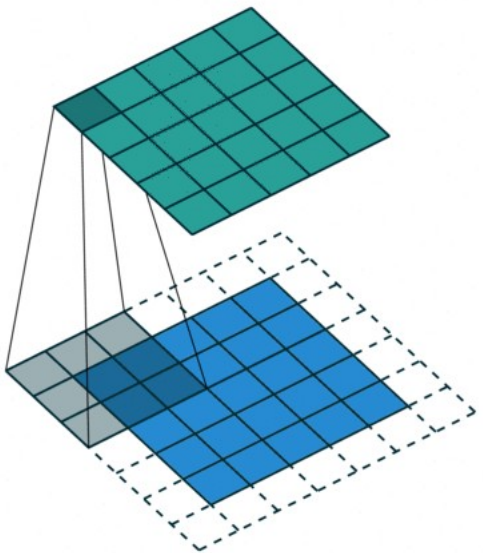

Original image

Stride
Step size
=> 1

Padding
Active (« same ») or not (« valid »)
=>Active

Kernel
=> 3x3

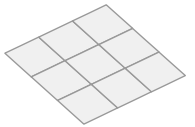

Illustration : An Introduction to different Types of Convolutions in Deep Learning
https://towardsdatascience.com/types-of-convolutions-in-deep-learning-717013397f4d

Convolutions

tf.keras.layers.Conv2D

Convolution layer
(2 x 2)

Original image
(5 x 5)

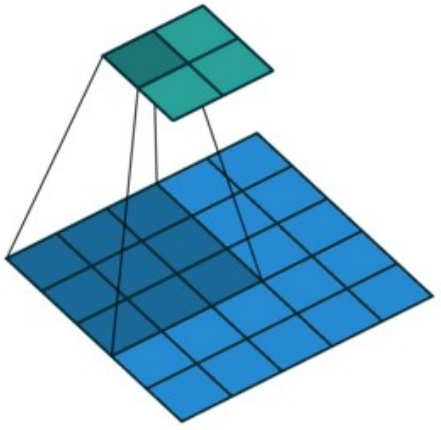

Stride
=>2

Padding
=> Desactivated (« valid »)

Kernel
=> (3,3)

Illustration : An Introduction to different Types of Convolutions in Deep Learning
https://towardsdatascience.com/types-of-convolutions-in-deep-learning-717013397f4d

Variational Autoencoder

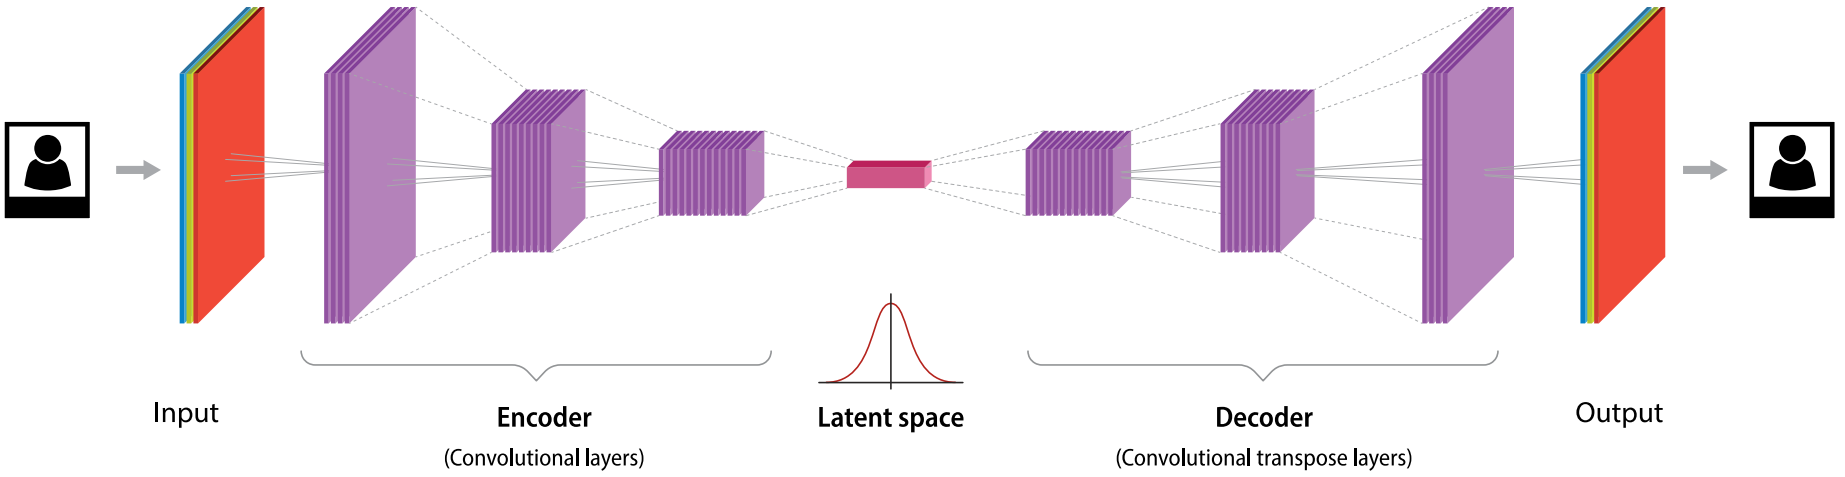

Transposed Convolutions

tf.keras.layers.Conv2DTranspose

Original image
(5 x 5)

