

Deep learning for medical imaging school 2021

Hands-on session

2021

Autoencoders

By

Pierre-Marc Jodoin, Nathan Painchaud

With the support of

Thomas Grenier, Olivier Bernard, Marco A. Armenta, Thierry Judge.

Setup your Floyd hub environment

First, login on floydhub.com

fastest way to build, train, and deploy
deep learning models



You focus on the **science**

We'll handle the **infrastructure**



[Sign Up for Free](#)

[Team? Talk to us](#)



Filter by team... ▾

Filter by job state... ▾

Filter by tags...

New project

New dataset

New team

No jobs found

Please **select your project (you should normally have one)**. You may skip the next page.

If you do not already have a project, **please create one**.





Create a new project

A project contains all your jobs, workspaces, and APIs associated with a particular deep learning goal.

i You can start a Project with a [Template](#) preconfigured for a specific deep learning task



Owner

ge-insa-lyon ▾

Project name

autoencoderproject

Description (optional)



Visibility

Private Private

Create project



1. Give your project a **name**
2. Select the right **owner**
3. Click "**Create project**"



Find projects, datasets, and people.

Jobs

Workspaces

Projects

Datasets

Templates



Help



ge-insa-lyon / projects / autoencoderproject Private

★ Star 0

Overview

Workspaces

Jobs

Model API

Settings

Get started

Choose an option below to start working on your project

Run a Jupyter notebook

- JupyterLab IDE to create and run Jupyter notebooks
- One-click access to GPUs
- Start and stop when you need
- Terminal access to run scripts

 Create workspace



Train a model

- Kick off long-running training scripts
- Choose a suitable machine and framework
- View real-time logs and metrics

 Run a job

Create a model API

- Deploy trained models as a REST API
- Scale by adding more instances
- Monitor real-time API metrics for your model

Deploy model API

Now create a **workspace**



Find projects, datasets, and people.

Jobs Workspaces Projects Datasets Templates

+ Help [Icon]

ge-insa-lyon / projects / autoencoderproject Private

★ Star 0

Overview Workspaces Jobs Model API Settings

0 workspaces

Create Workspace

No workspaces found.

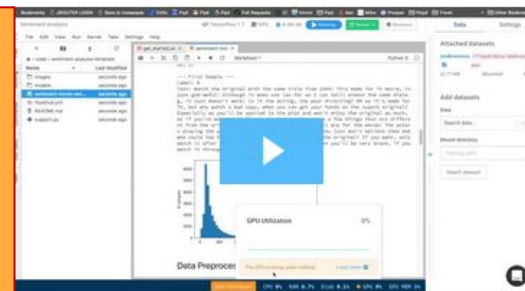


Build models with Workspaces

Workspaces are configurable, interactive development environments built for deep learning and machine learning.

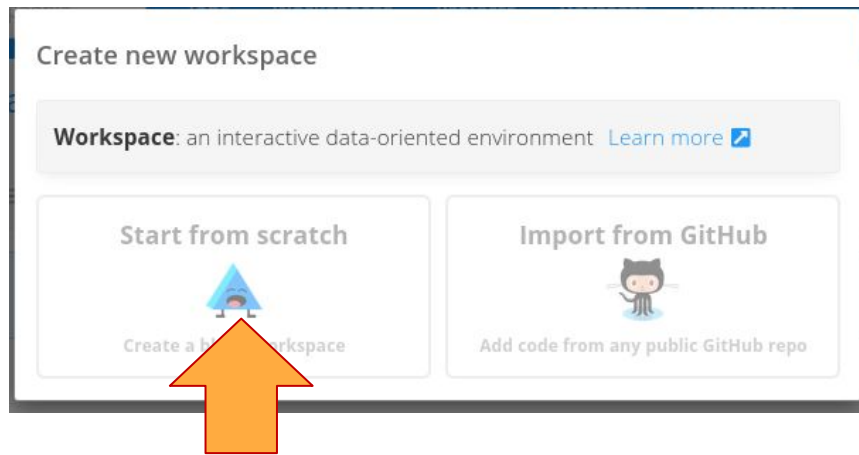
- Create and run Jupyter notebooks
- Toggle between GPU and CPU-powered machines
- Attach datasets to your workspace
- Terminal access to run scripts
- Start and stop when you need

Click on **“Create workspace”**

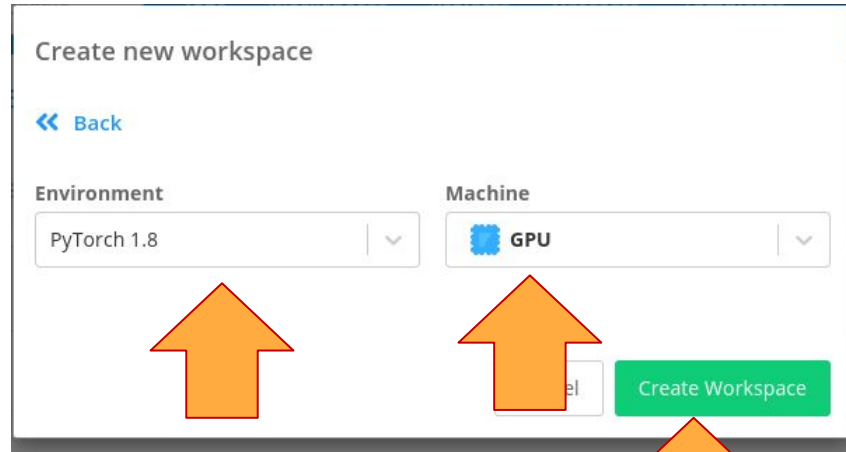


Explore more features of Workspaces

Select **“Start from scratch”**



Then **“PyTorch 1.8”**, **“GPU”**
and click **“Create
Workspace”**





Find projects, datasets, and people.

