VARIATIONAL AUTOENCODERS

LUC HENDRIKS
RADBOUD UNIVERSITY, NIJMEGEN (NL)

VARIATIONAL AUTOENCODERS

- Conceptual talk about VAEs
- VAEs as a tool to do:
    - Anomaly / outlier detection
    - Noise reduction
    - Generative modelling
    - Event generation with a density buffer (Sydney's talk)

VARIATIONAL AUTOENCODERS

- Conceptual talk about VAEs
- VAEs as a tool to do:
    - Anomaly / outlier detection
    - Noise reduction
    - Generative modelling
    - Event generation with a density buffer (Sydney's talk)
- Topics
    - Normal AEs
    - The concept of latent spaces
    - VAEs
    - β-VAEs

AUTOENCODERS

- Class of deep learning algorithms
- Output = input
- Unsupervised learning (no labels needed)

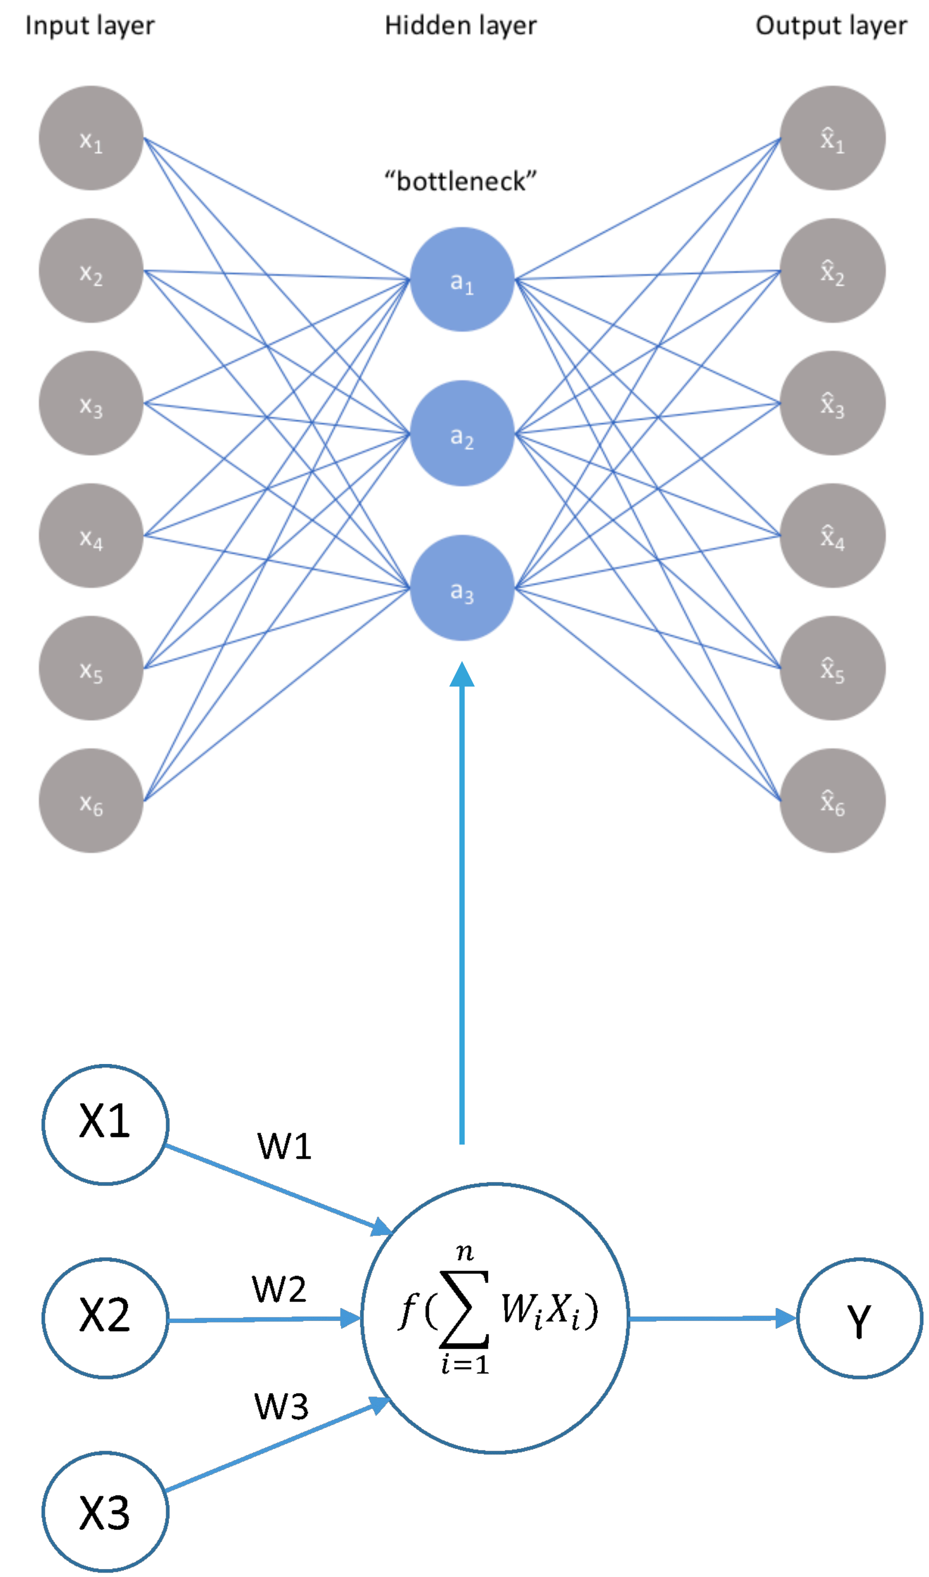

AUTOENCODERS

- Class of deep learning algorithms
- Output = input
- Unsupervised learning (no labels needed)

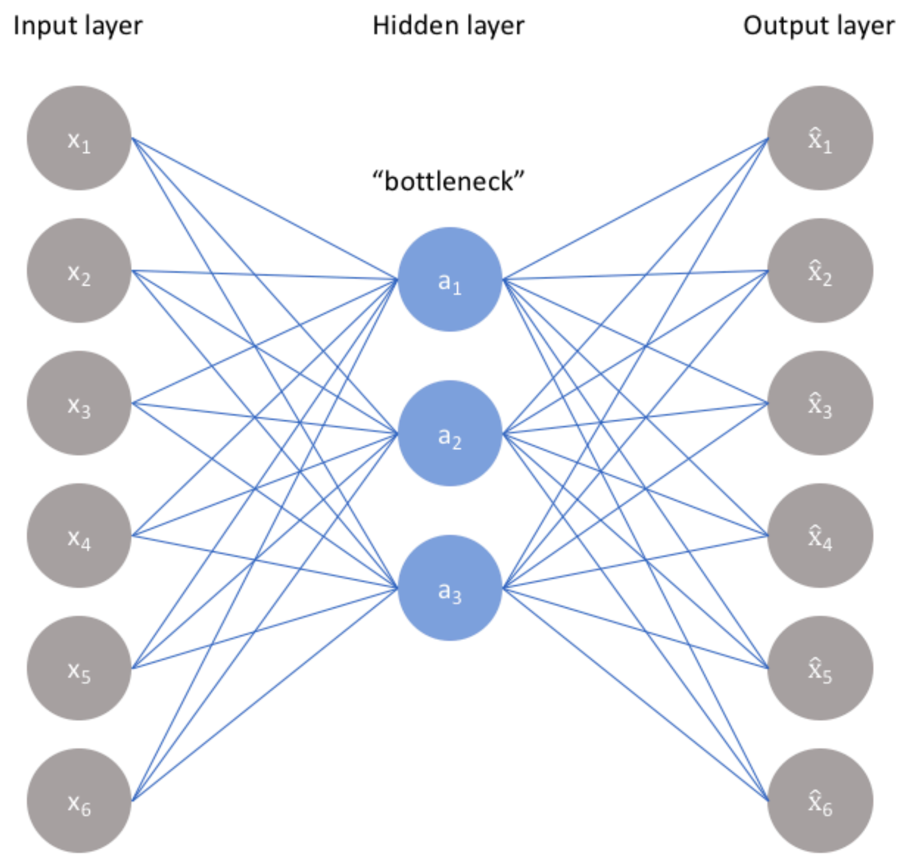

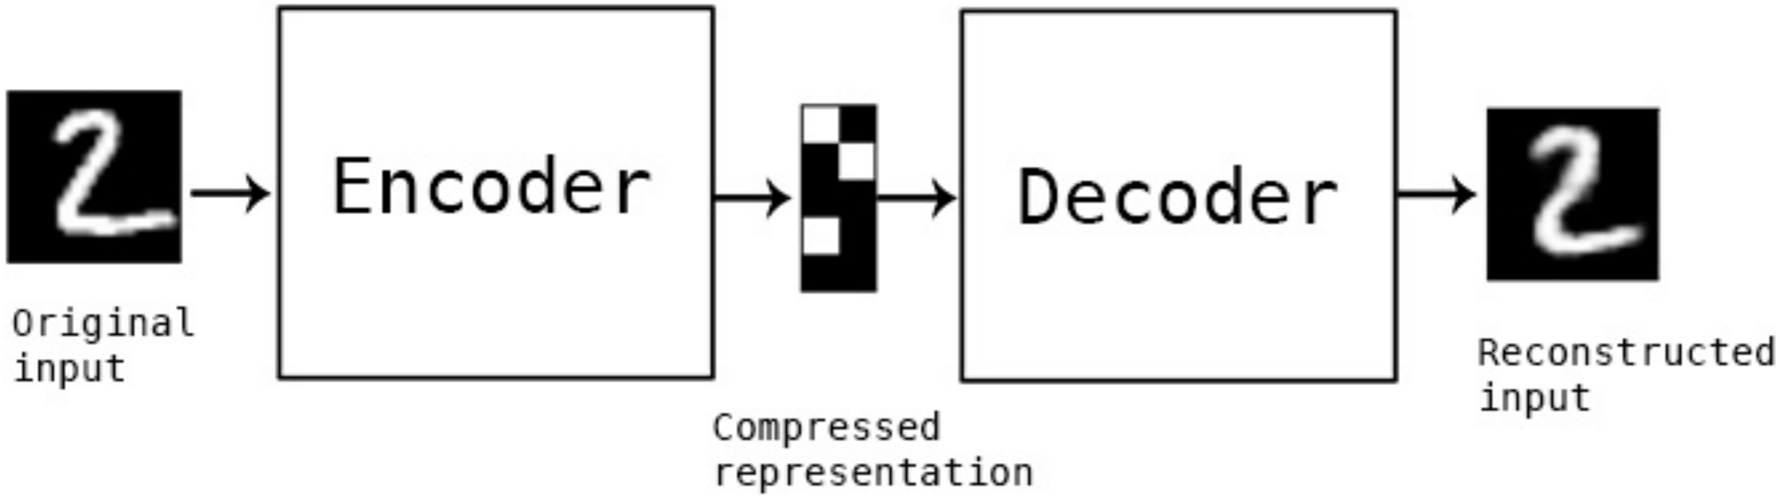

AUTOENCODERS

- Class of deep learning algorithms
- Output = input
- Unsupervised learning (no labels needed)