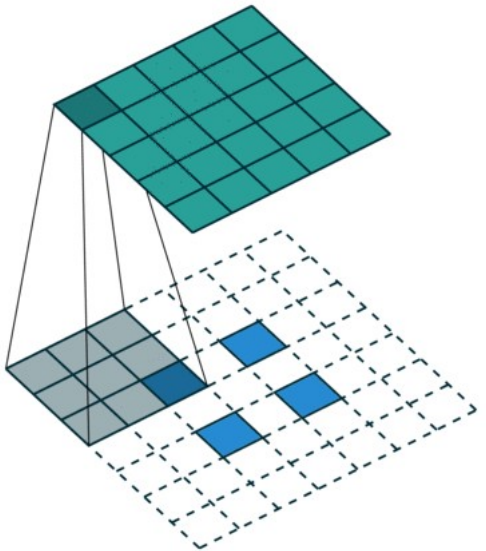

Convolution layer
(2 x 2)

Stride
Step size

Dilatation rate
Image expansion

Padding
Active or not

Kernel
Kernel size

Illustration : An Introduction to different Types of Convolutions in Deep Learning
https://towardsdatascience.com/types-of-convolutions-in-deep-learning-717013397f4d

Variational Autoencoder

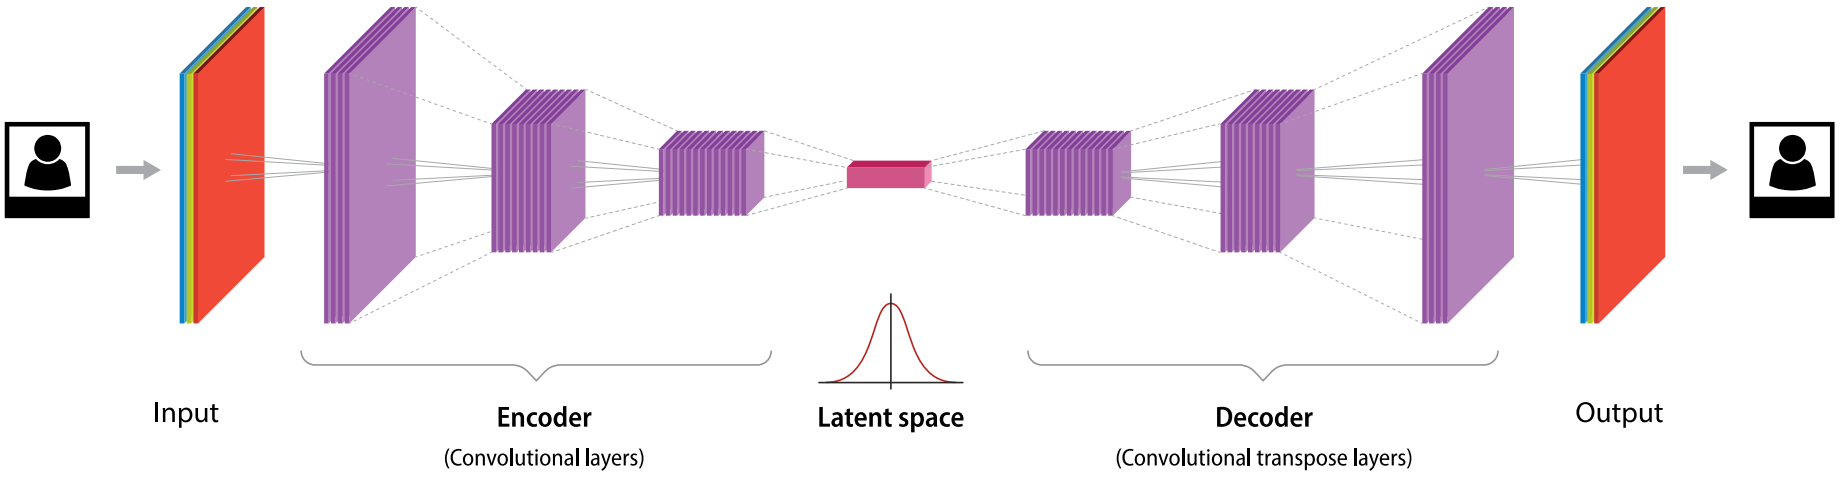