Jobs

Workspaces

Projects

Datasets

Templates



Help



ge-insa-lyon / projects / autoencoderproject Private

★ Star 0

Overview

Workspaces 1

Jobs

Model API

Settings

1 workspaces

Create Workspace

▶ Running

⊖ Shutdown

autoencoderproject workspace

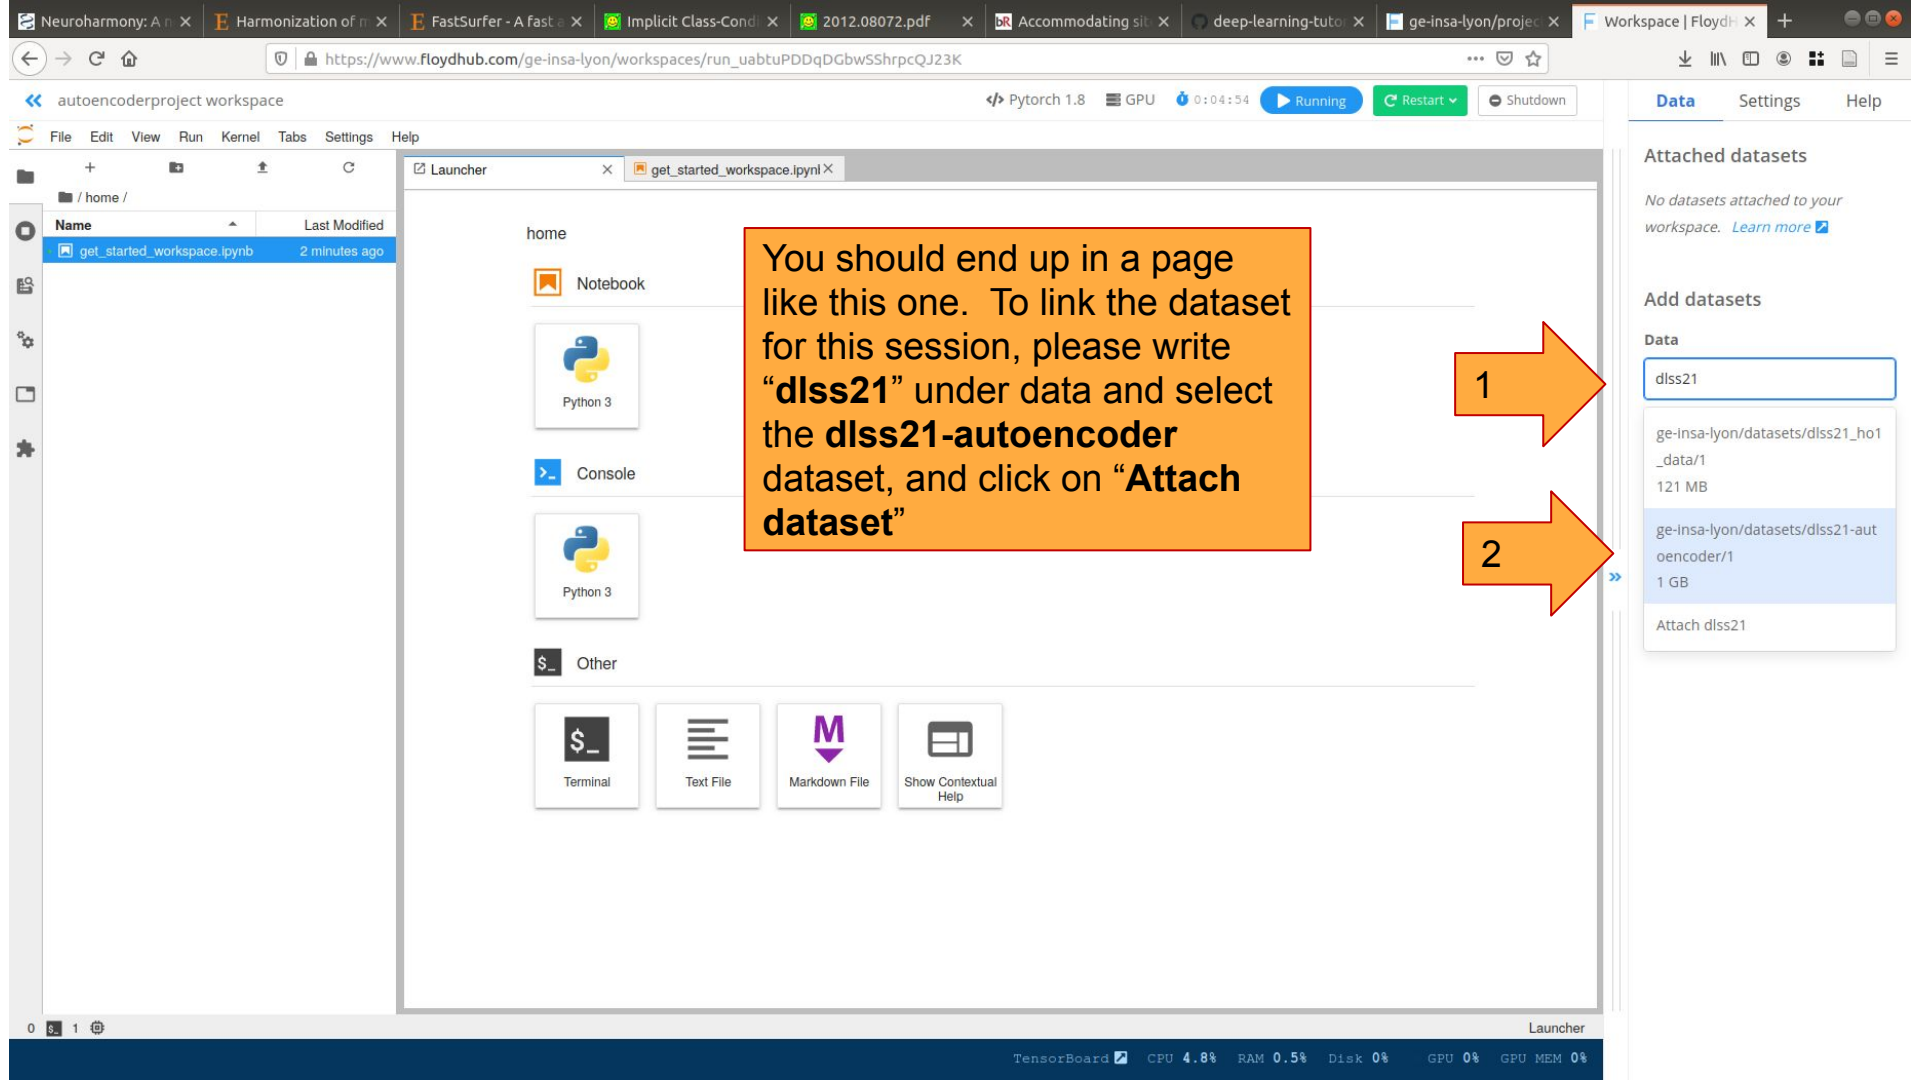
submitted 21 seconds ago

Pytorch 1.8



GPU



You should see your new workspace. So **click on it.**



You should end up in a page like this one. To link the dataset for this session, please write “**dlss21**” under data and select the **dlss21-autoencoder** dataset, and click on “**Attach dataset**”