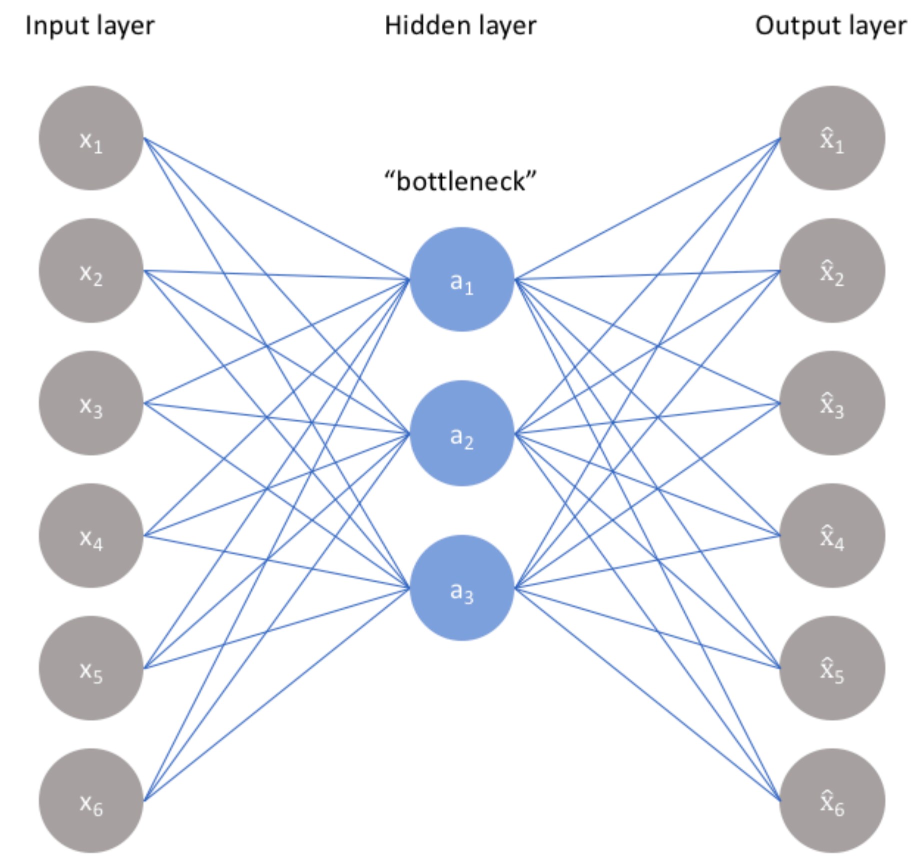

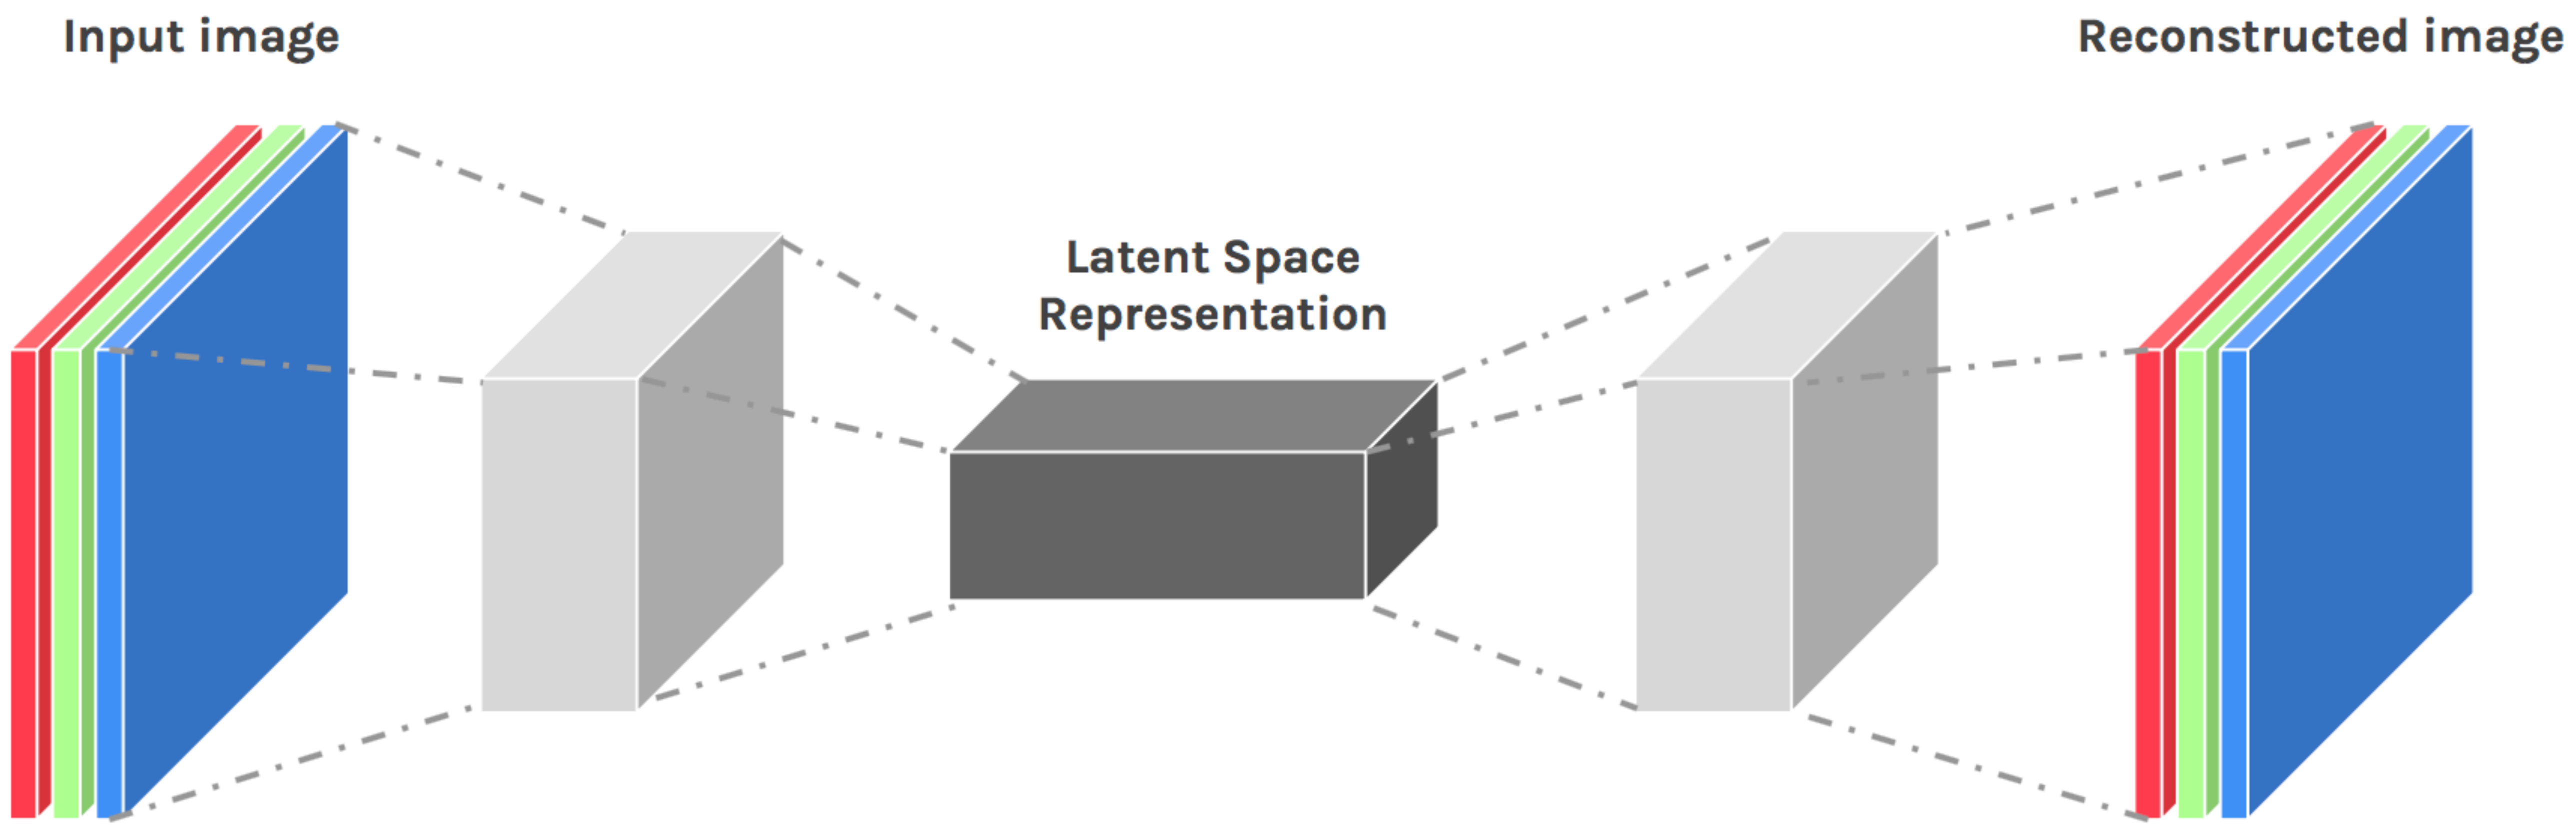

AUTOENCODERS

- Reconstruction very good –> compression algorithm
- Noise reduction
- Outlier detection:
    - Put in something that the AE never saw –> bad reconstruction
    - Reconstruction loss = variable for outlier detection

