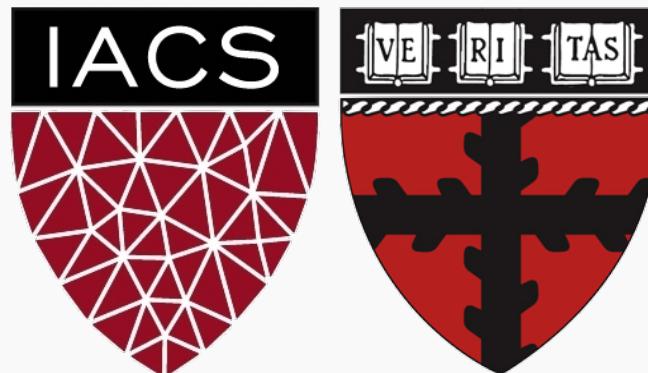


# Lecture 19: Variational Autoencoders

CS109B Data Science 2  
Pavlos Protopapas and Mark Glickman



# Outline

---

Motivation for Variational Autoencoders (VAE)

Mechanics of VAE

Separability of VAE

The math behind everything

Generative models

# Outline

---

Motivation for Variational Autoencoders (VAE)

Mechanics of VAE

Separability of VAE

The math behind everything

Generative models

# Outline

---

**Motivation for Variational Autoencoders (VAE)**

Mechanics of VAE