Attached datasets

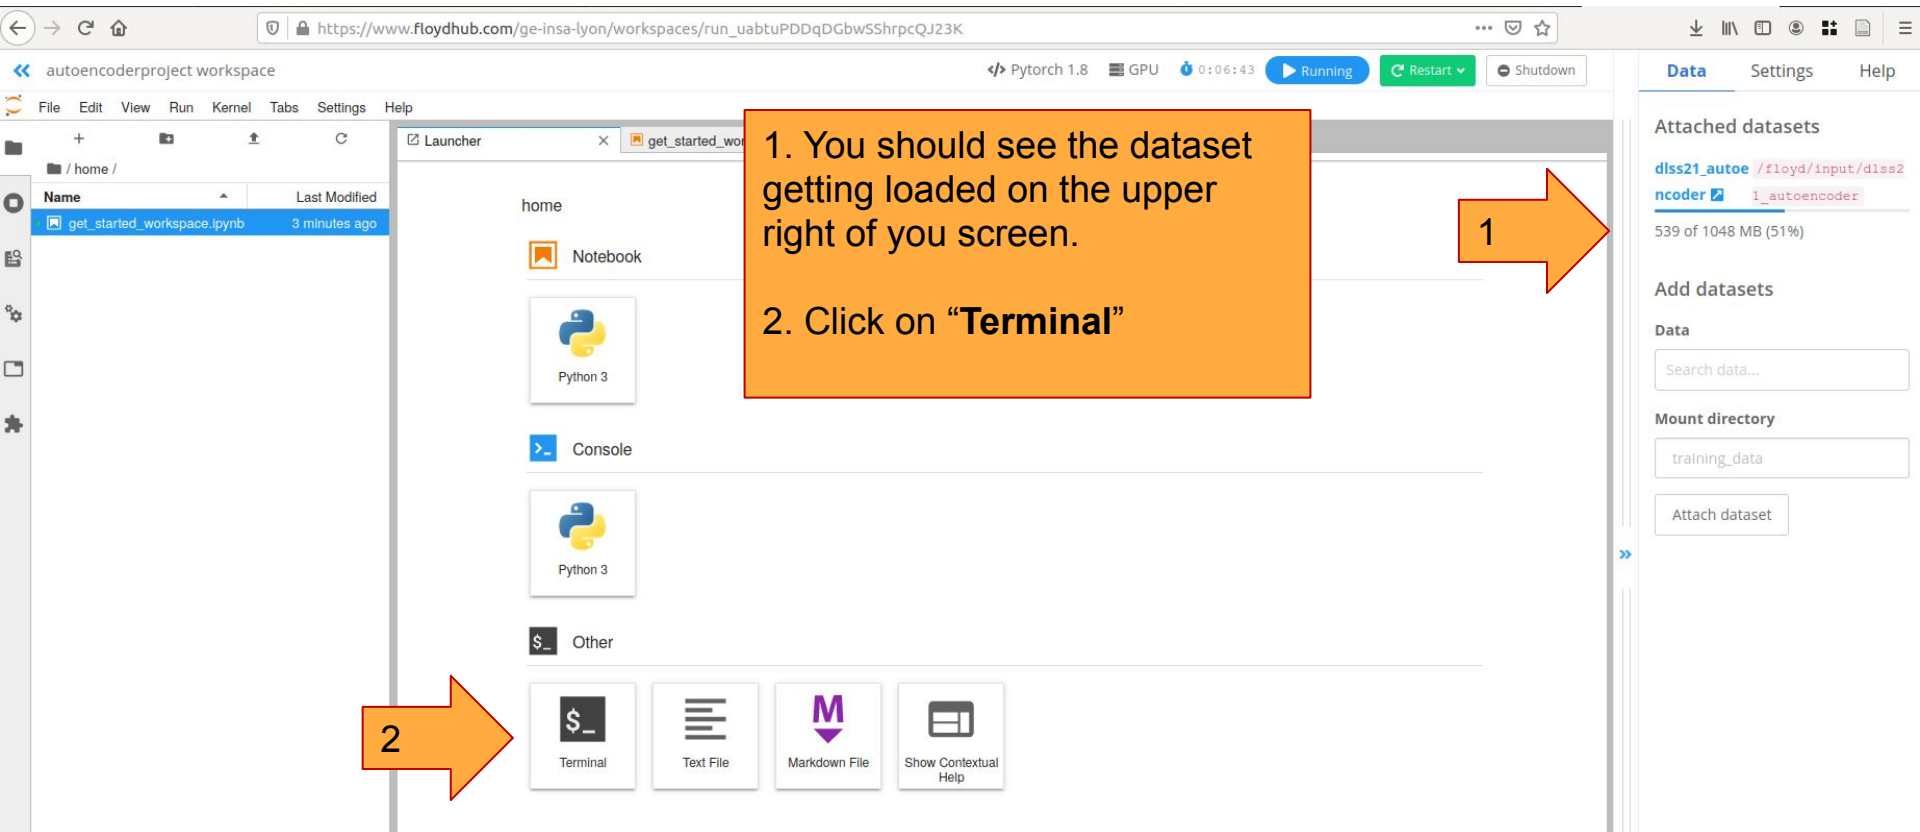
No datasets attached to your workspace. [Learn more](#)

Add datasets

Data


ge-insa-lyon/datasets/dlss21_ho1_data/1	121 MB
ge-insa-lyon/datasets/dlss21-autoencoder/1	1 GB

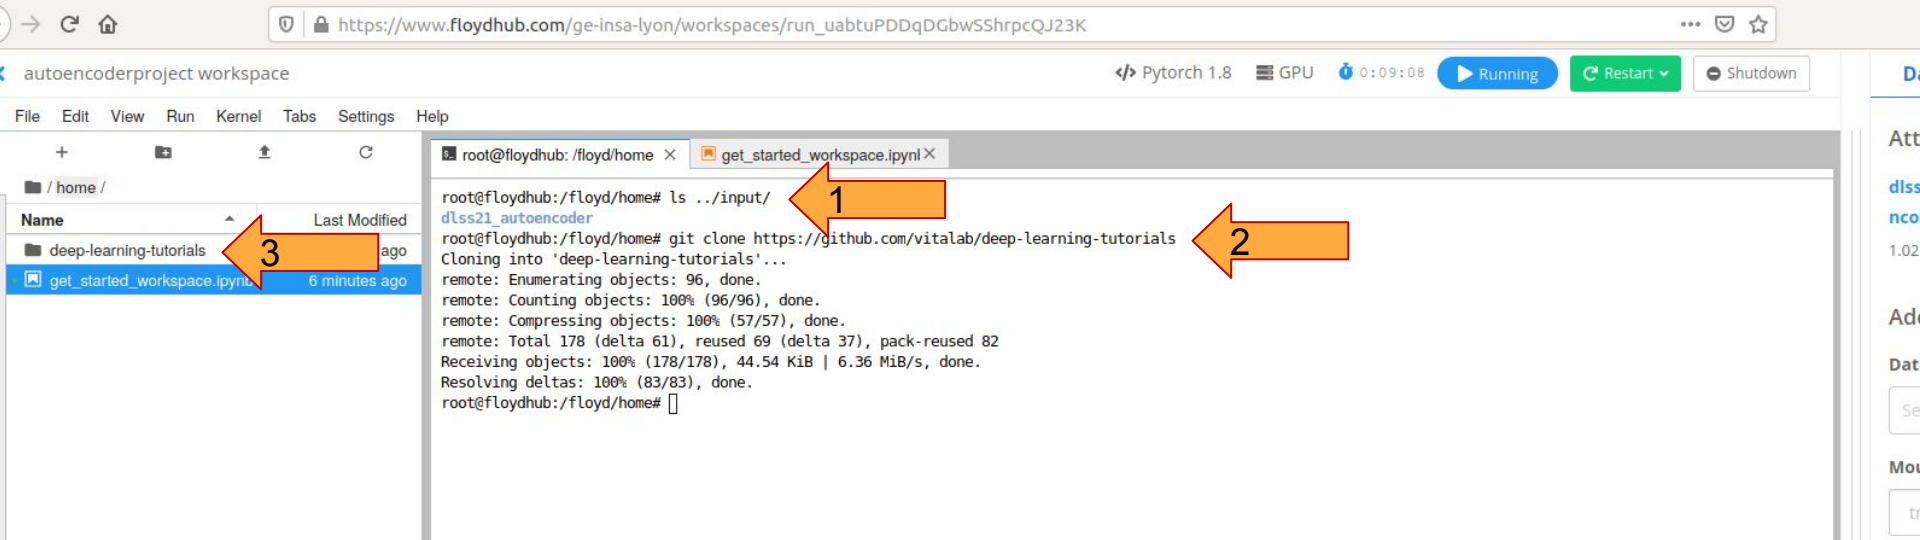
Attach dlss21



1. You should see the dataset getting loaded on the upper right of your screen.

2. Click on **“Terminal”**





1. Type `ls ../input/` in the terminal. You should see the autoencoder dataset.

2. Type `git clone https://github.com/vitalab/deep-learning-tutorials` to download the code

3. A new `deep-learning-tutorials` folder should appear. Click on that folder

📁 / home / deep-learning-tutorials /

Name	Last Modified
src	seconds ago
tutorials	seconds ago
dev.txt	seconds ago
environment.yml	seconds ago
LICENSE	seconds ago
pyproject.toml	seconds ago
README.md	seconds ago
requirements.txt	seconds ago
setup.cfg	seconds ago
setup.py	seconds ago

```
root@floydhub: /floyd/home# ls ../input/  
dls21_autoencoder  
root@floydhub: /floyd/home# git clone https://github.com/vitalab/deep-learning-tutorials  
Cloning into 'deep-learning-tutorials'...  
remote: Enumerating objects: 96, done.  
remote: Counting objects: 100% (96/96), done.  
remote: Compressing objects: 100% (57/57), done.  
remote: Total 178 (delta 61), reused 69 (delta 37), pack-reused 82  
Receiving objects: 100% (178/178), 44.54 KiB | 6.36 MiB/s, done.  
Resolving deltas: 100% (83/83), done.  
root@floydhub: /floyd/home#
```



This is the code for this hands-on session.