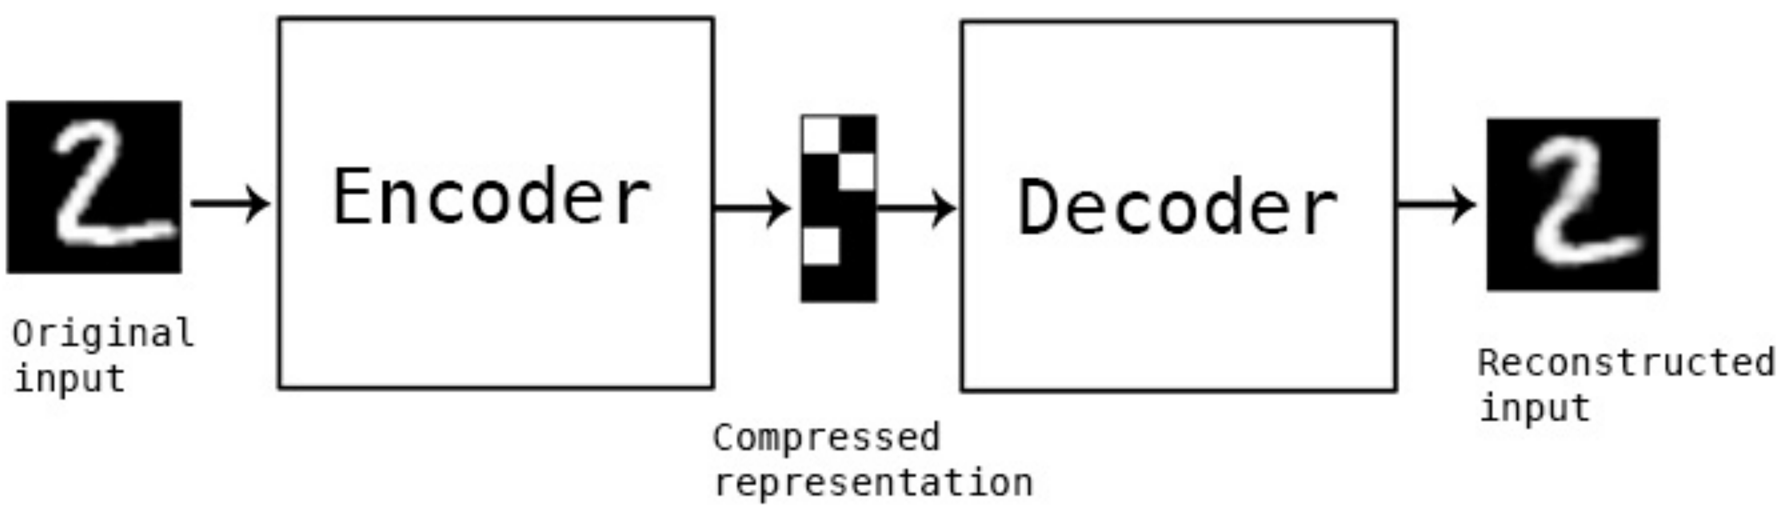

- Outlier: credit card fraud detection

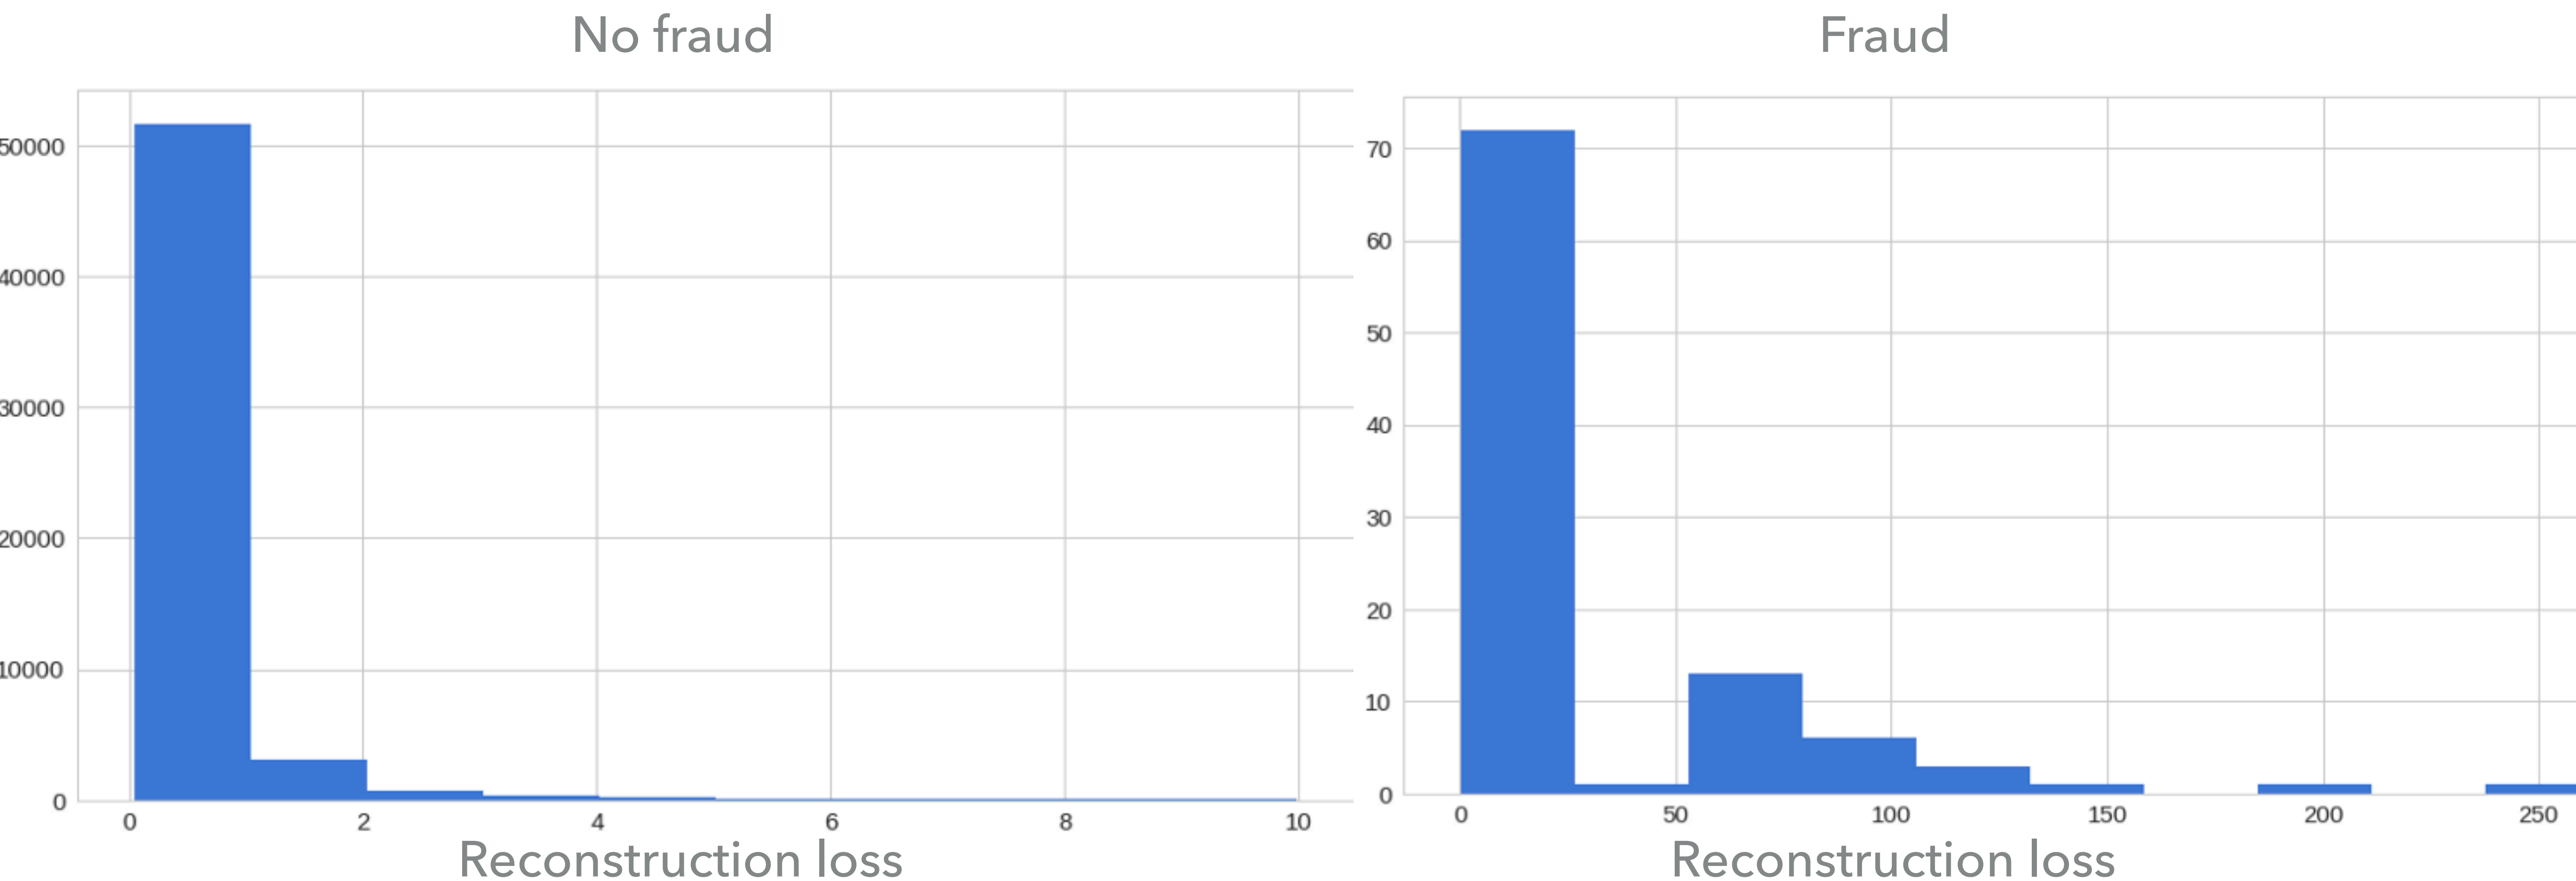

- Outlier: credit card fraud detection

No fraud

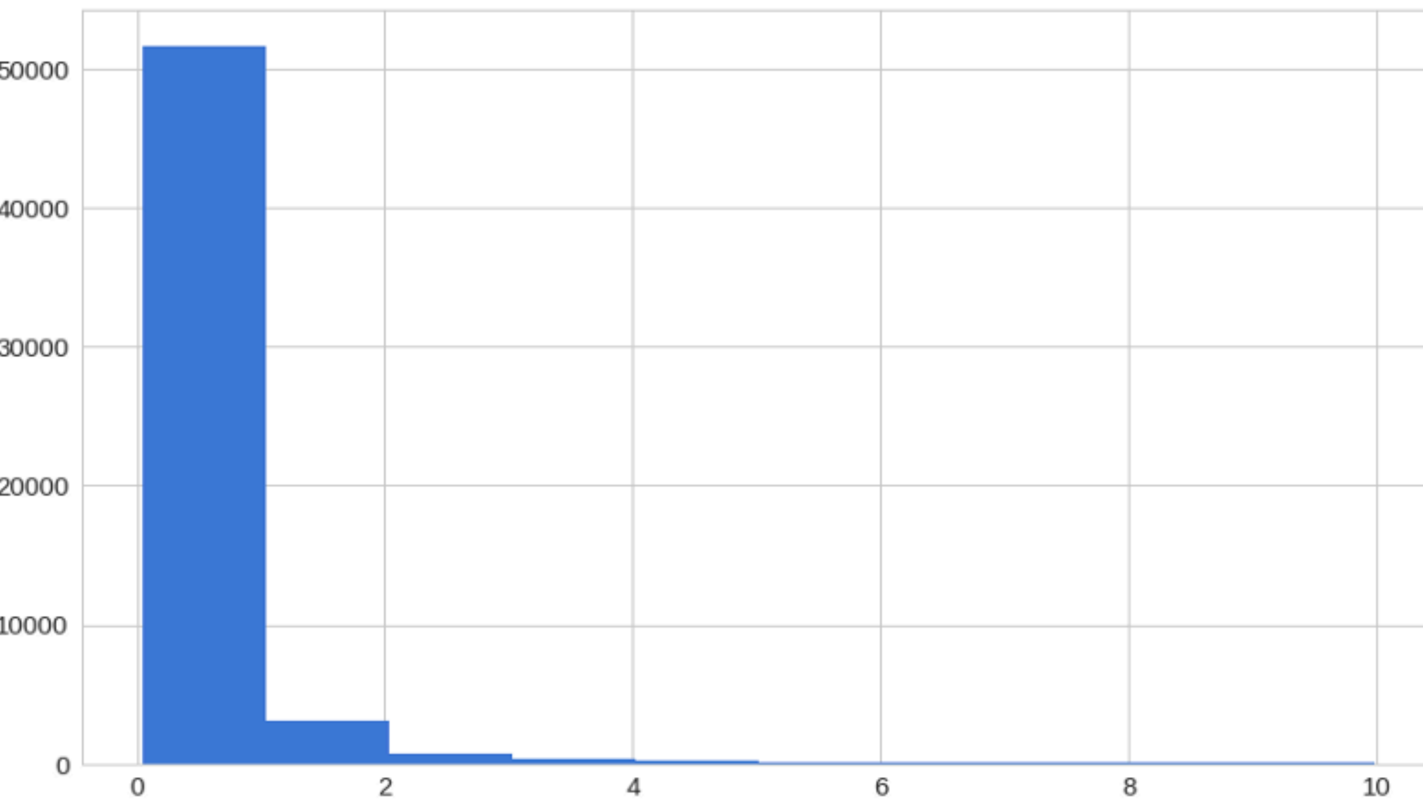

Reconstruction loss

Fraud

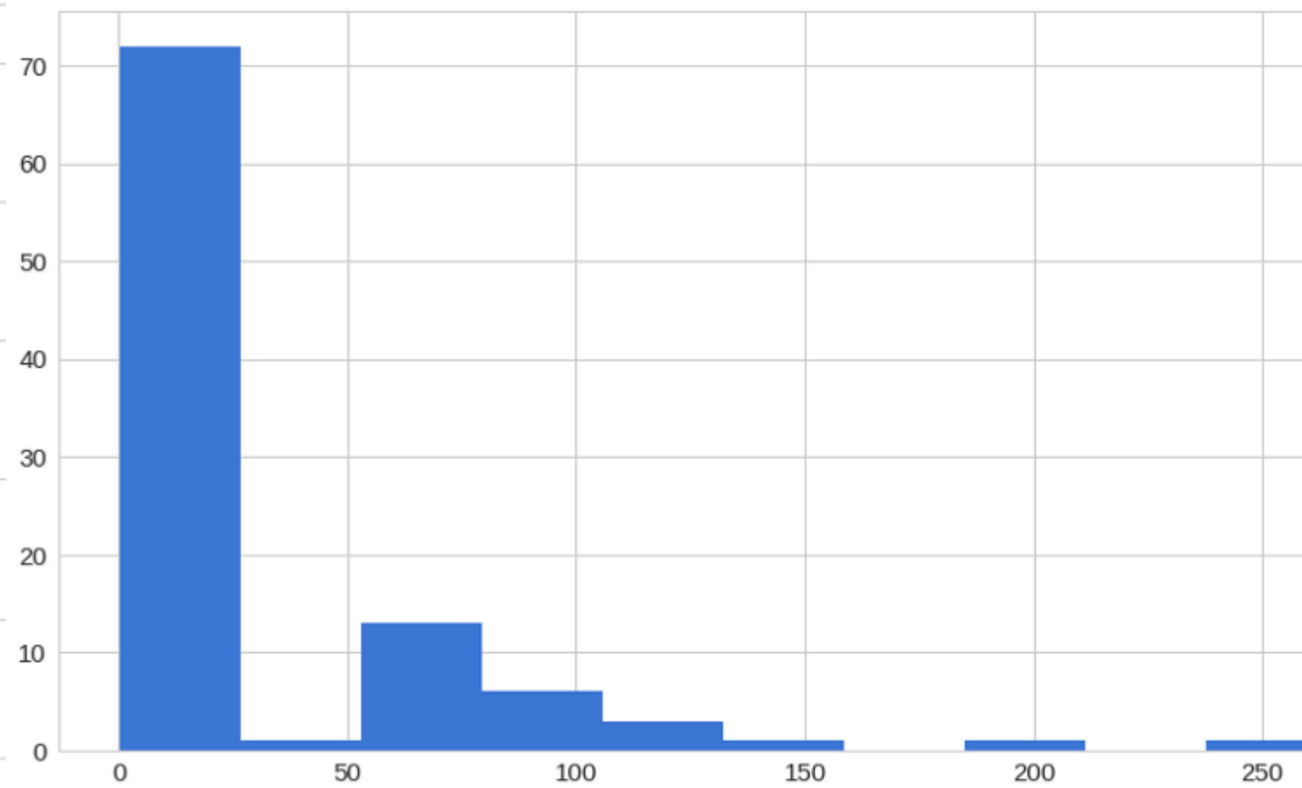

Reconstruction loss

- Noise reduction: MNIST noisy

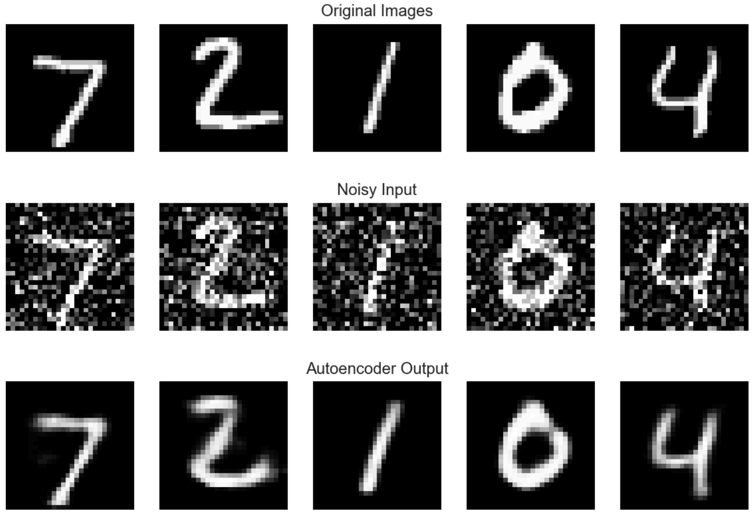

AUTOENCODERS

- No ordering in latent space

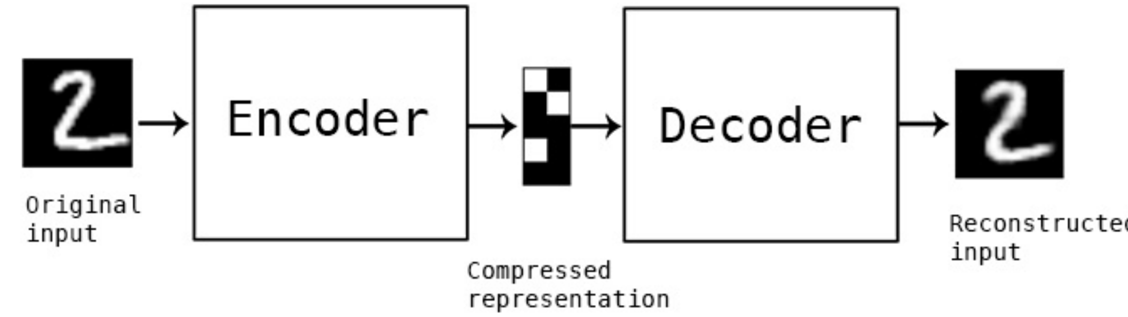

Assume 2D easy viz.

AUTOENCODERS

- No ordering in latent space

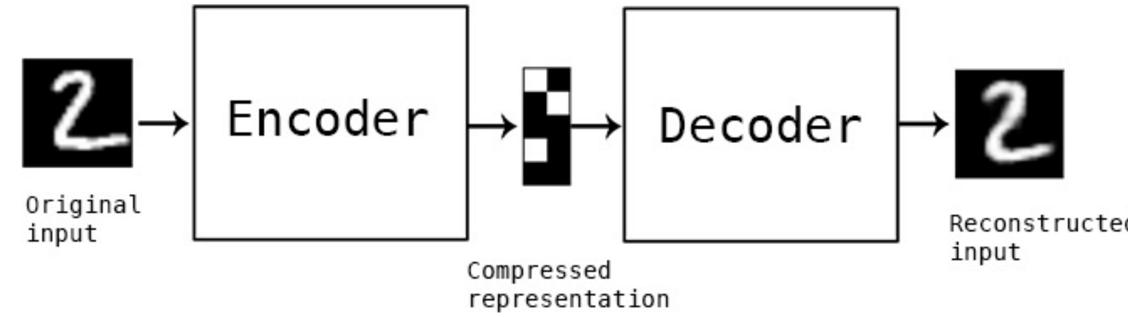

Assume 2D
easy viz.