#1

MNIST dataset

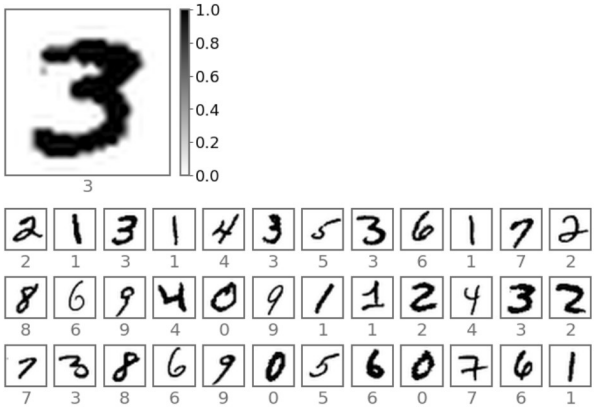

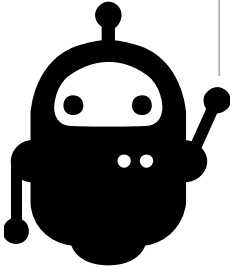

VAE with MNIST

INPUT

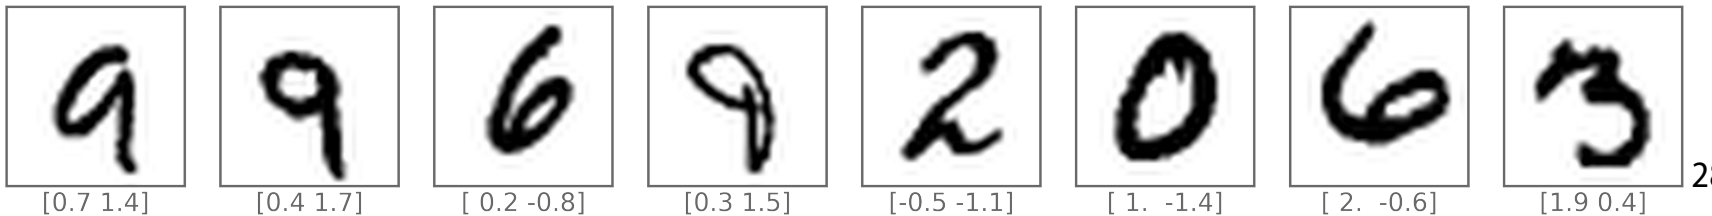

28x28

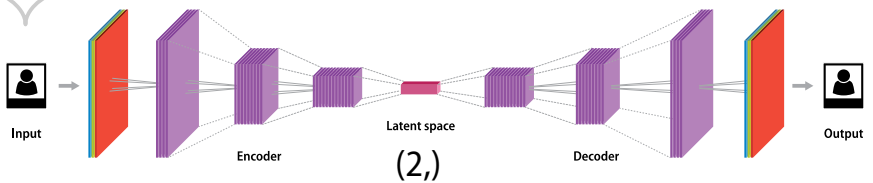

OUTPUT

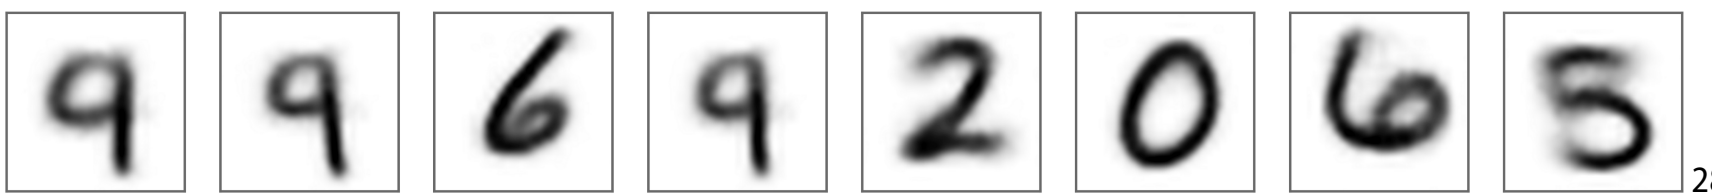

28x28

Dataset : 60.000 images, 28x28x1, 50 Mo - Run time : 6’ 30 on VT100 GPU

VAE with MNIST

INPUT

Latent Space
(u,v)