Click on **“tutorials”**

At
dls
ncc
1.0:
Ad
Dal
Si
Mo
t
/


The screenshot shows a JupyterLab interface. The top navigation bar includes 'File', 'Edit', 'View', 'Run', 'Kernel', 'Tabs', 'Settings', and 'Help'. The main area is split into a file browser on the left and a terminal window on the right. The file browser shows a directory structure with 'mnist-autoencoders.ipynb' highlighted in blue, and a red arrow pointing to it from the right. The terminal window shows the following output:

```
root@floydhub: /floyd/home# ls ../input/  
dlss21_autoencoder  
root@floydhub: /floyd/home# git clone https://github.com/vitalab/deep-learning-tutorials  
Cloning into 'deep-learning-tutorials'...  
remote: Enumerating objects: 96, done.  
remote: Counting objects: 100% (96/96), done.  
remote: Compressing objects: 100% (57/57), done.  
remote: Total 178 (delta 61), reused 69 (delta 37), pack-reused 82  
Receiving objects: 100% (178/178), 44.54 KiB | 6.36 MiB/s, done.  
Resolving deltas: 100% (83/83), done.  
root@floydhub: /floyd/home#
```


Here you have 2 notebooks. Start with the **mnist notebook** and once you are done, move to the **cardiac ACDC notebook**

Overview of the Autoencoder hands-on


(very) quick recap




Annotated dataset: $D = \{(\vec{x}_1, t_1), (\vec{x}_2, t_2), \dots, (\vec{x}_N, t_N)\}$

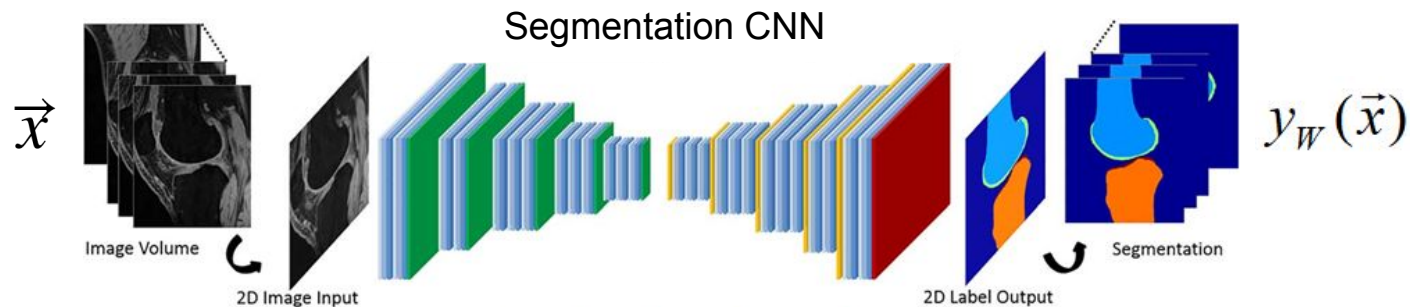




Annotated dataset: $D = \{(\vec{x}_1, t_1), (\vec{x}_2, t_2), \dots, (\vec{x}_N, t_N)\}$




Annotated dataset: $D = \{(\vec{x}_1, t_1), (\vec{x}_2, t_2), \dots, (\vec{x}_N, t_N)\}$






GAN

(Latent space)




sampling

\vec{z}



$G(\vec{z})$

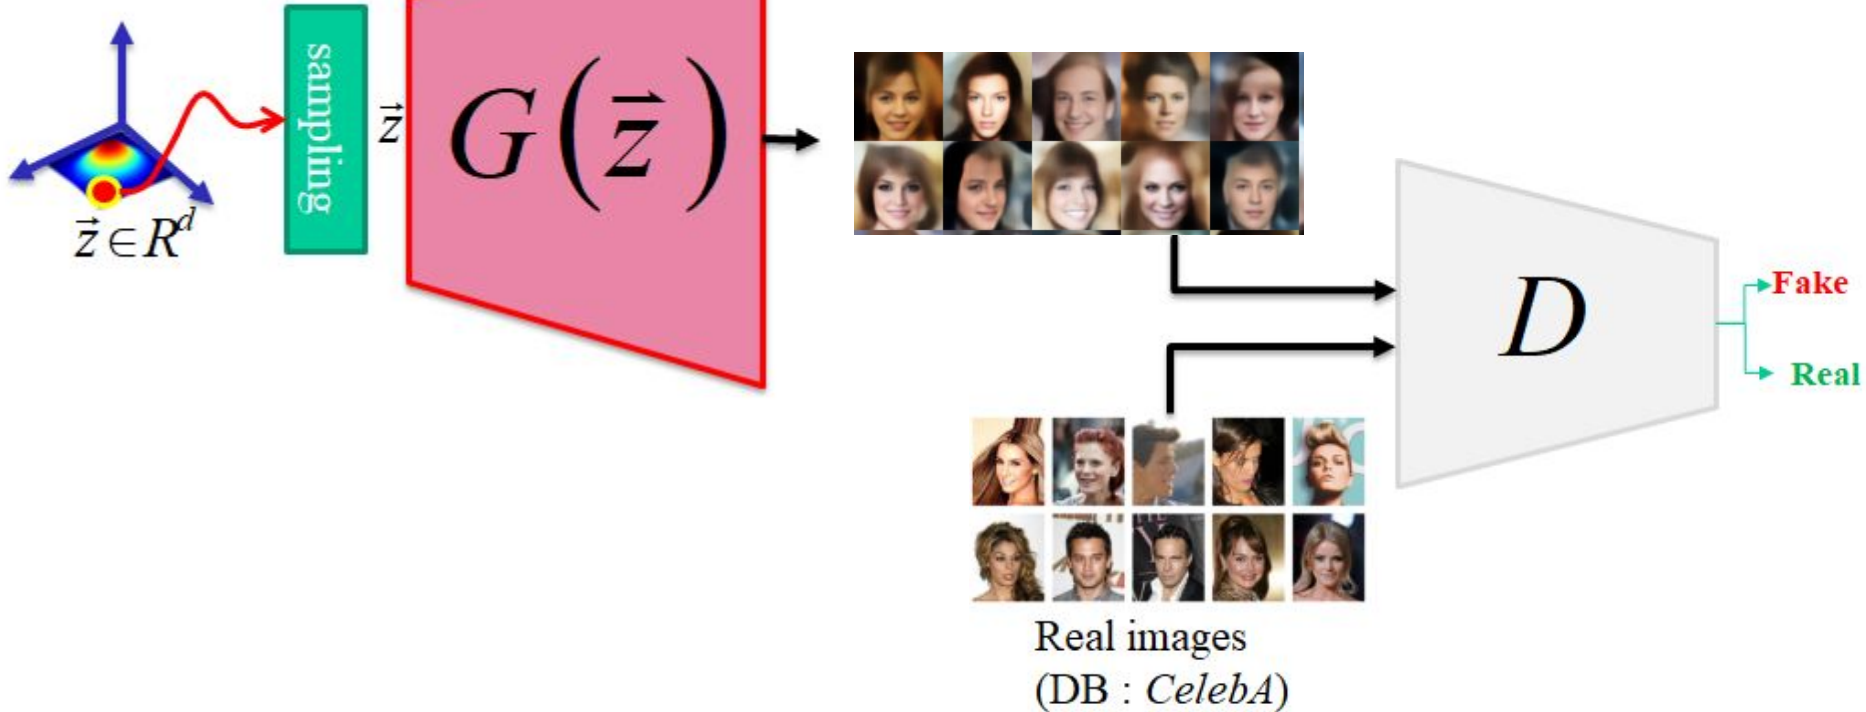


Loss?