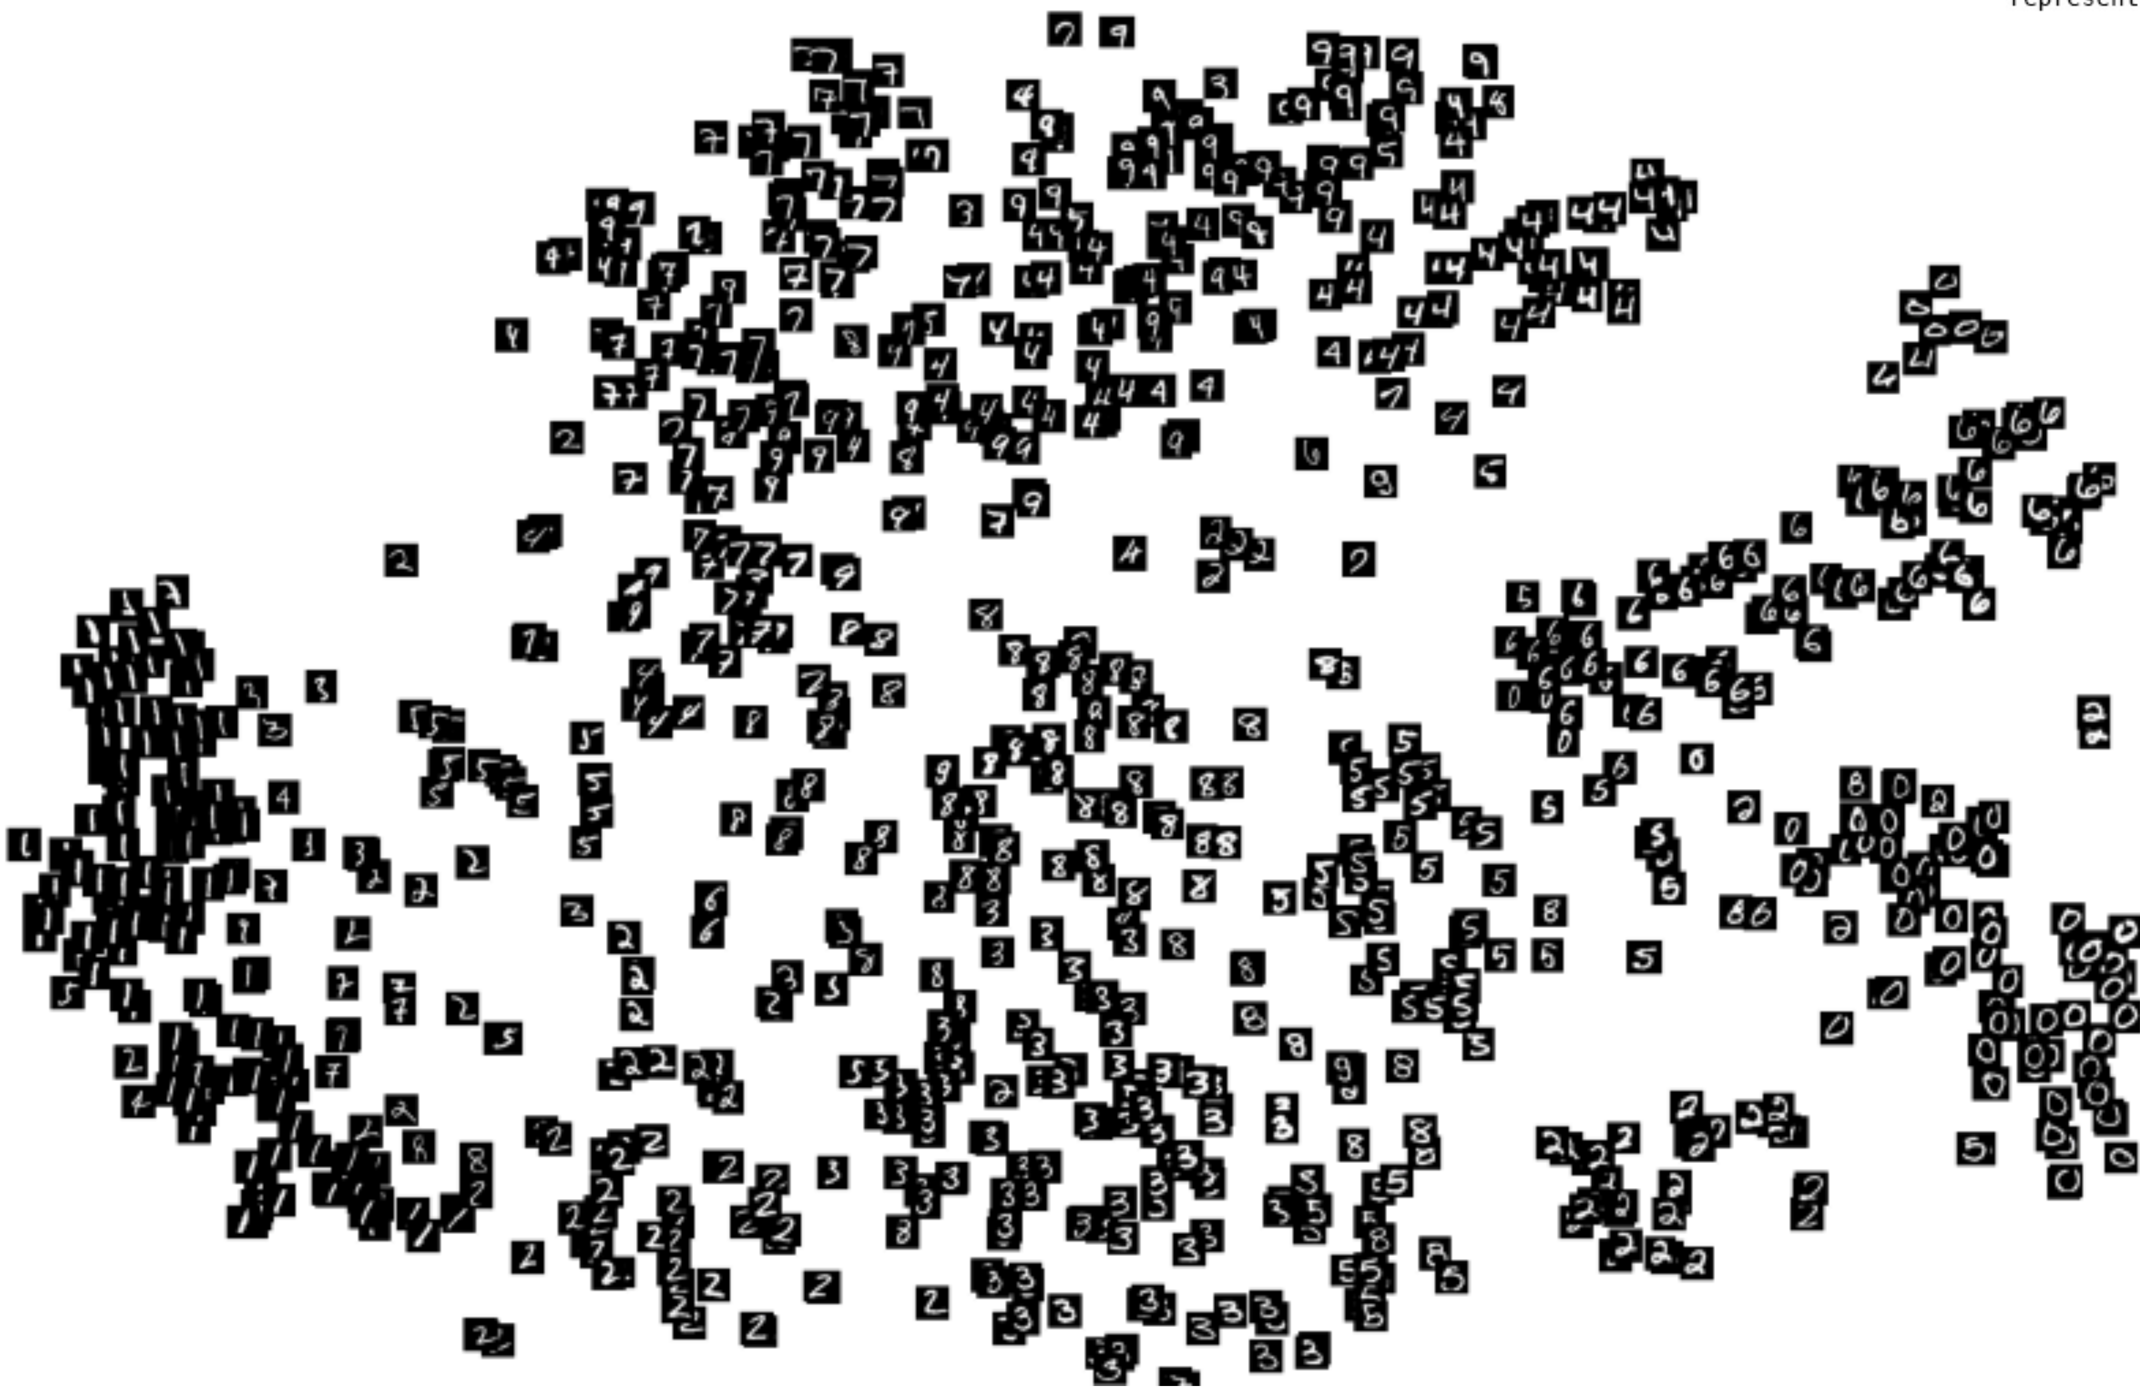

Latent dim 2

Latent dim 1

AUTOENCODERS

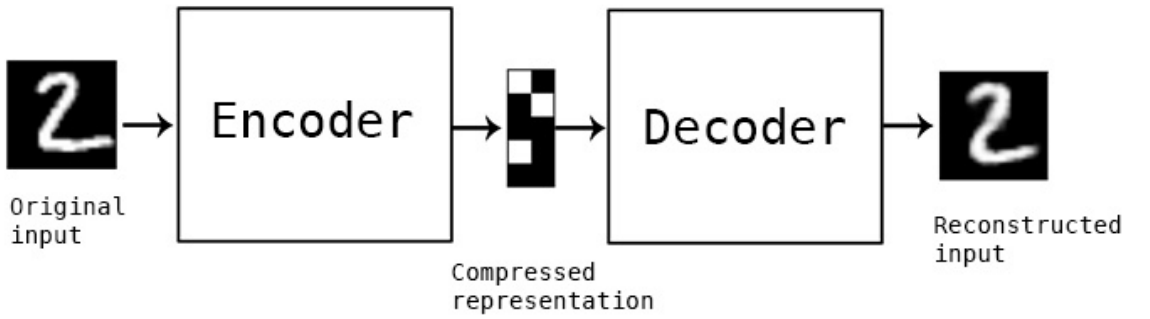

- Input slightly different than training set –> reconstruction loss high, because latent space is ill-defined there
- Not robust
- What is between the data points?

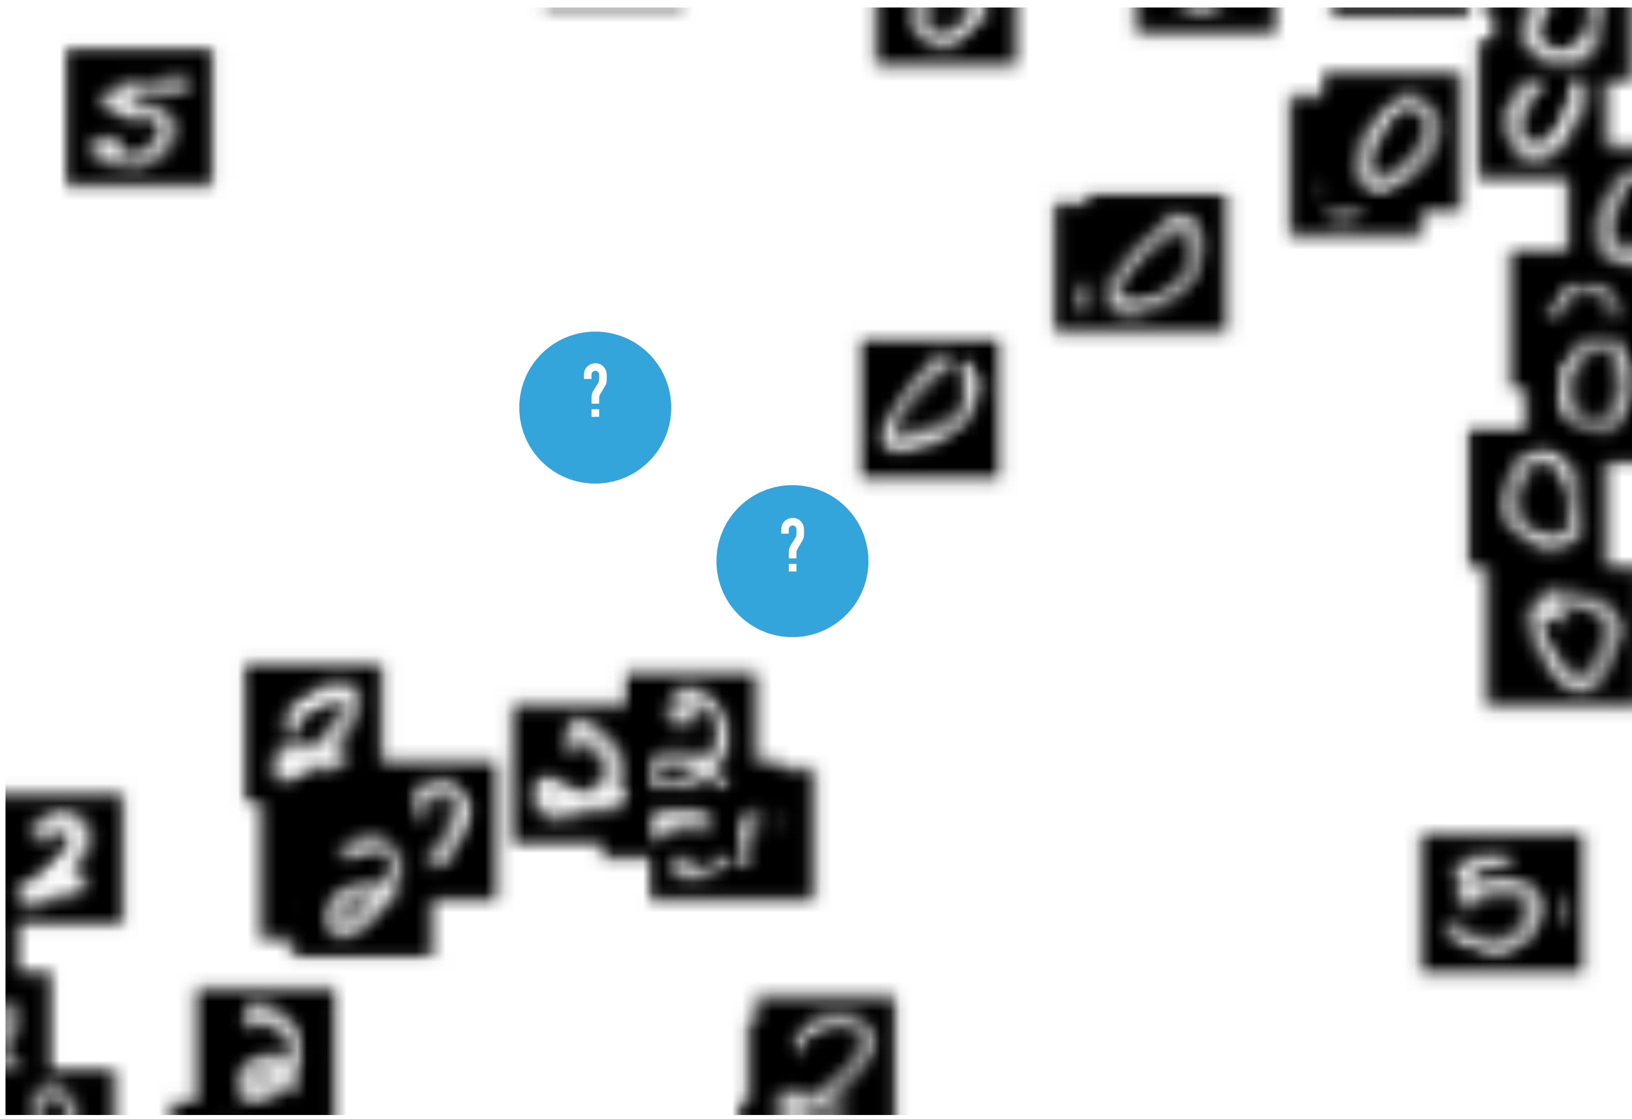

AUTOENCODERS

- If only the points could be grouped together…
- Unsupervised clustering, interpolation between data points …