OUTPUT

Latent space

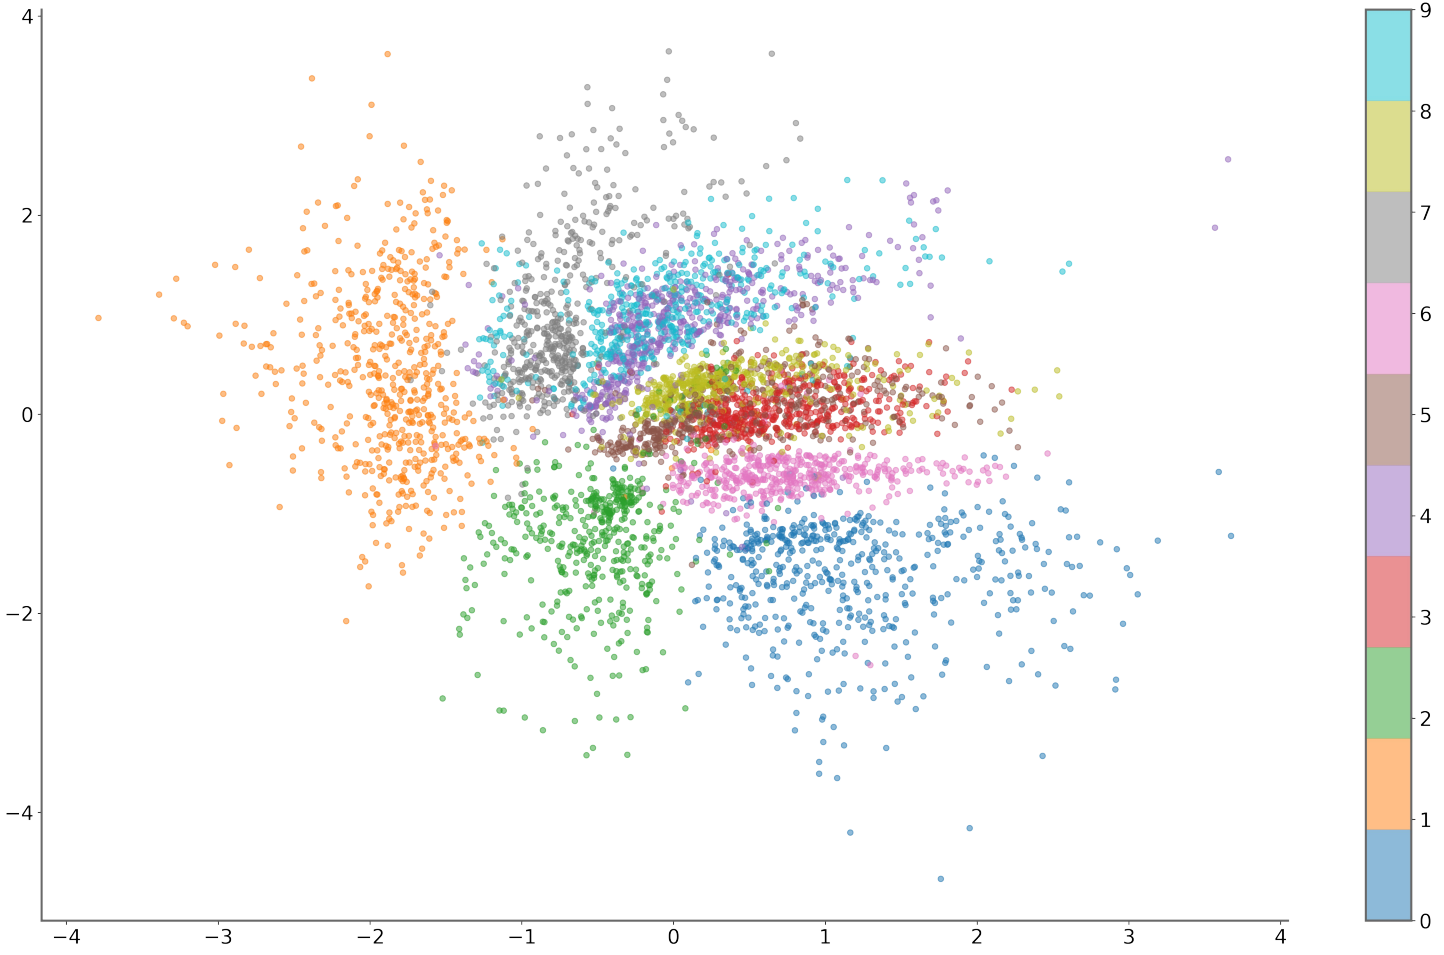

Latent space

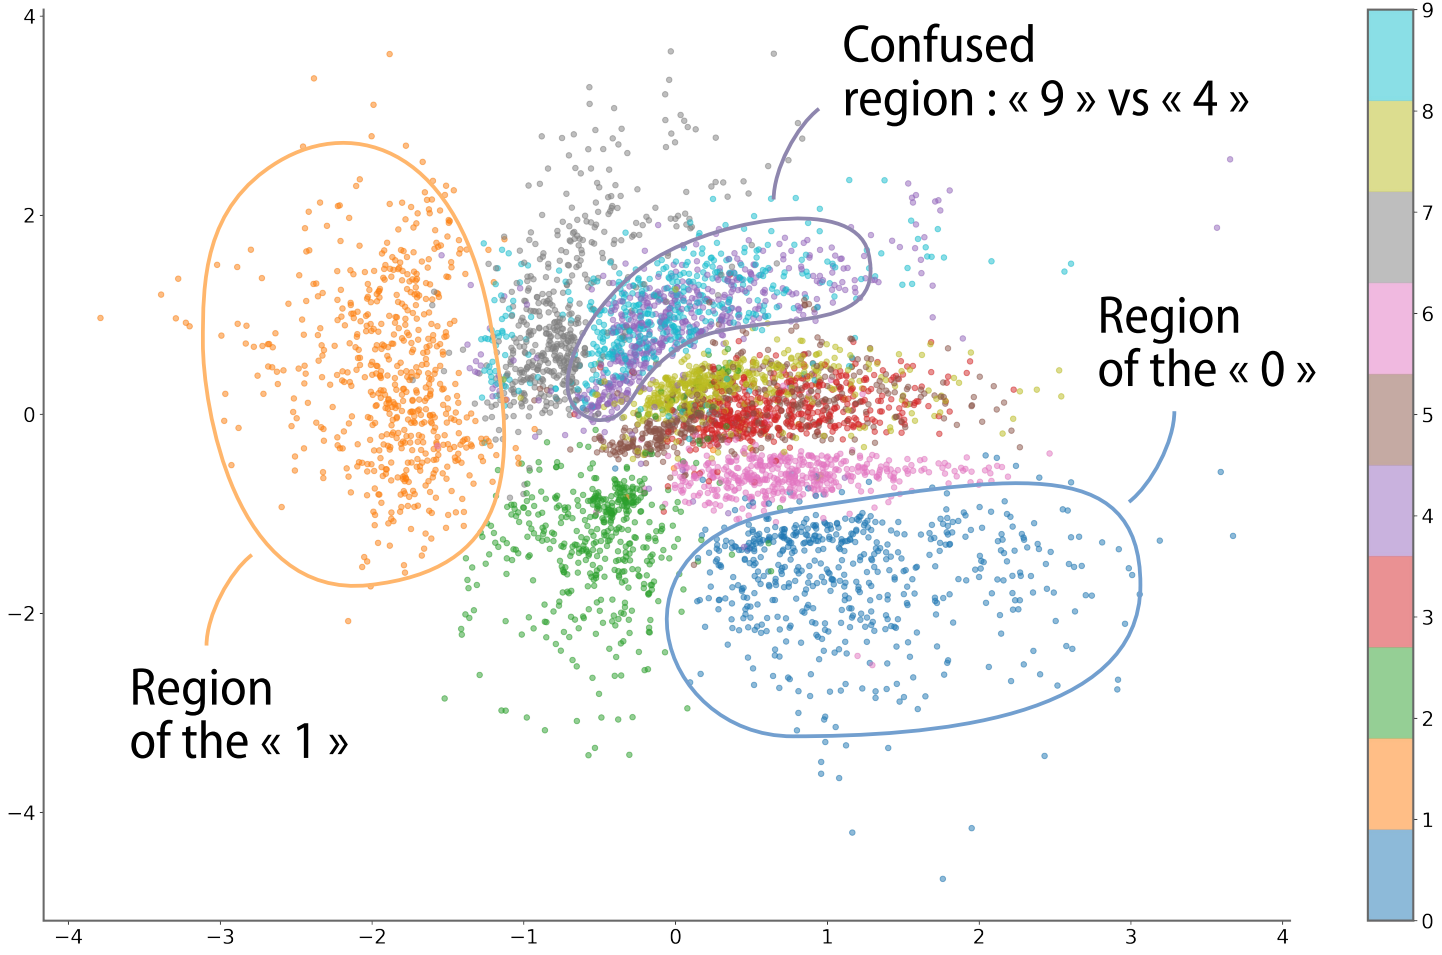

Latent space

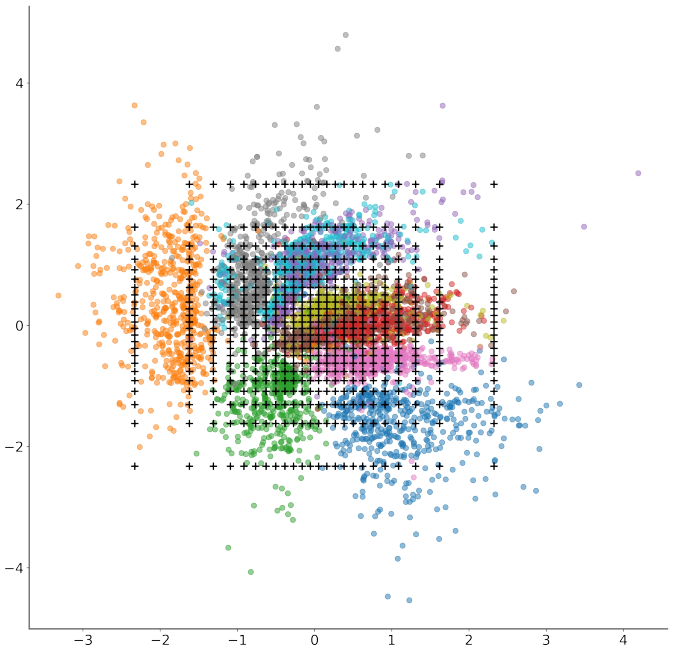

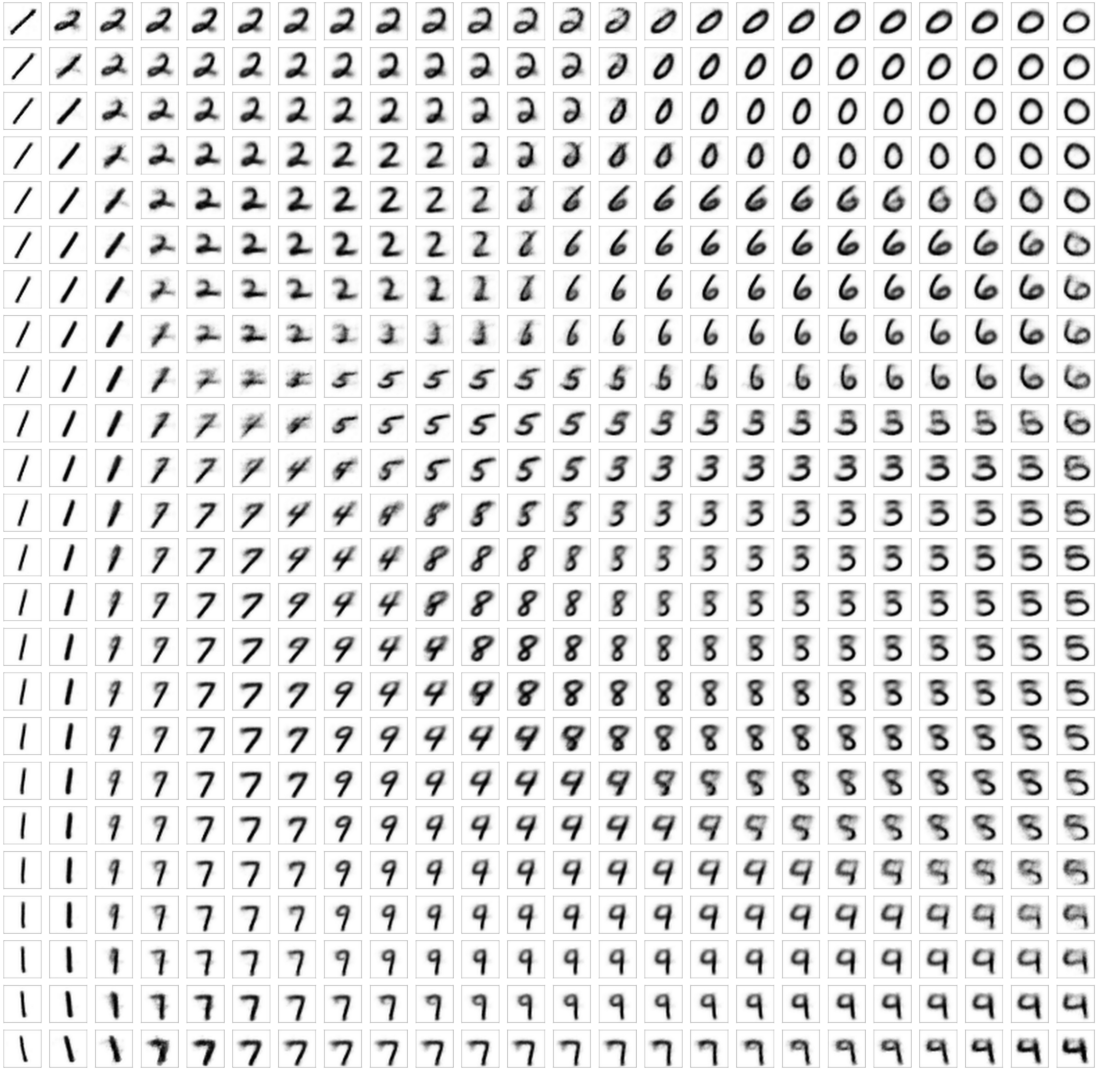

#2

CelebA dataset

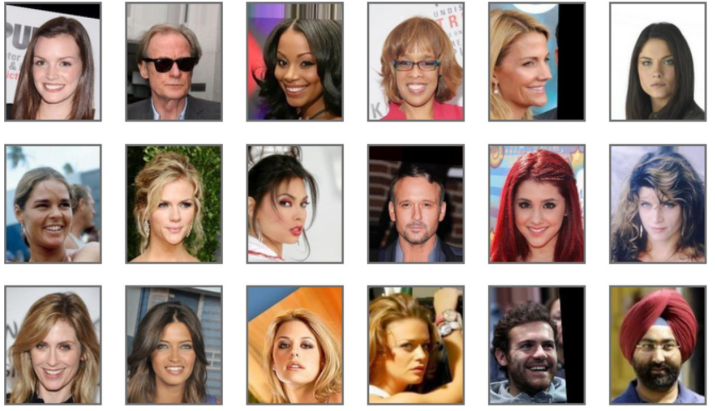