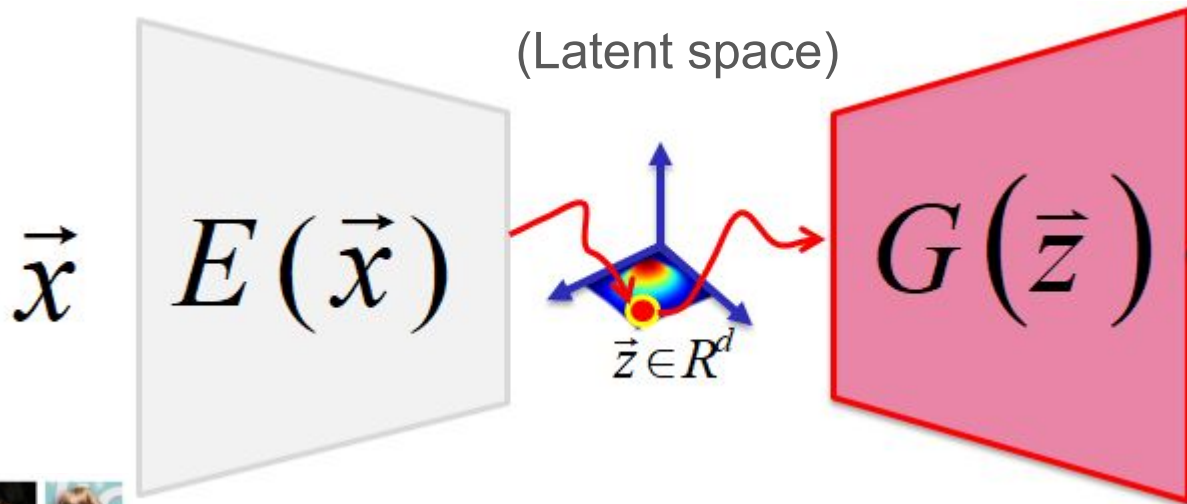
Annotated dataset?

GAN

(Latent space)




Autoencoders



Autoencoders (once training is over)

(Latent space)




Summary

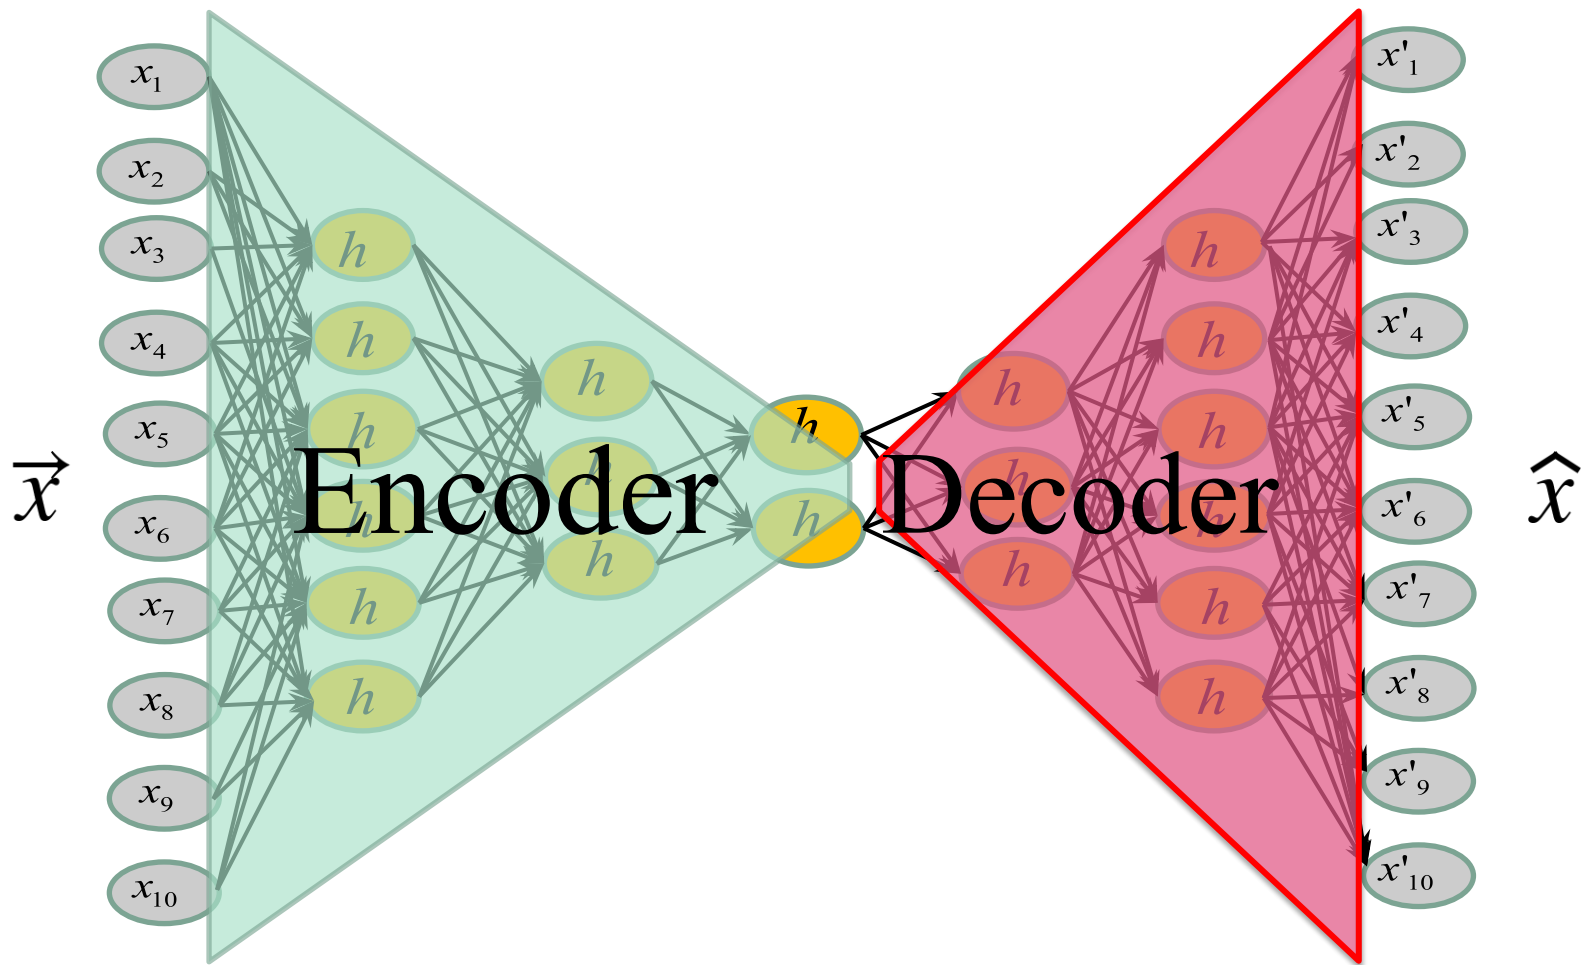
1. What are autoencoders and variational autoencoders?
2. How do they apply to MNIST (grayscale images)?
3. How do they apply to segmentation maps (ACDC cardiac labels)?