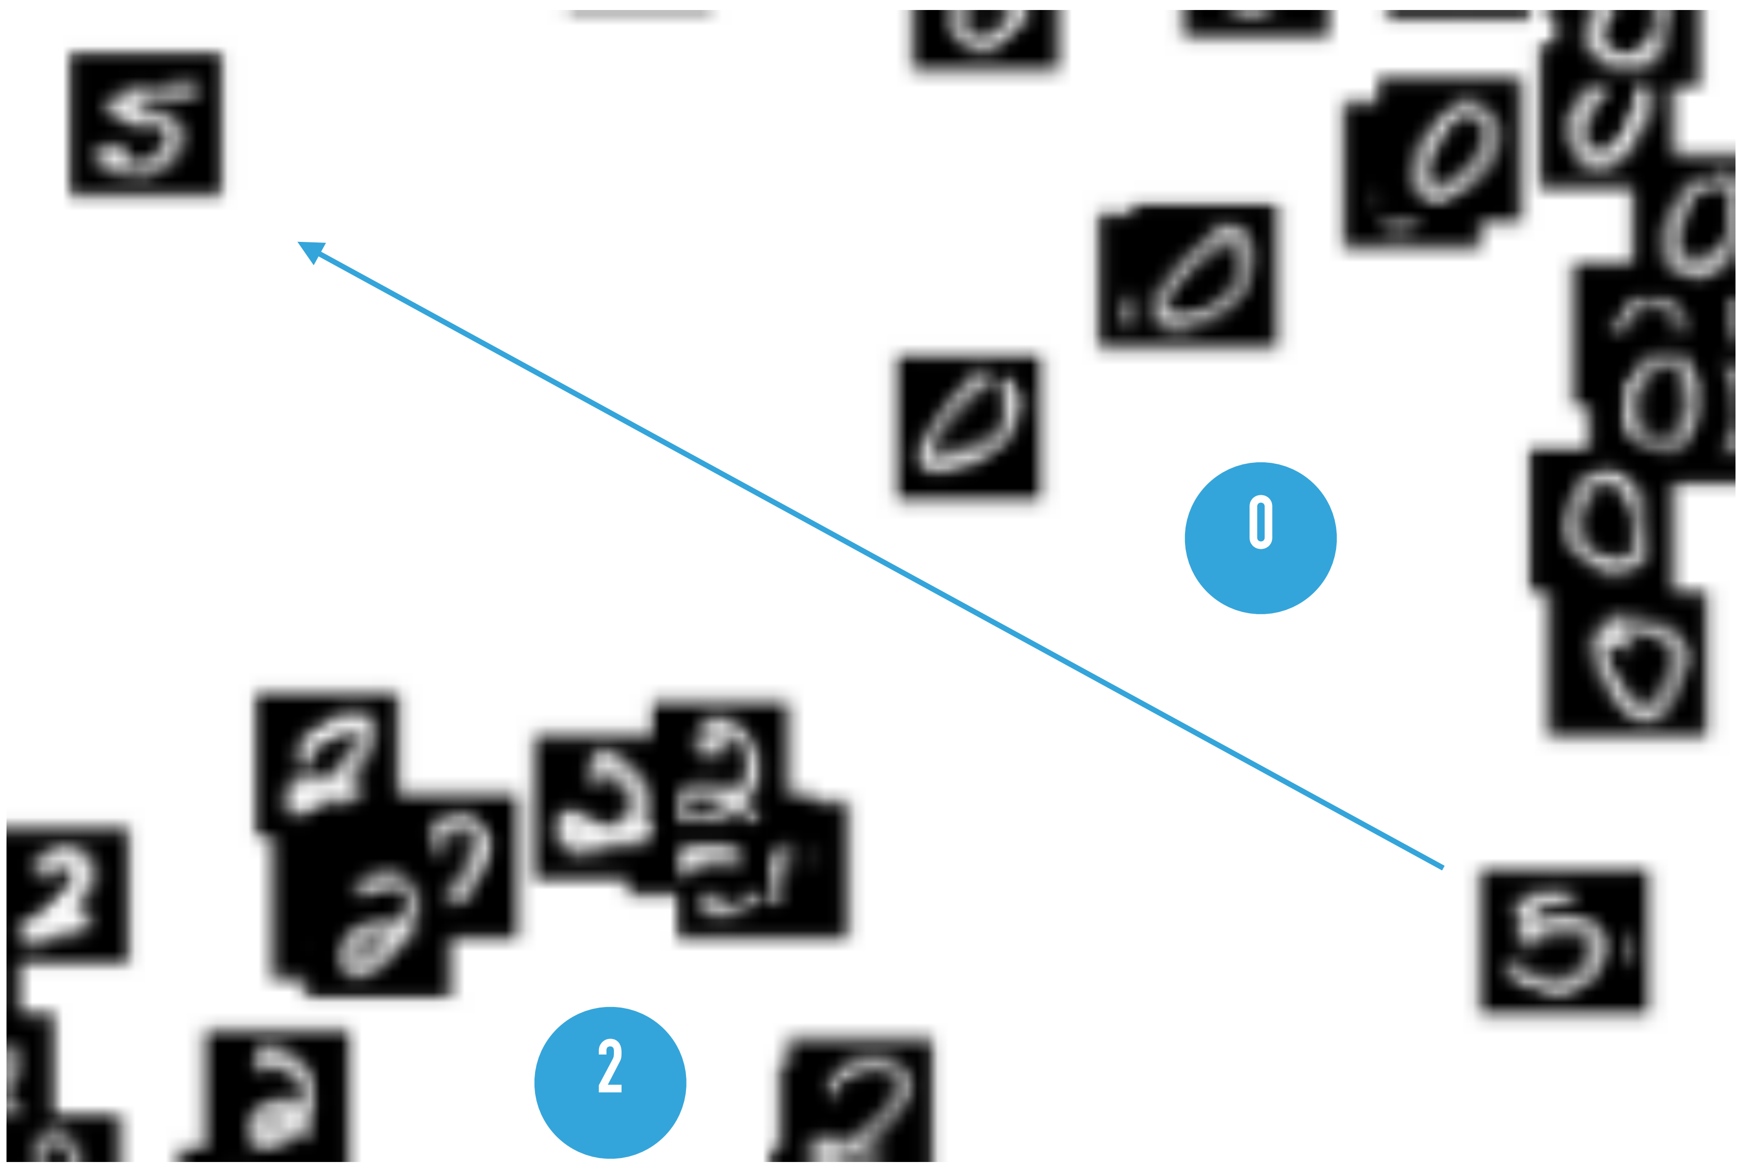

VARIATIONAL AUTOENCODERS

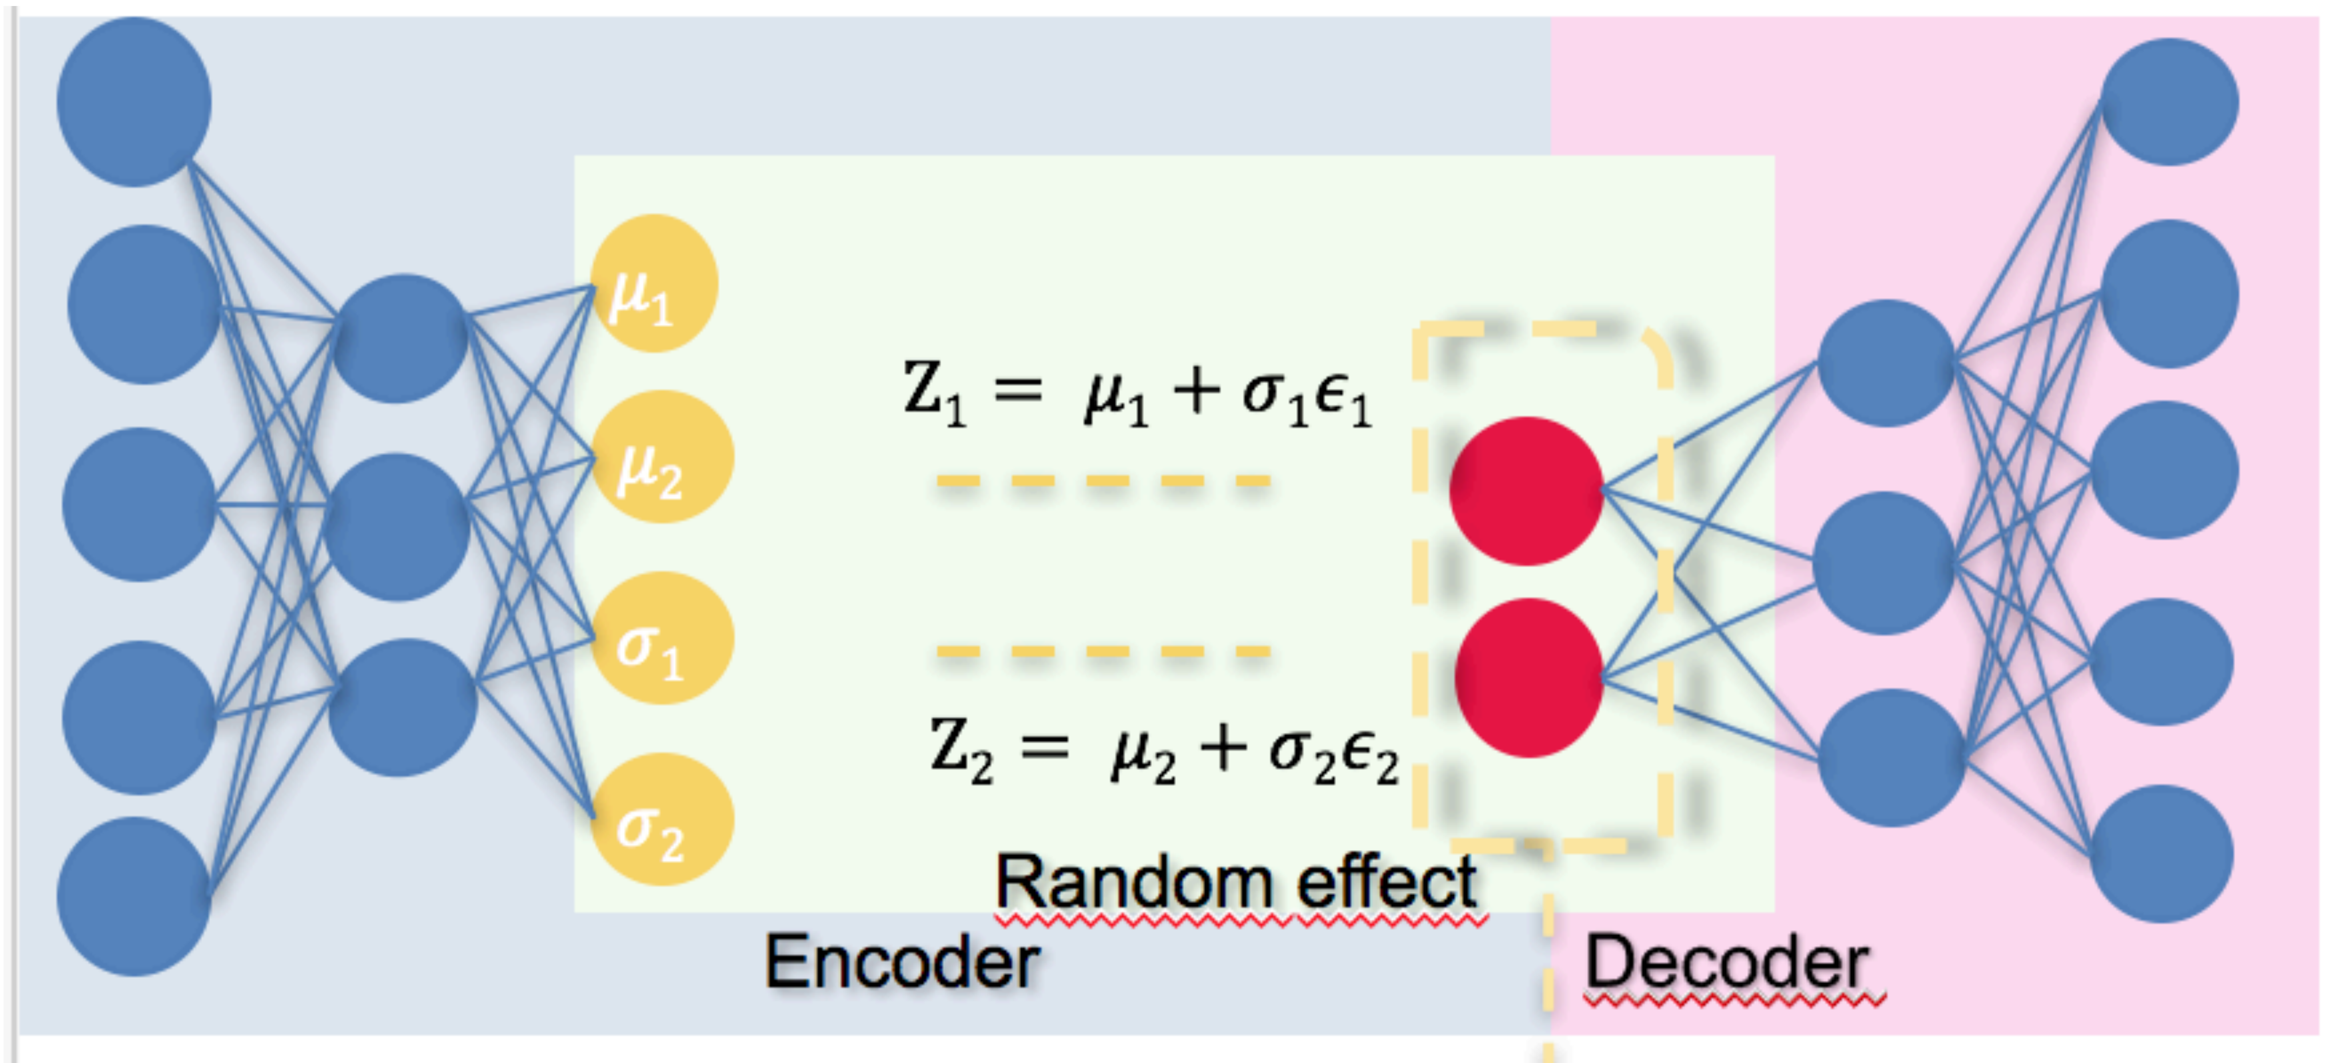

VAE

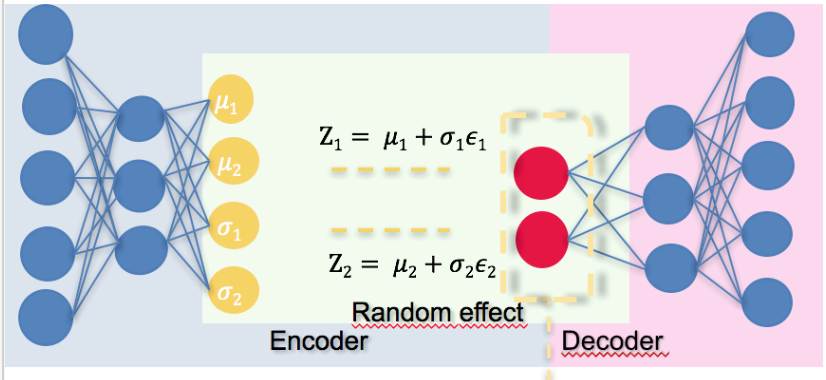

- Force ordering in latent space
- During training, you are minimising some loss function
- For regression (normal AE): MSE(*output - input*)

VAE

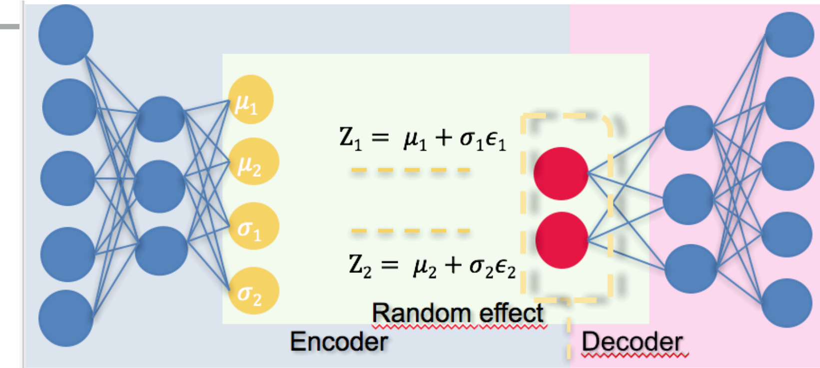

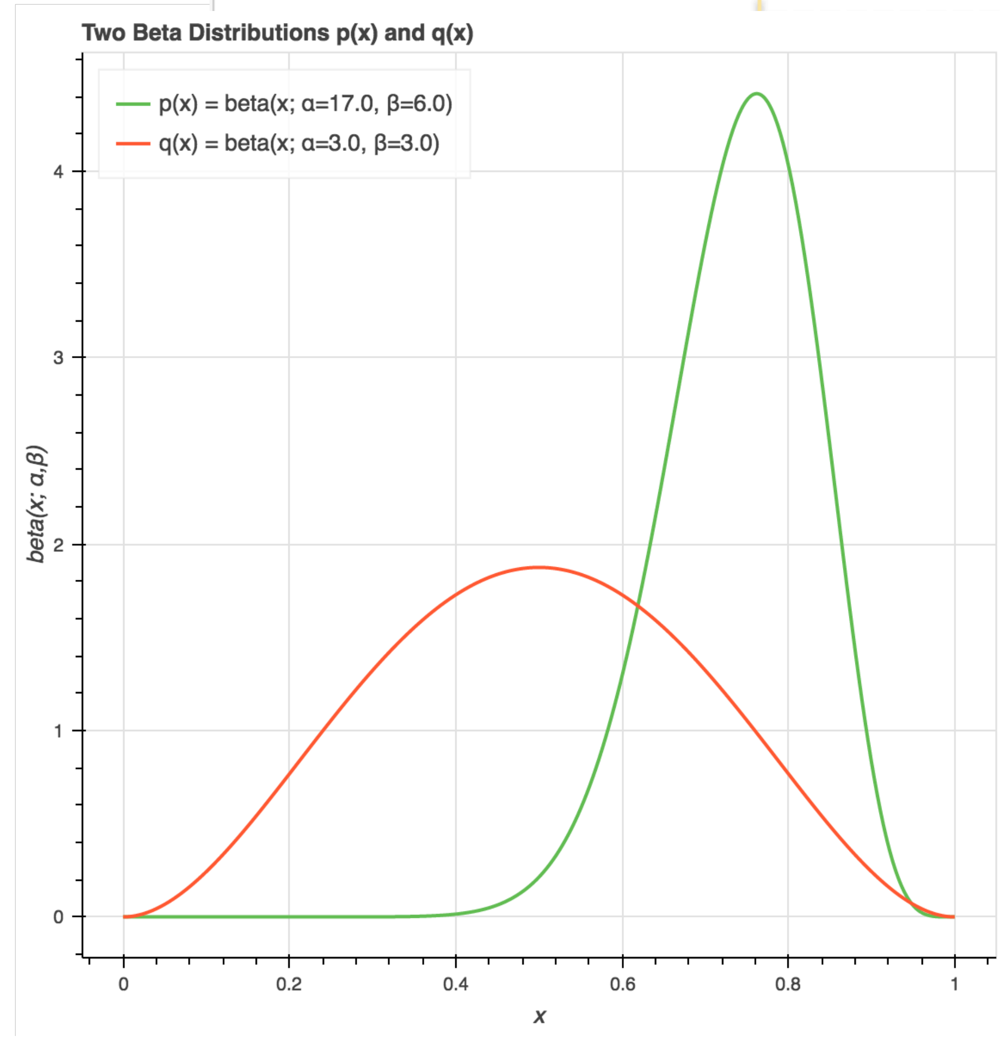

- Force ordering in latent space
- During training, you are minimising some loss function
- For regression (normal AE): MSE(*output - input*)
- Add KL-divergence term: $\Sigma_i \mathrm{KL}(\mathcal{N}(\mu_i, \sigma_i), \mathcal{N}(0,1)) := \mathrm{KL}(\mu,\sigma)$
- So $\mathcal{L} = \mathrm{MSE}(output - input) + \mathrm{KL}(\mu,\sigma)$