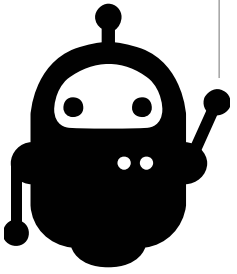

VAE with CelebA

INPUT

192x160

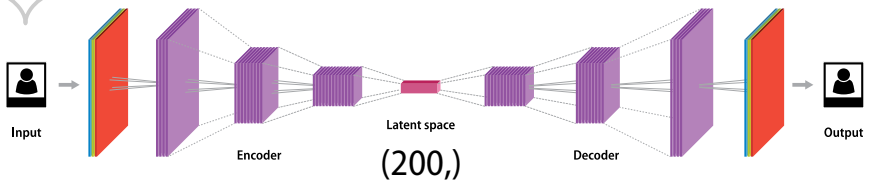

OUTPUT

192x160

Dataset : 200.000 images, 192x160x3, 140 GB - Run time : 1h50 on VT100 GPU

Using decoder as a generative network

Pick a
Random position
in latent space

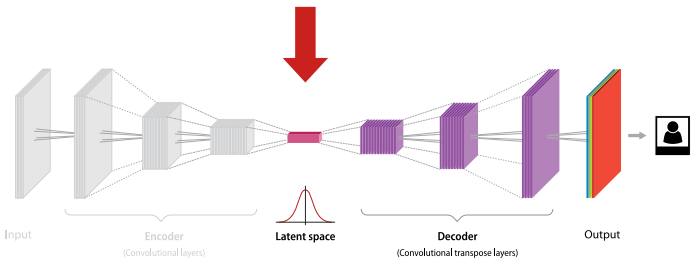

Face generation

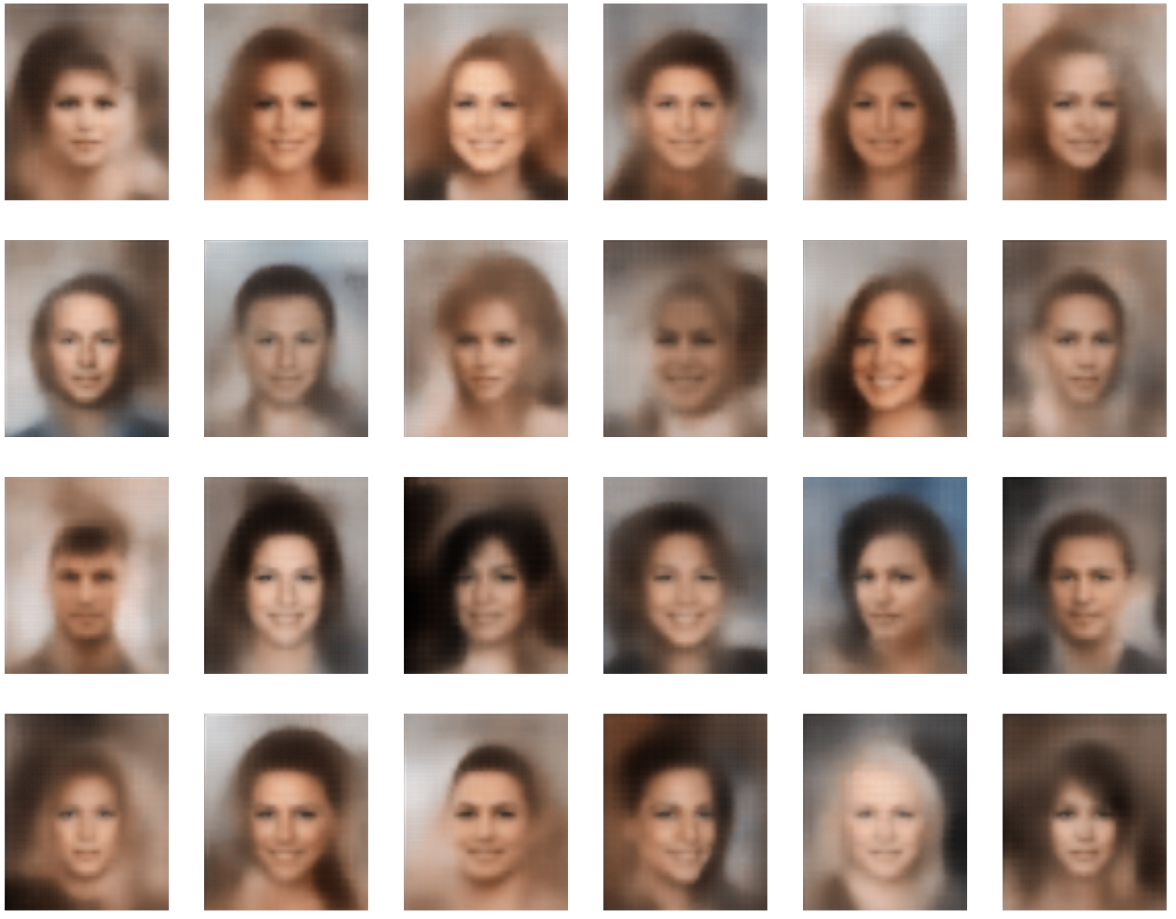

Morphing from latent space

Woman
Blond hair
Smiling

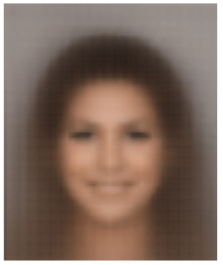

Woman
Black hair
Smiling