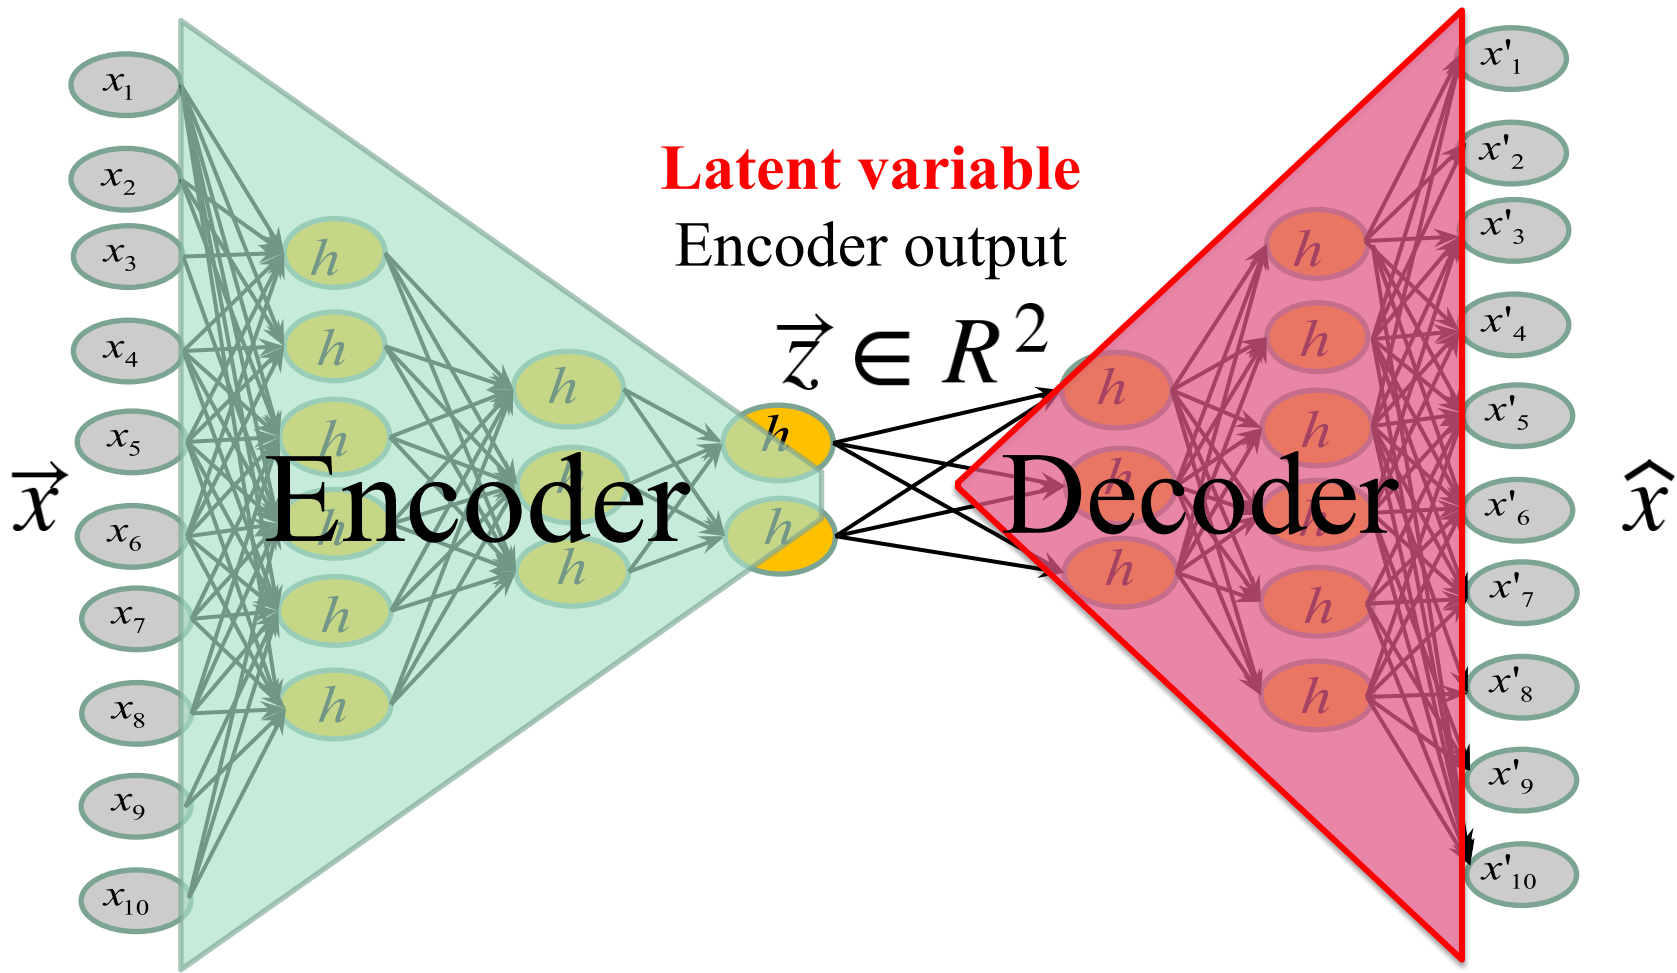
Goal : learn the latent representation of a set of data


How : by training a Neural Net to output its own ...
input!

Note : if you are familiar with AE and VAE, you may skip the next couple of slides.











Fully-Connected Layers

h : activation function




Very often...

No activation function at the output
of the encoder and the decoder



See **why?**


Sometimes **sigmoid** to predict pixel values between 0 and 1 or **ReLU** when the pixel values can be large but never negative.



The number of neurons

Decrease (or stay the same)


from a layer to another



The number of neurons


Increase (or stay the same)

from a layer to another



Very often...

The structure of the encoder is the **dual** of that of the decoder



Simple MNIST Autoencoder

```
class autoencoder(nn.Module):  
    def __init__(self):  
        super(autoencoder, self).__init__()  
        self.encoder = nn.Sequential(  
            nn.Linear(28 * 28, 128), nn.ReLU(True),  
            nn.Linear(128, 64), nn.ReLU(True),  
            nn.Linear(64, 12), nn.ReLU(True),  
            nn.Linear(12, 2))  
        self.decoder = nn.Sequential(  
            nn.Linear(2, 12), nn.ReLU(True),  
            nn.Linear(12, 64), nn.ReLU(True),  
            nn.Linear(64, 128), nn.ReLU(True),  
            nn.Linear(128, 28 * 28))  
  
    def forward(self, x):  
        z = self.encoder(x)  
        x_prime = self.decoder(z)  
        return x_prime
```

Latent space 2D



Simple MNIST Autoencoder


```
class autoencoder(nn.Module):  
    def __init__(self):  
        super(autoencoder, self).__init__()  
        self.encoder = nn.Sequential(  
            nn.Linear(28 * 28, 128), nn.ReLU(True),  
            nn.Linear(128, 64), nn.ReLU(True),  
            nn.Linear(64, 12), nn.ReLU(True),  
            nn.Linear(12, 2))  
        self.decoder = nn.Sequential(  
            nn.Linear(2, 12), nn.ReLU(True),  
            nn.Linear(12, 64), nn.ReLU(True),  
            nn.Linear(64, 128), nn.ReLU(True),  
            nn.Linear(128, 28 * 28))  
  
    def forward(self, x):  
        z = self.encoder(x)  
        x_prime = self.decoder(z)  
        return x_prime
```

symmetry




MNIST latent space (for 1000 images)