- The KL divergence punishes latent space values far away from the center
- Also, every point has a variance that is pushed to 1
- Balance MSE and KL –> group similar structures around the center while keeping RL in check

LATENT SPACE

- Same example, but now a VAE

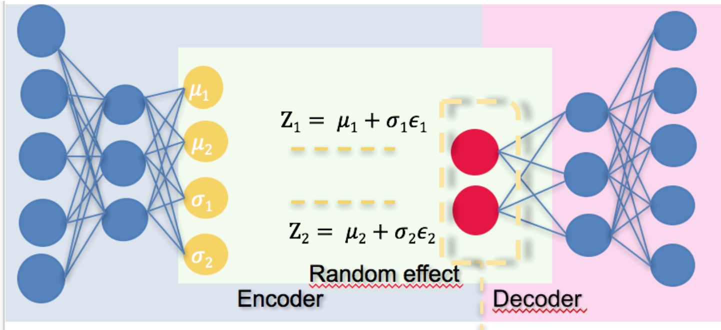

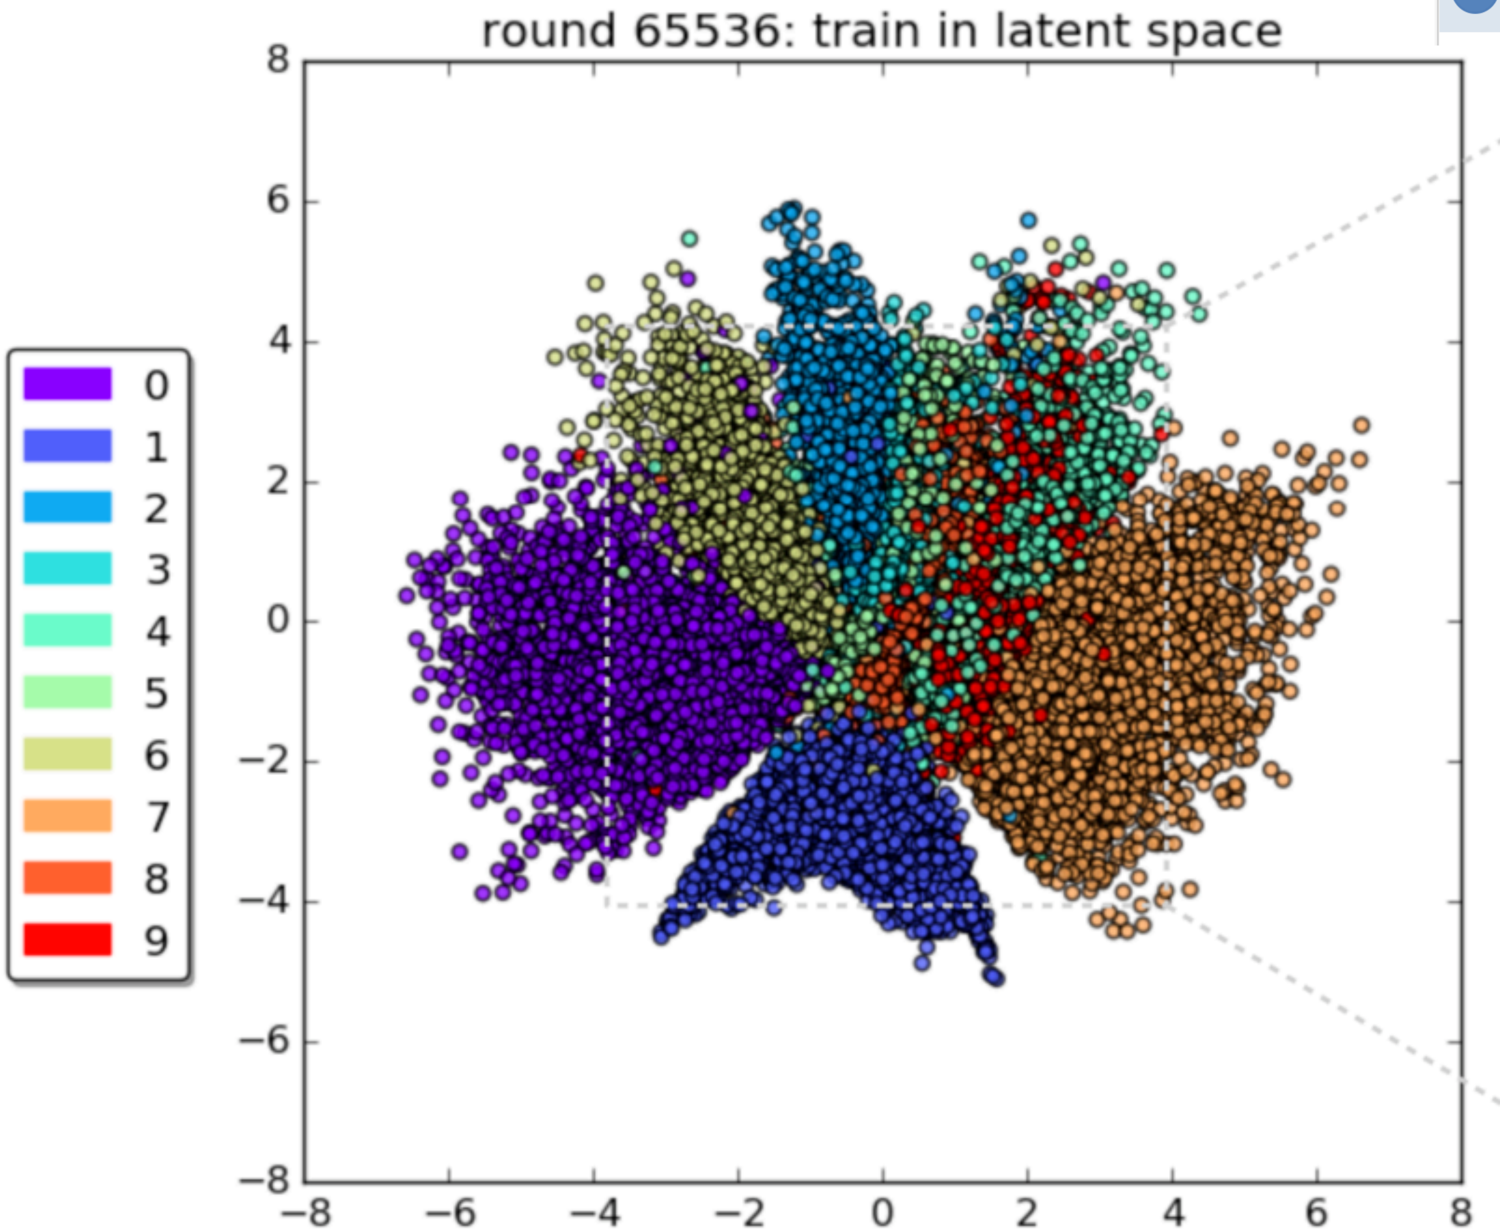

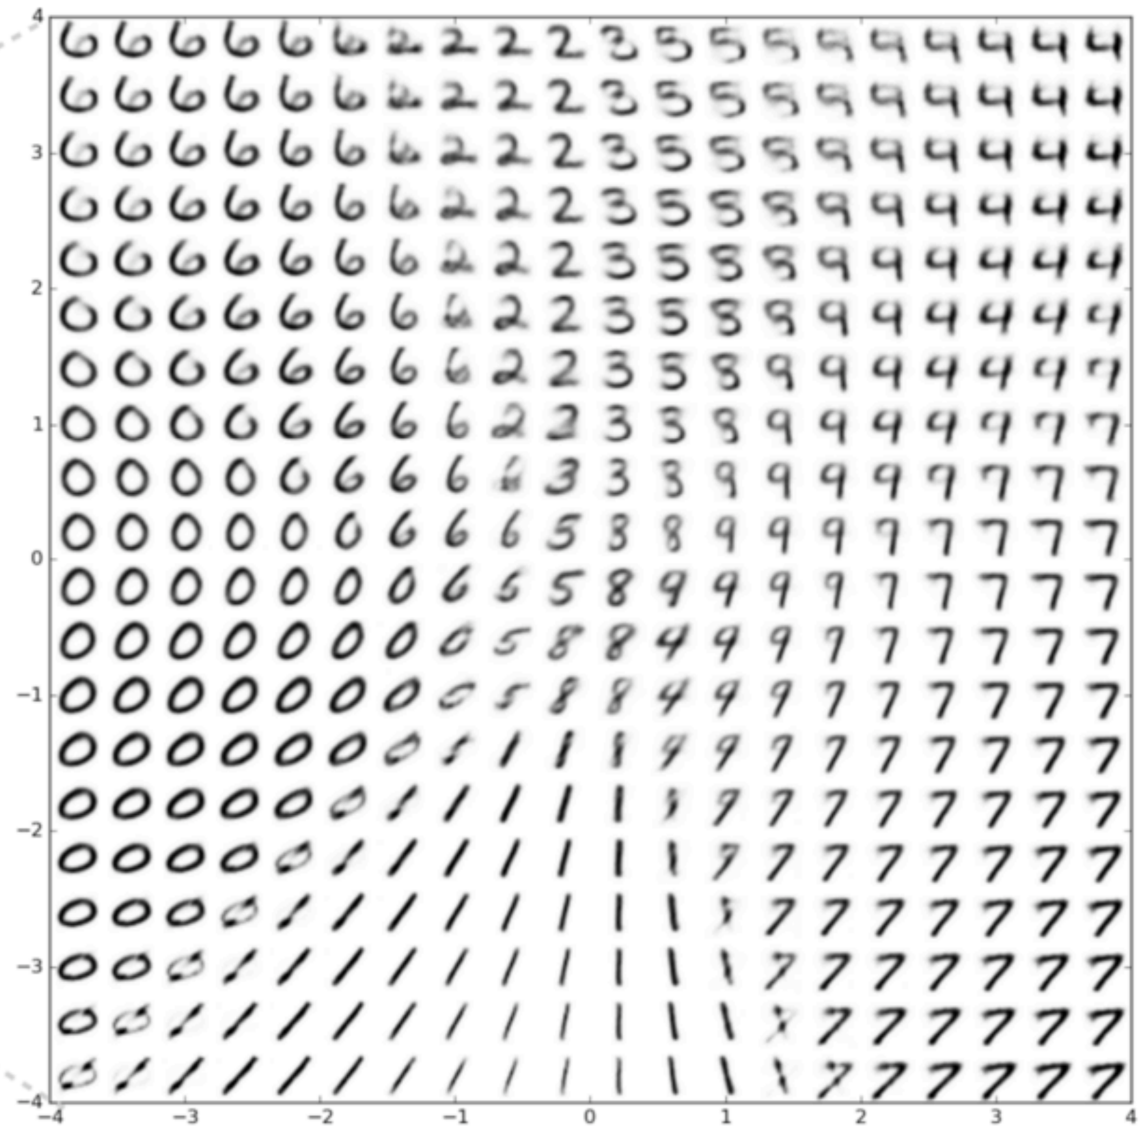

- Balancing MSE and KL is tricky
- Balance using another hyperparameter **β**
  - $\mathscr{L} = (1-\beta) * \mathrm{MSE}(\mathit{output} - \mathit{input}) + \beta * \mathrm{KL}(\mu, \sigma)$
- β-VAE

β	Avg var	Avg mean
1	1	1.89E-09
5E-01	0.99999905	2.35E-07
5E-02	0.86448085	…
5E-03	0.554529	
5E-04	0.3784553	
5E-05	0.09676677	
5E-06	0.008932933	
0	0.0000442	

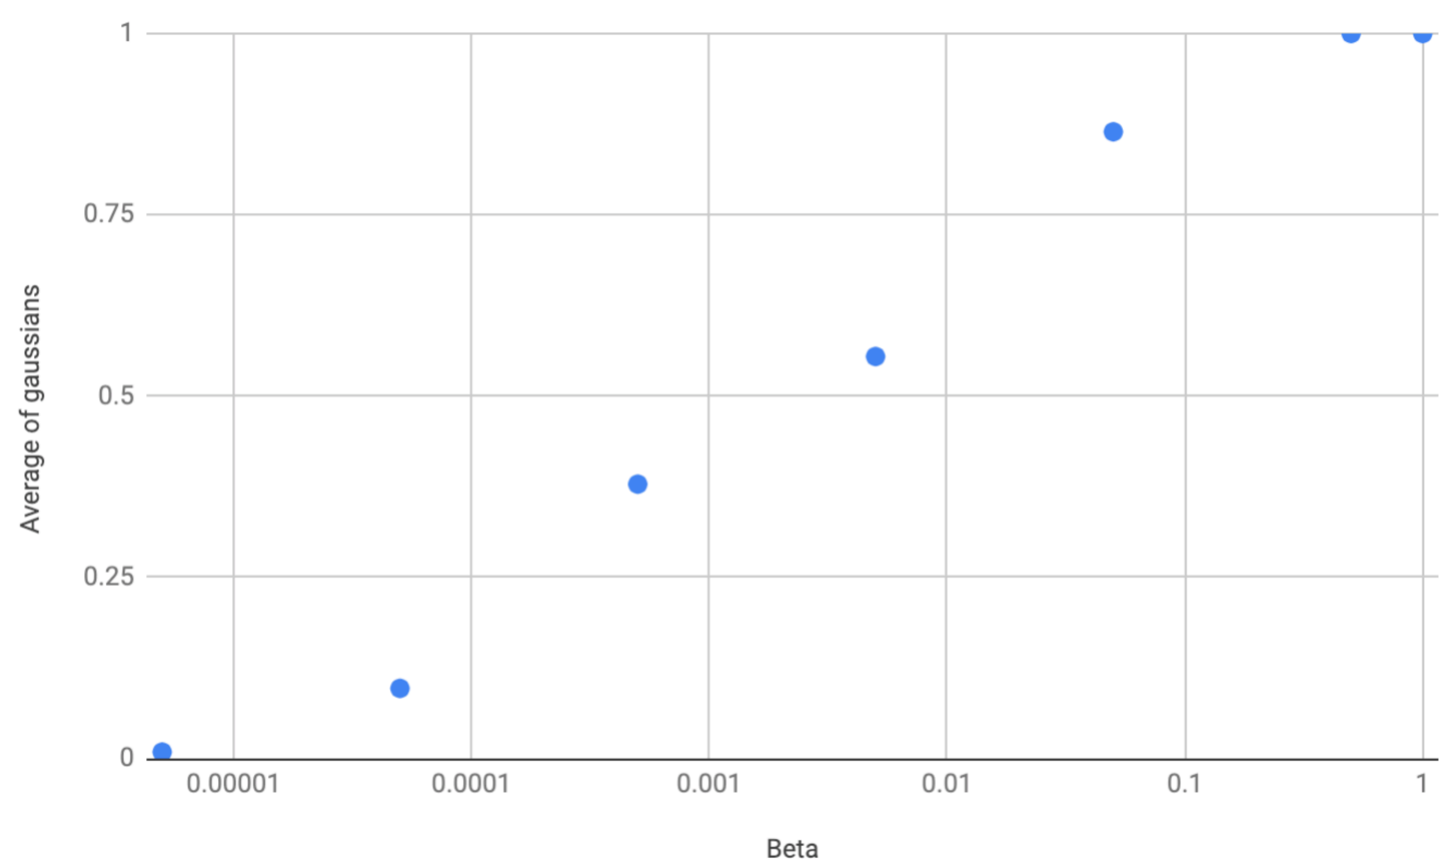

VAE

- Use the latent space and decoder as generative model\
- Explore the latent space!