Separability of VAE

The math behind everything

Generative models



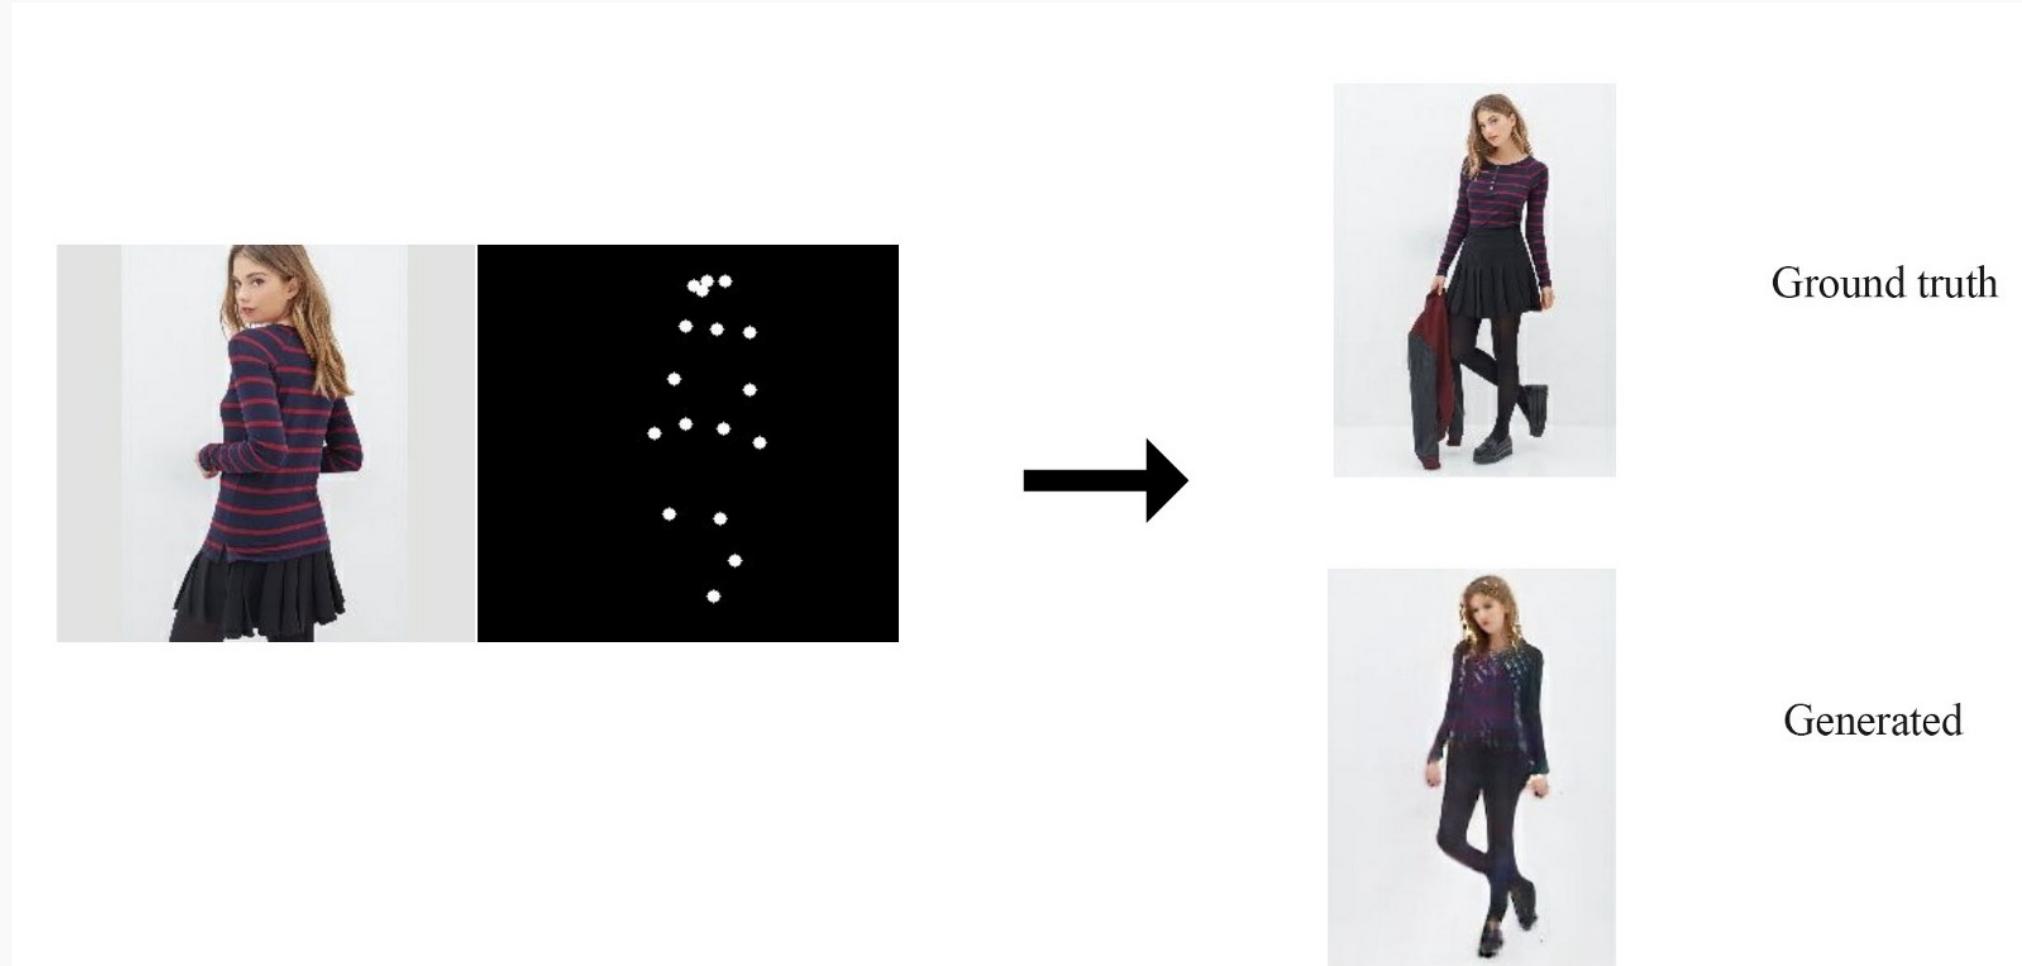
# Generating Data (is exciting)



Figure 7: Generated samples

<https://arxiv.org/pdf/1708.05509.pdf>

# Generating Data (is exciting)



# Generating Data (is exciting)

Zebras ↘ Horses



zebra → horse



horse → zebra



Figure 5:  $1024 \times 1024$  images generated using the CELEBA-HQ dataset. See Appendix F for a larger set of results, and the accompanying video for latent space interpolations.