Each 2D point corresponds to an image




Conv layers





Basic image-based autoencoder



\vec{x}




$\vec{z} \in R^d$

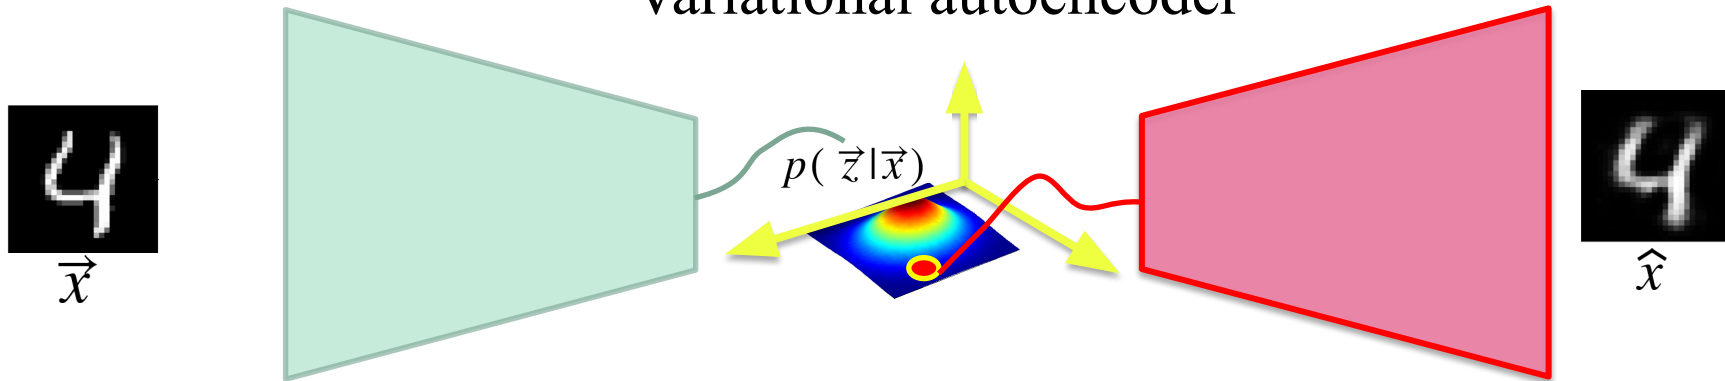


\hat{x}

Basic autoencoder




Variational autoencoder



Variational autoencoder



\vec{x}




The encoder outputs a **distribution** $p(\vec{z}|\vec{x})$
and not just a **vector** \vec{z}

Variational autoencoder




\vec{x}




Random sample $\vec{z} \sim P(\vec{z}|\vec{x})$

Variational autoencoder




...and rebuild \hat{x}

$$p(\vec{z}|\vec{x}) \sim \text{Gaussian}$$



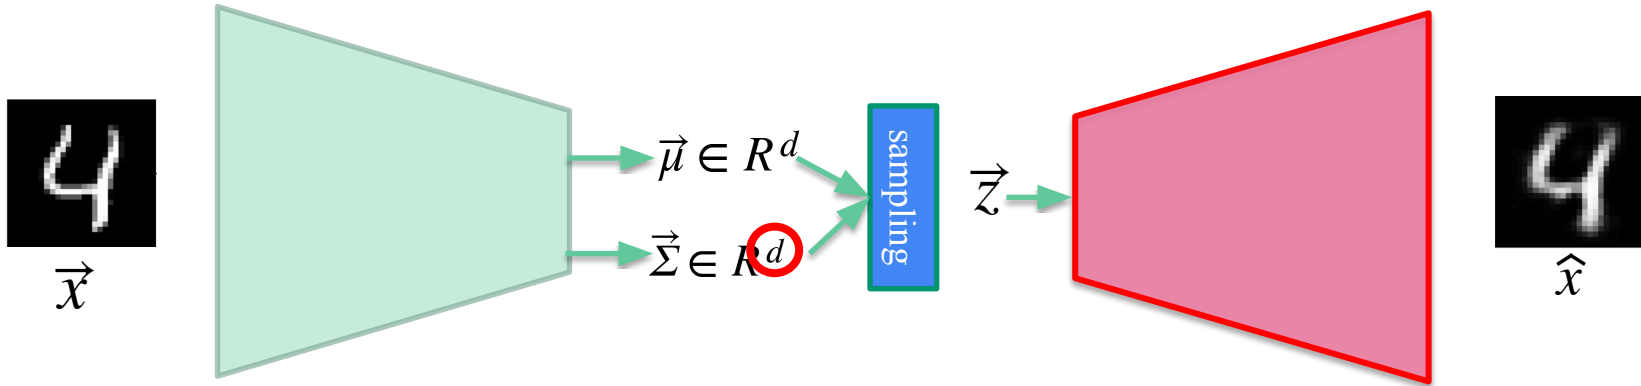
$$p(\vec{z}|\vec{x}) \sim \text{Gaussian}$$



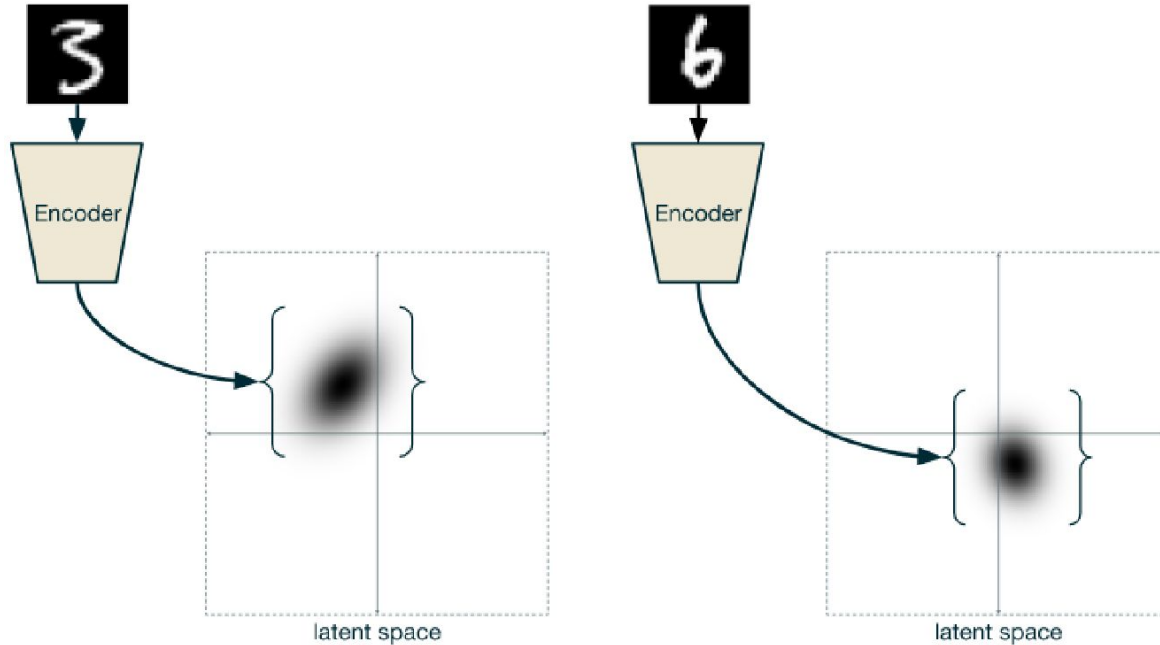
To simplify things, we assume Σ is **diagonal**

$$\Sigma = \begin{pmatrix} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_1^2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \sigma_d^2 \end{pmatrix}$$


Variational autoencoder




Another way of seeing things...



Another way of seeing things...