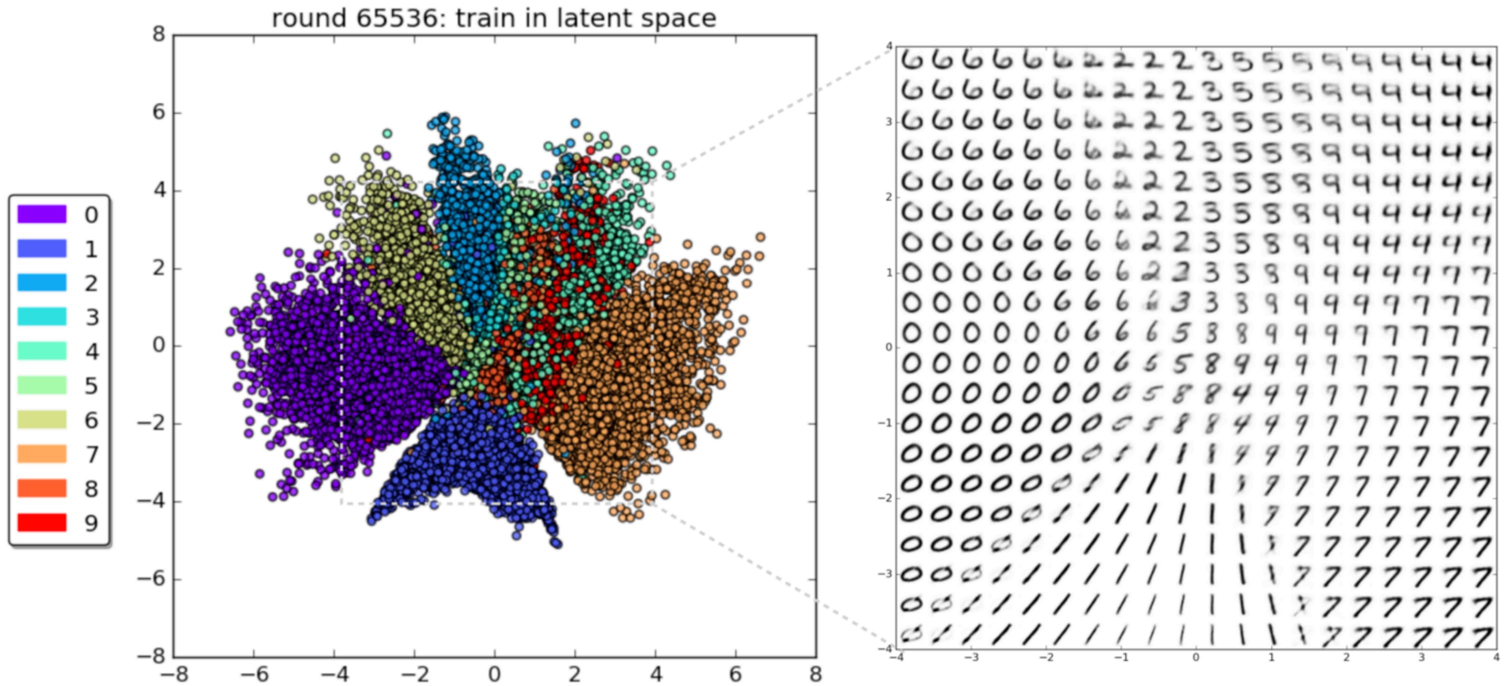

PCA on the latent variables

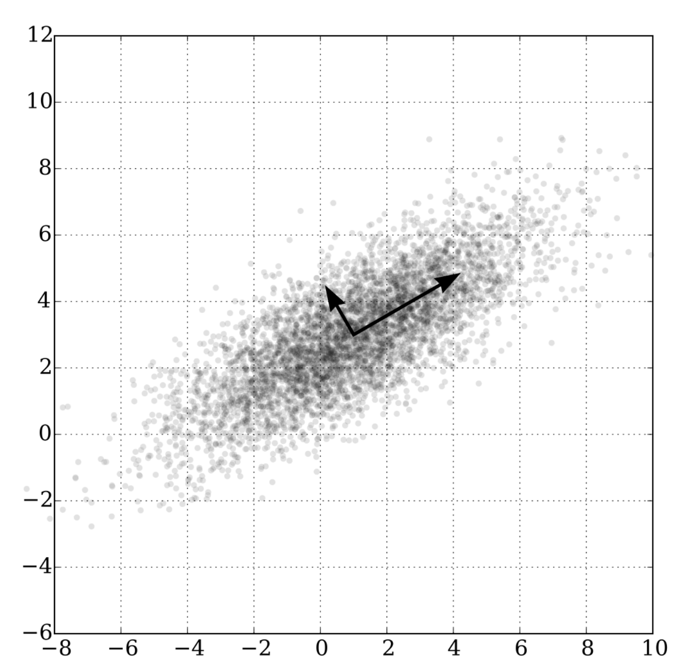

PLAYING WITH LATENT SPACES

- Train VAE on face images
- Change the latent space variables

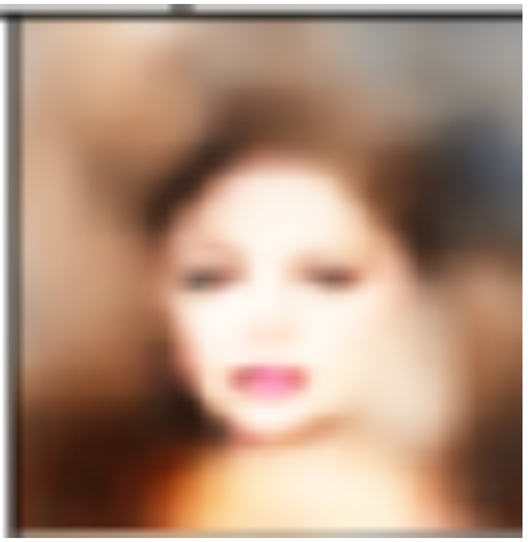

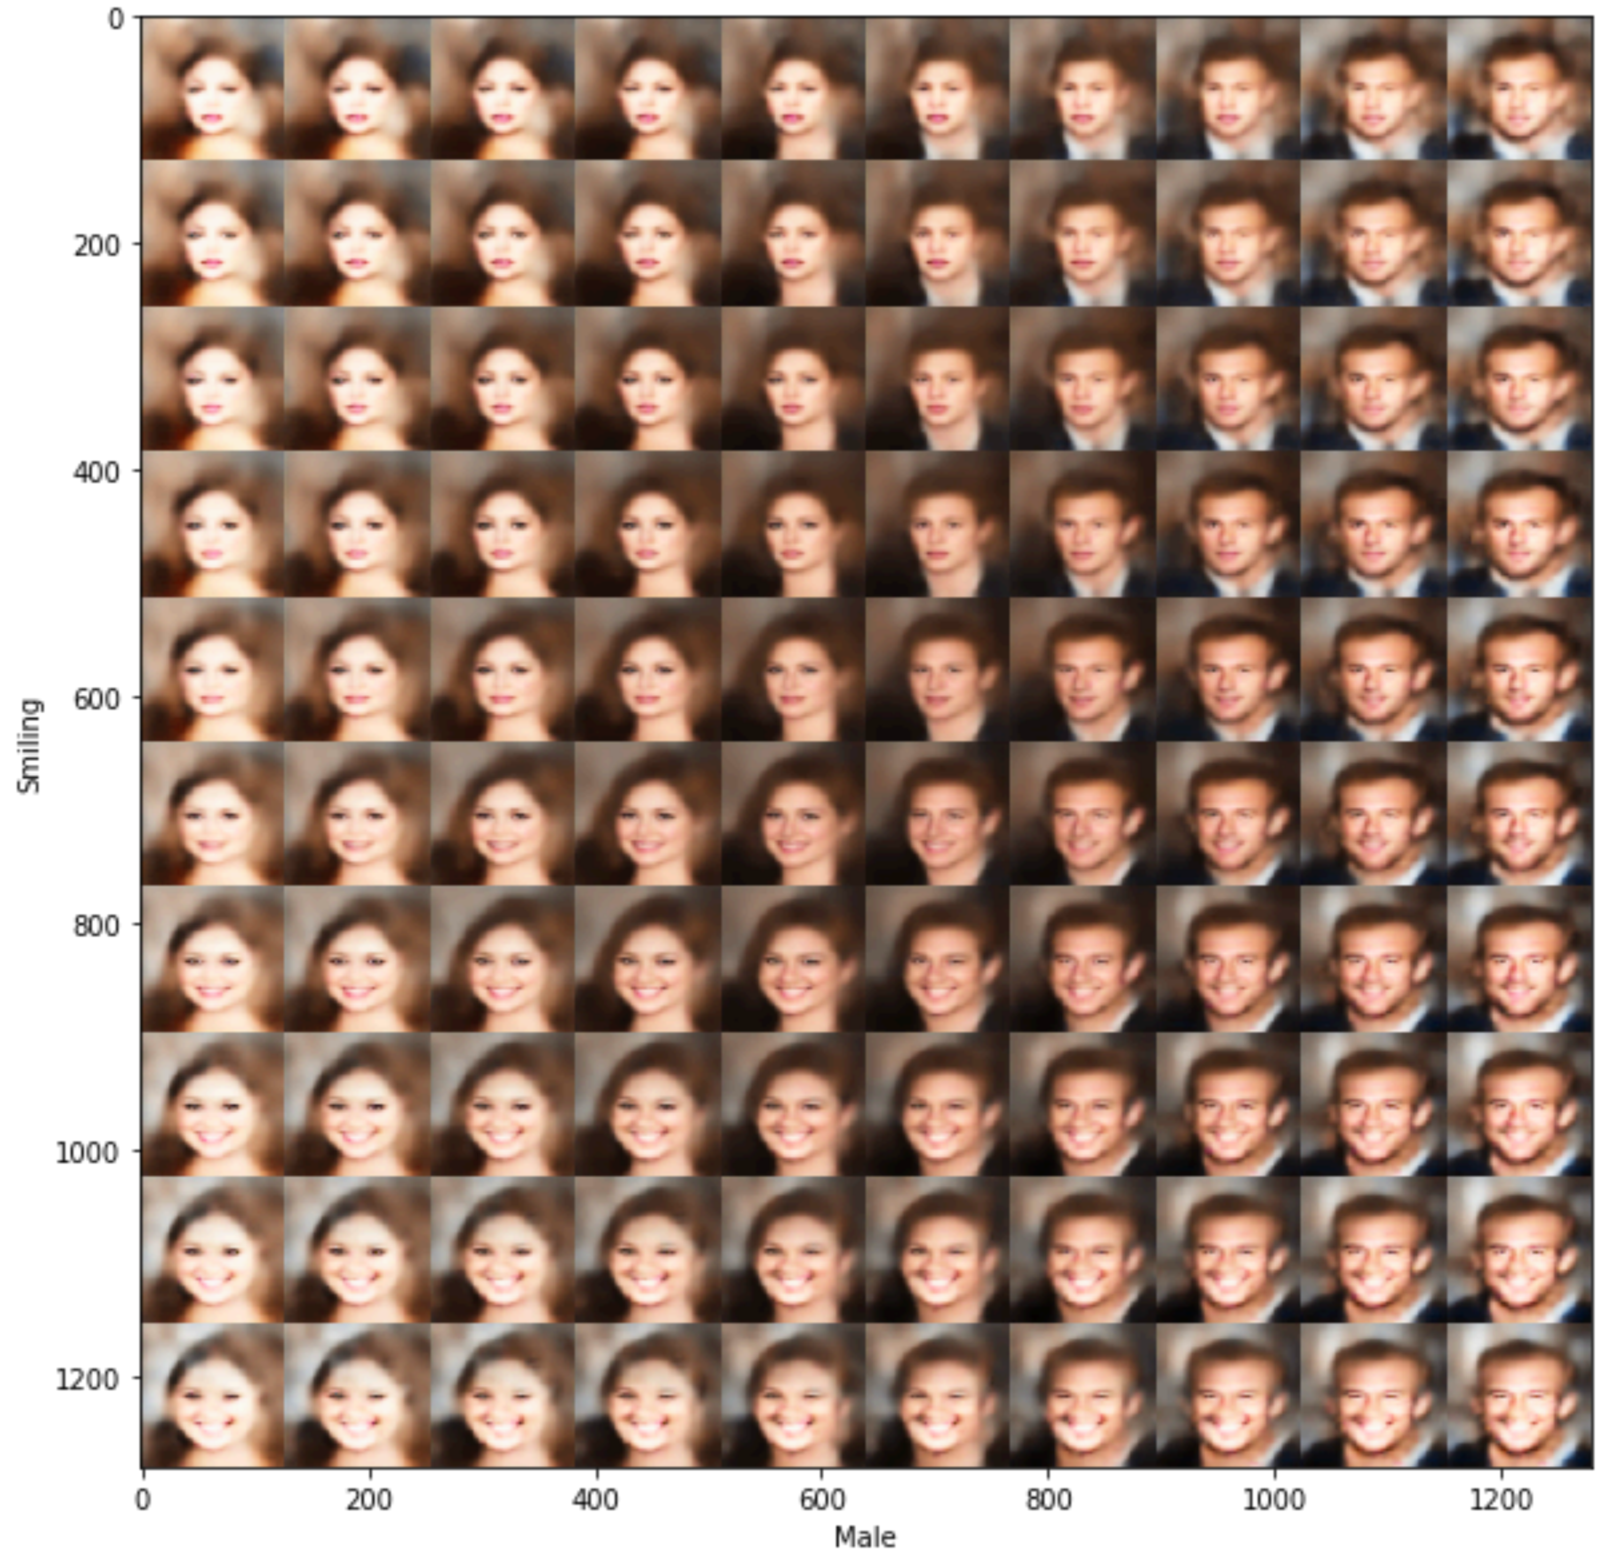

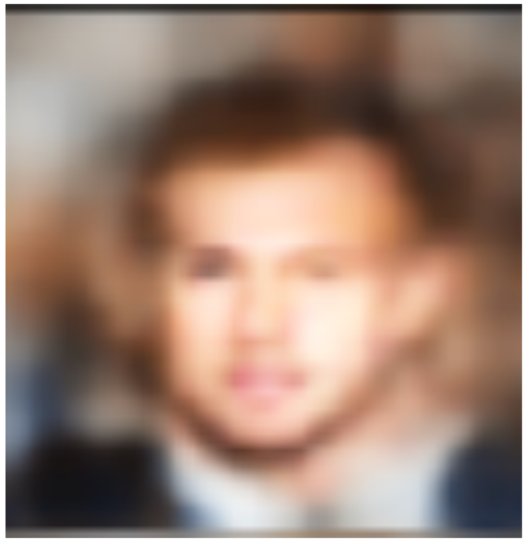

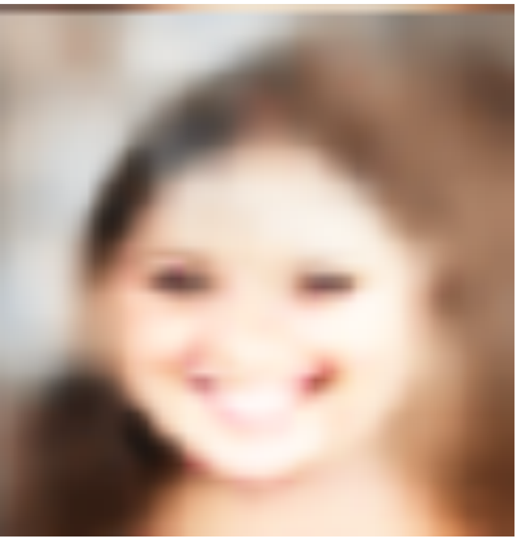

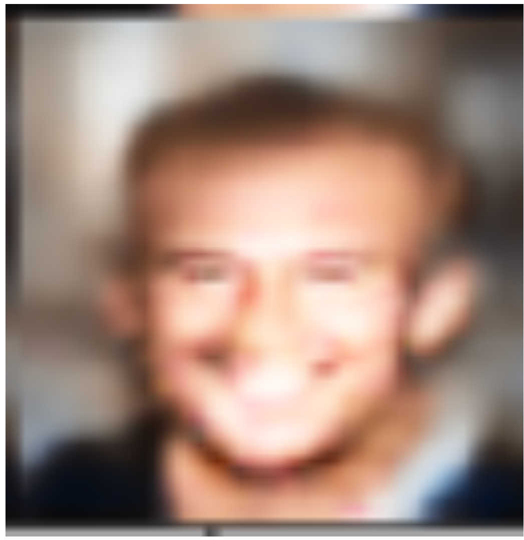

PLAYING WITH LATENT SPACES

- Or 3D objects

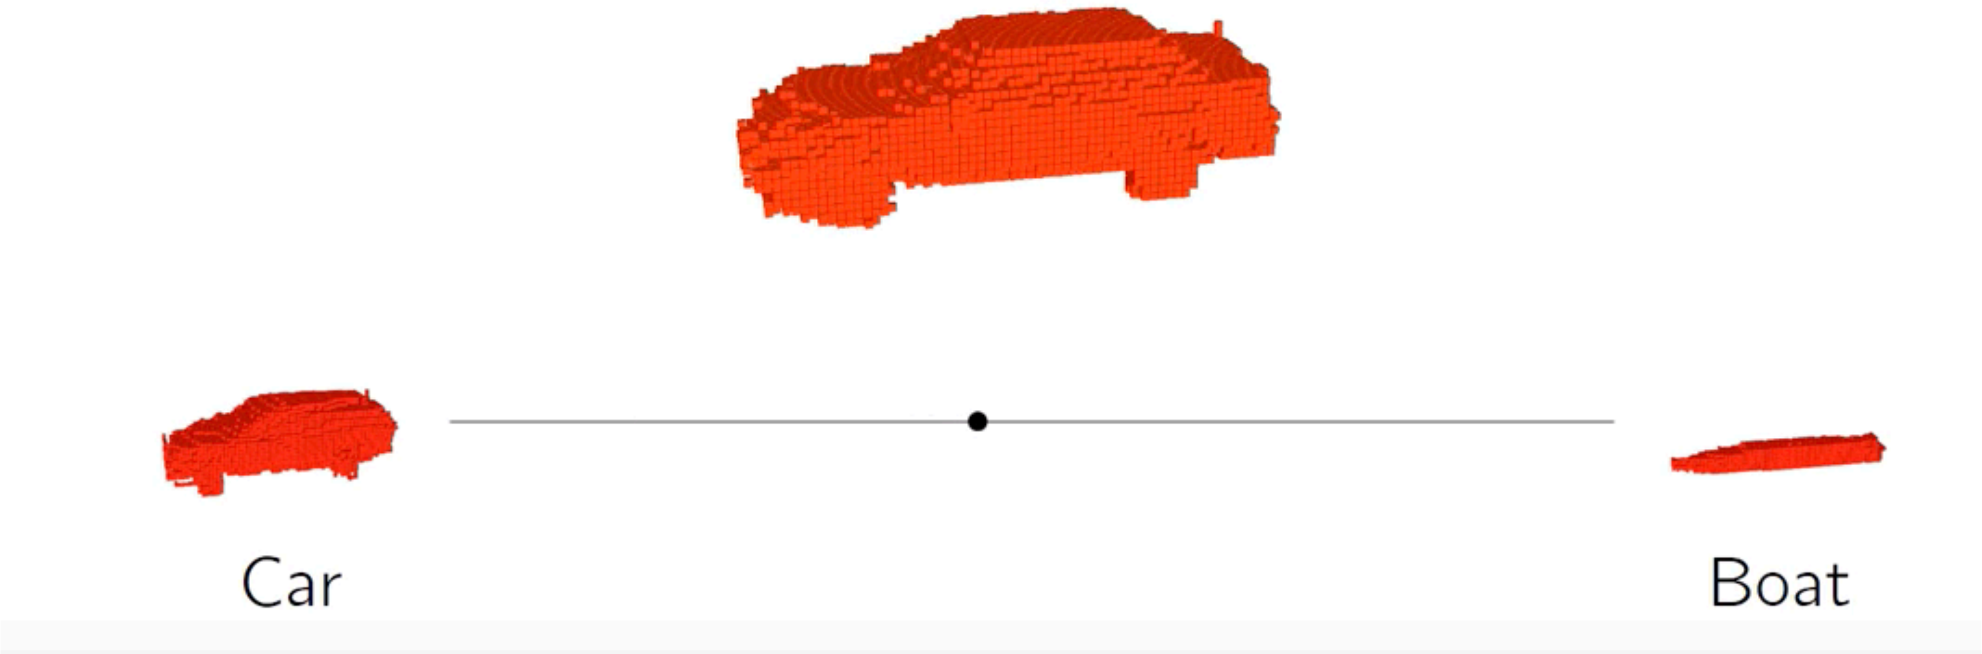

PLAYING WITH LATENT SPACES

- Or 3D objects

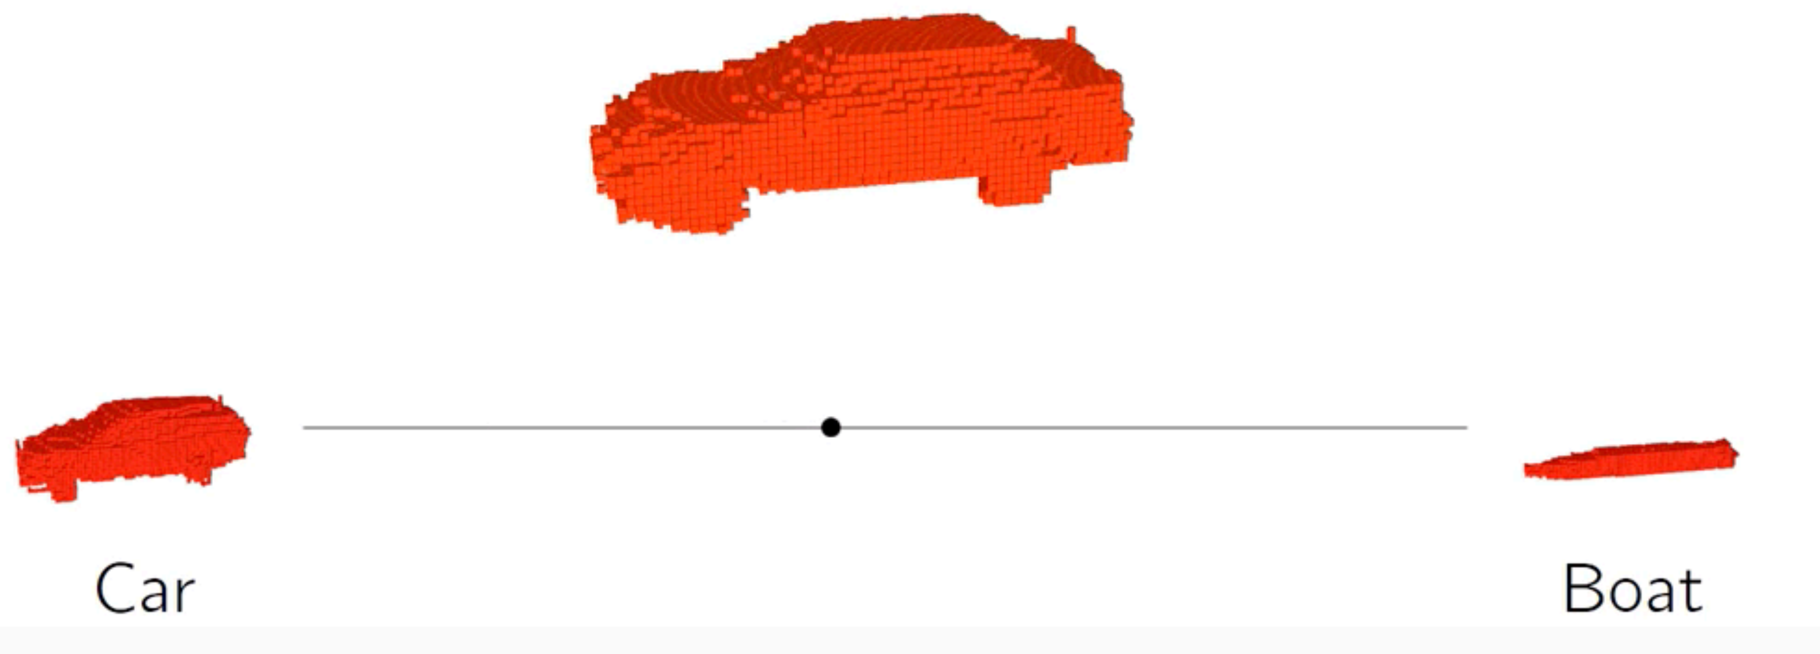

- Latent space = abstract representation of your data

- Encoder maps input to gaussians in latent space
  = Gaussian mixture –> you can do lots of stuff

Event Generation and Statistical Sampling with Deep Generative Models

Sydney Otten, Sascha Caron, Wieske de Swart, Melissa van Beekveld, Luc Hendriks, Caspar van Leeuwen, Damian Podareanu, Roberto Ruiz de Austri, Rob Verheyen

(Submitted on 3 Jan 2019)

We present a study for the generation of events from a physical process with generative deep learning. To simulate physical processes it is not only important to produce physical events, but also to produce the events with the right frequency of occurrence (density). We investigate the feasibility to learn the event generation and the frequency of occurrence with Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) to