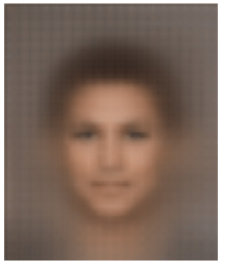

Man
Black hair
Not smiling

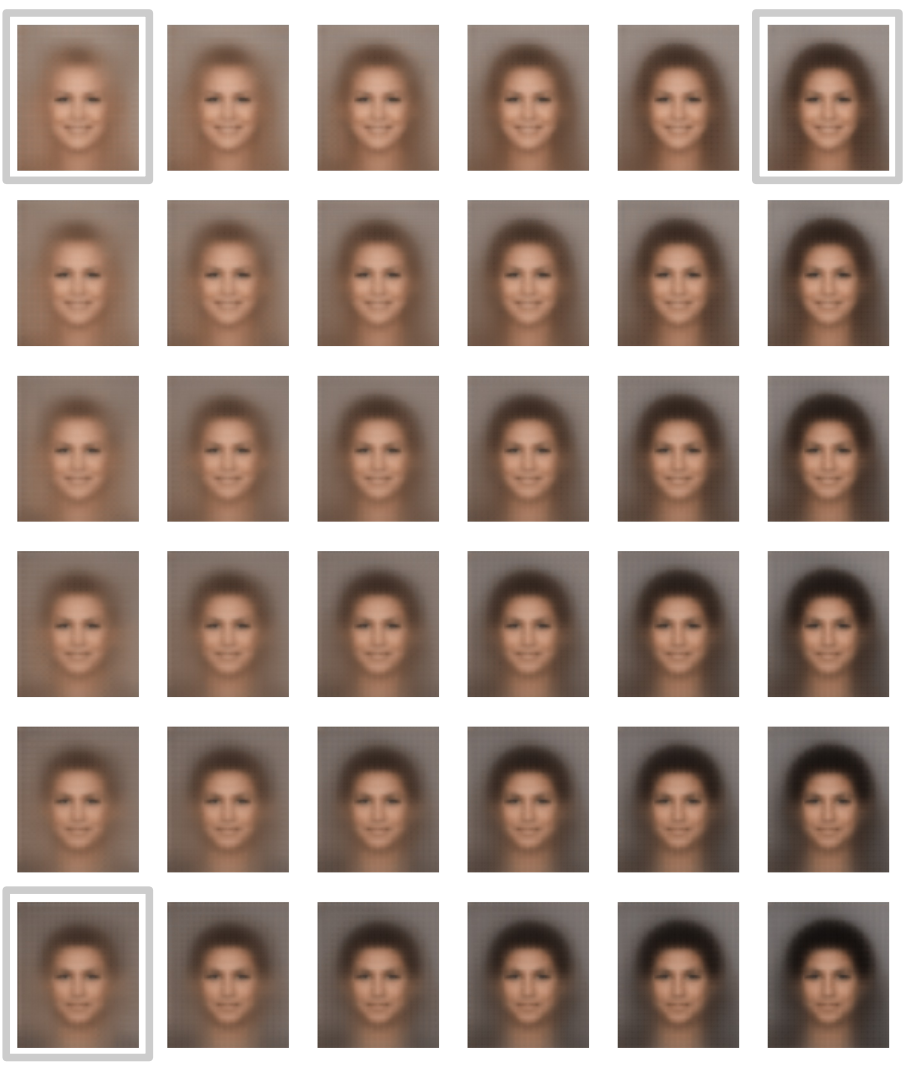

Variational Autoencoder

Use cases :

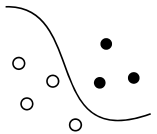

Unsupervised classification

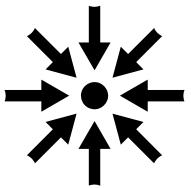

Dimension reduction

Data génération

Image Processing (denoising, ...)

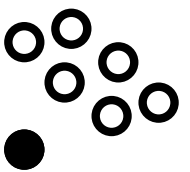

Anomaly detection

...

Notebooks and more…

GitLab

https://gricad-gitlab.univ-grenoble-alpes.fr/talks/fidle

http://bit.ly/fidle432

FIDLE


Merci !