ELBO loss : Evidence Lower Bound loss



$$Loss = \underbrace{(\hat{x} - \hat{\hat{x}})^2}_{\text{Loss decoder}} + \lambda \underbrace{KL(\mathcal{N}(\vec{z}; \vec{0}, \vec{1}), \mathcal{N}(\vec{z}; \vec{\mu}, \vec{\Sigma}))}_{\text{Loss encoder}}$$

Other loss (in case the output is binary)




$$Loss = \underbrace{CrossEntropy(\vec{x}, \hat{x})}_{\text{Loss decoder}} + \lambda \underbrace{KL(\mathcal{N}(\vec{z}; \vec{0}, \vec{1}), \mathcal{N}(\vec{z}; \vec{\mu}, \vec{\Sigma}))}_{\text{Loss encoder}}$$

Loss decoder


Loss encoder

Reparametrization trick instead of sampling

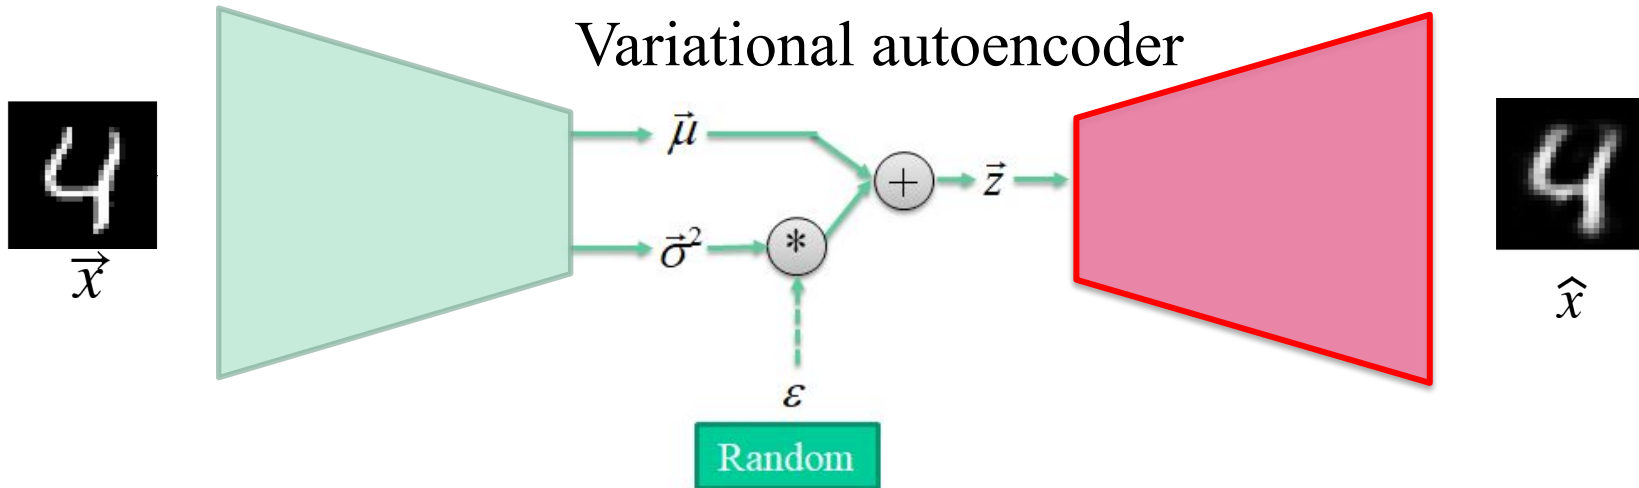


MNIST

Basic autoencoder

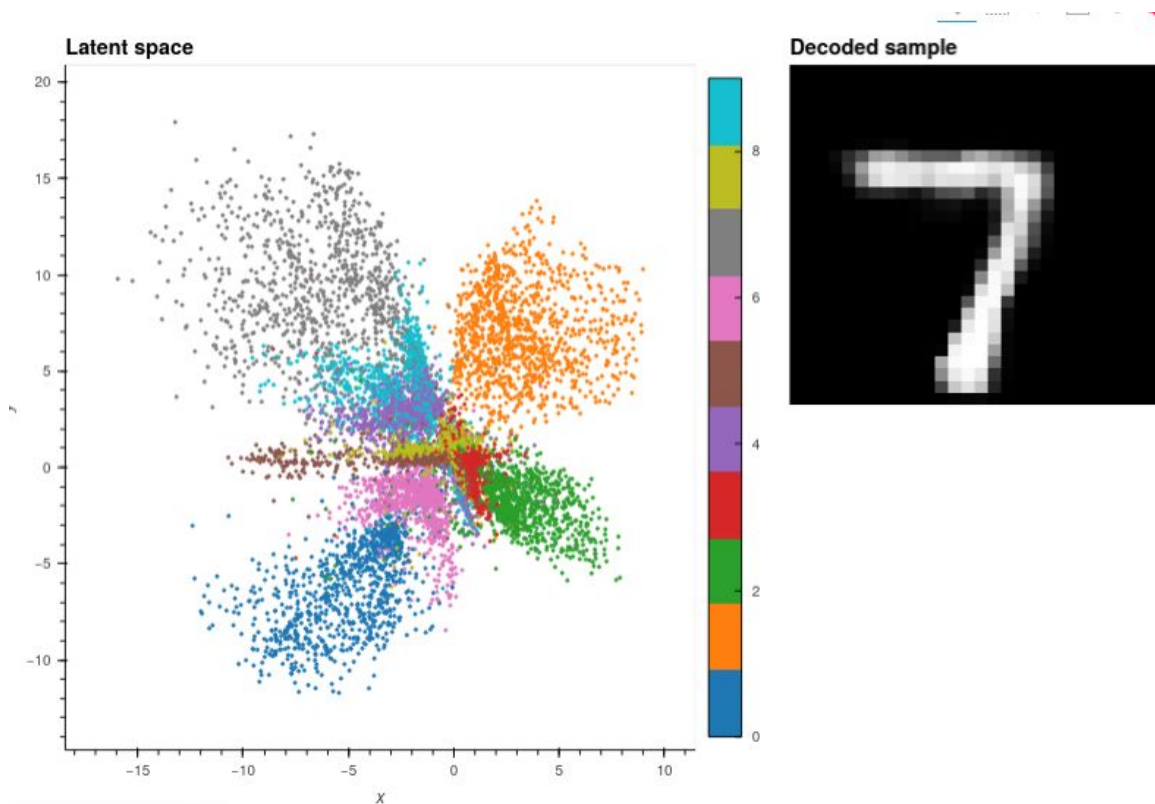


Variational autoencoder




MNIST

Visualize the latent space (autoencoder)



MNIST

Visualize the latent space (variational autoencoder)



MNIST

Code is in the file:

`mnist-autoencoders.ipynb`


Automated Cardiac Diagnosis Challenge



150 exams (all from different patients) divided into 5 groups:

- dilated cardiomyopathy (DCM),
- hypertrophic cardiomyopathy (HCM),
- myocardial infarction (MINF),
- abnormal right ventricle (RV)
- patients without cardiac disease (NOR).

website: creatis.insa-lyon.fr/Challenge/acdc



Overview
Contest
Participation
Databases
Evaluation
Code
Paper Submission
Contact

Overview

[General context](#) [Scientific interests](#) [Organizers](#)

The goal of this contest is two-fold:

- compare the performance of automatic methods on the segmentation of the left ventricular endocardium and epicardium as the right ventricular endocardium for both end diastolic and end systolic phase instances;
- compare the performance of automatic methods for the classification of the examinations in five classes (normal case, heart failure with infarction, dilated cardiomyopathy, hypertrophic cardiomyopathy, abnormal right ventricle).

This will be done using a common database of 3D cine-MRI images acquired from 150 patients (30 per pathology plus 30 healthy subjects) and the associated manual references based on the analysis of a clinical expert.

NEWS

4th of April 2017


The full training dataset (100 patients) is available

13th of March 2017

Challenge accepted by the MICCAI Satellite Committee


10th of January 2017

The online evaluation platform is operational



Cardiac anatomy


Normal Heart



Chambers relax and fill,
then contract and pump.


Cardiac anatomy

Normal Heart




Chambers relax and fill,
then contract and pump.

Cross-section : **short axis view**






Short axis cardiac MRI




Short axis segmentation map






Other examples:



Convolutional Variational autoencoder



ACDC

Code is in the file:

cardiac-mri-autoencoders.ipynb