produce events like Monte Carlo generators. We study three toy models from high energy physics, i.e. a simple two-body decay, the processes $e^+e^- \to Z \to l^+l^-$ and $pp \to t\bar{t}$ including the decay of the top quarks and a simulation of the detector response. We show that GANs and the standard VAE do not produce the right distributions. By buffering density information of Monte Carlo events in latent space given the encoder of a VAE we are able to construct a prior for the sampling of new events from the decoder that yields distributions that are in very good agreement with real Monte Carlo events and are generated $\mathcal{O}(10^8)$ times faster. Applications of this work include generic density estimation and sampling, targeted event generation via a principal component analysis of encoded events in the latent space and the possibility to generate better random numbers for importance sampling, e.g. for the phase space integration of matrix elements in quantum perturbation theories. The method also allows to build event generators directly from real data events.

CONCLUSION

- VAEs can be used for
    - Outlier / anomaly detection
    - Noise reduction
    - Generative modelling
    - Data compression

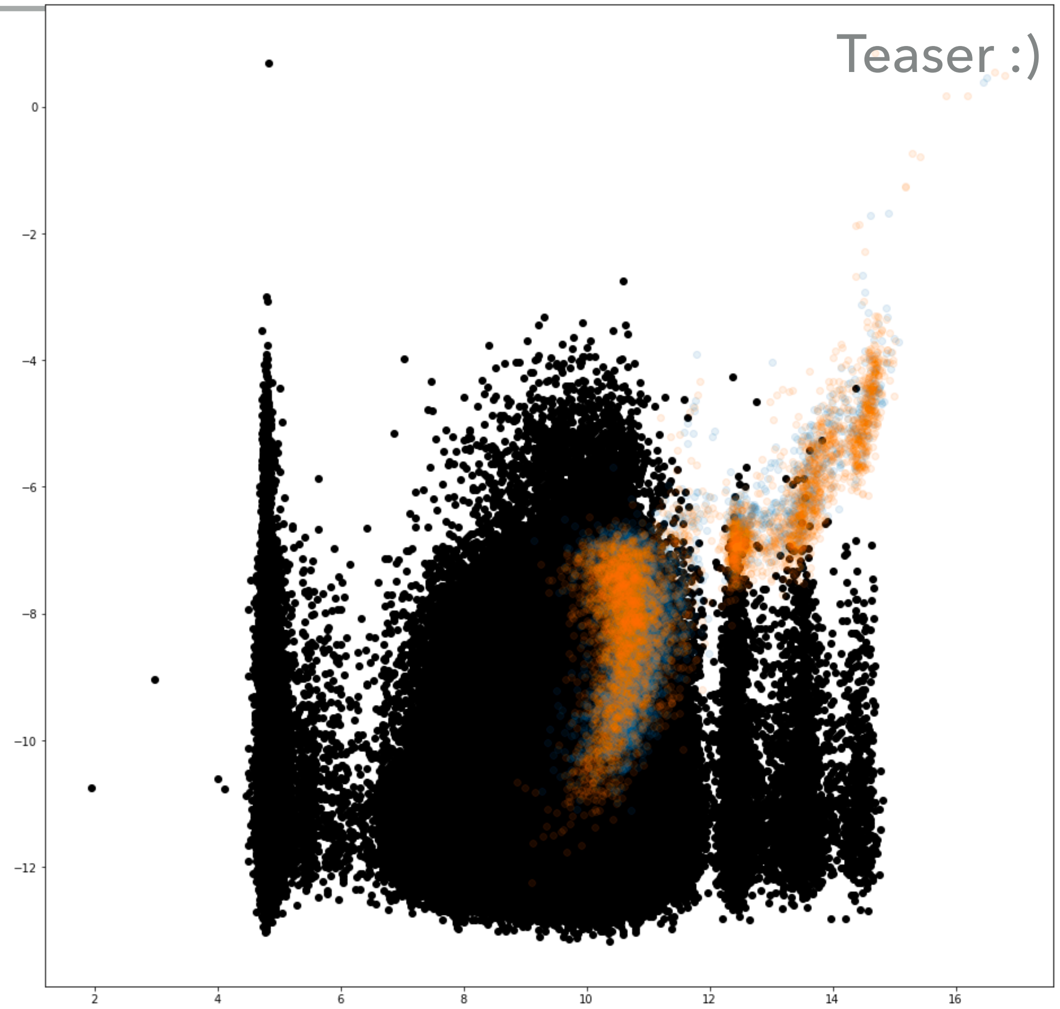

- Exploration of latent space can give very interesting applications – event generation, hybrid models, density estimation, …