(a) MidiNet model 1



(b) MidiNet model 2



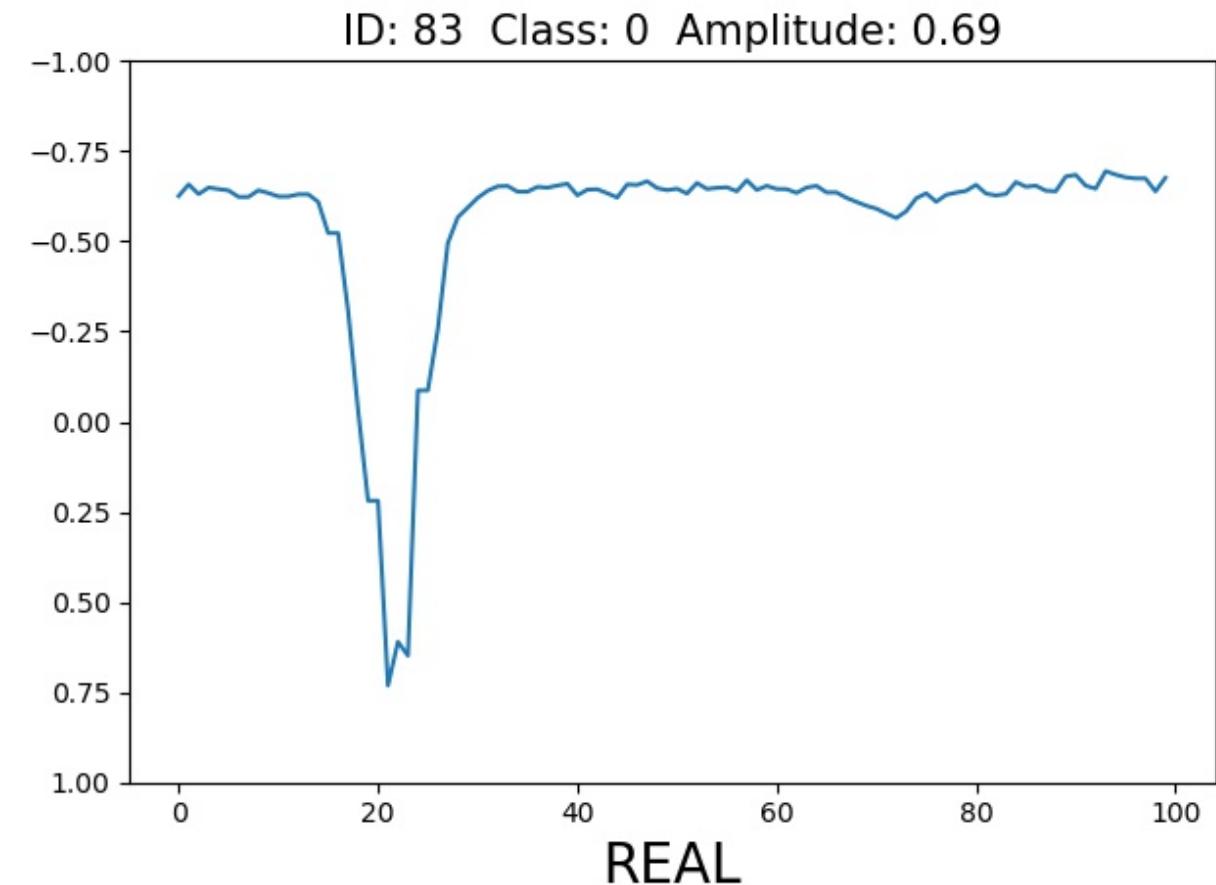
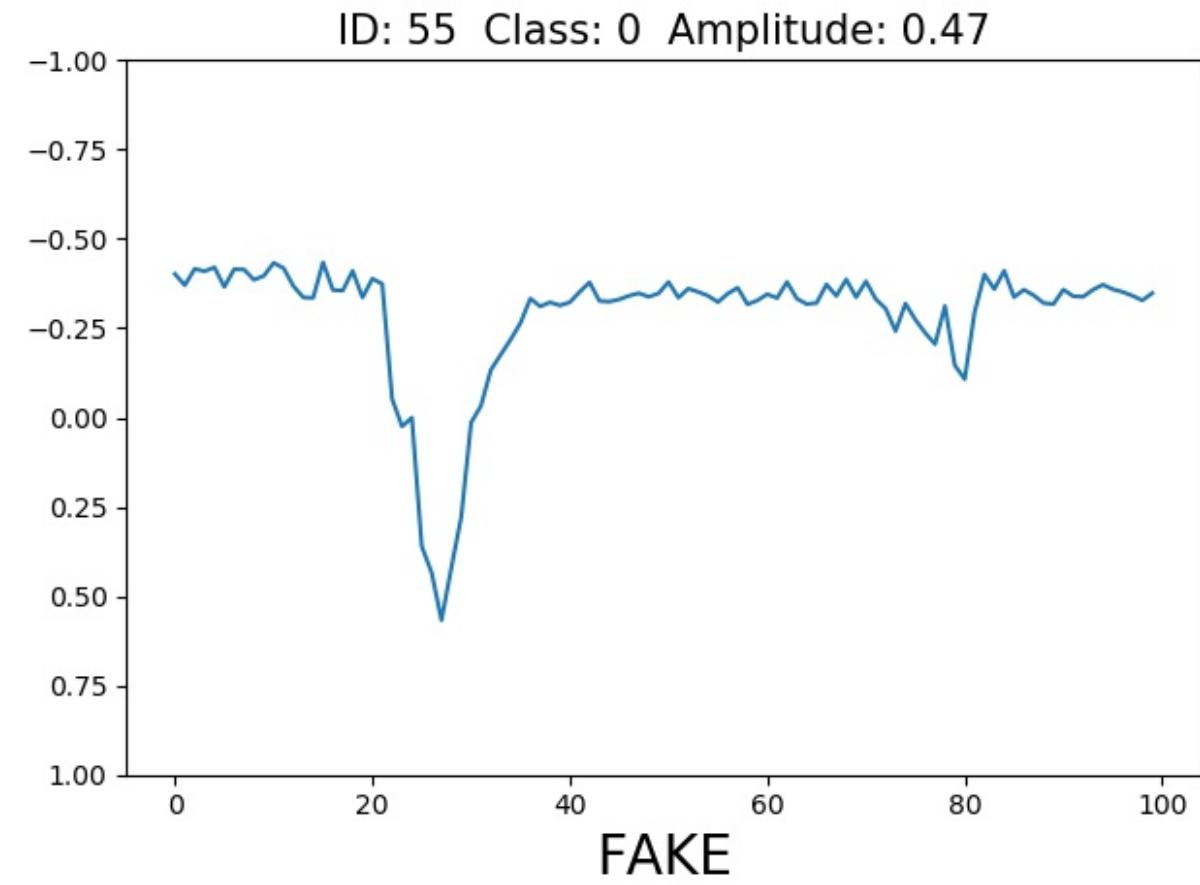
(c) MidiNet model 3

**Figure 3.** Example result of the melodies (of 8 bars) generated by different implementations of MidiNet.

Another use of generating new data is to give us ideas and options. Suppose we're planning a house. We can give the computer the space we have available, and its location. From this, the computer can give us some ideas.

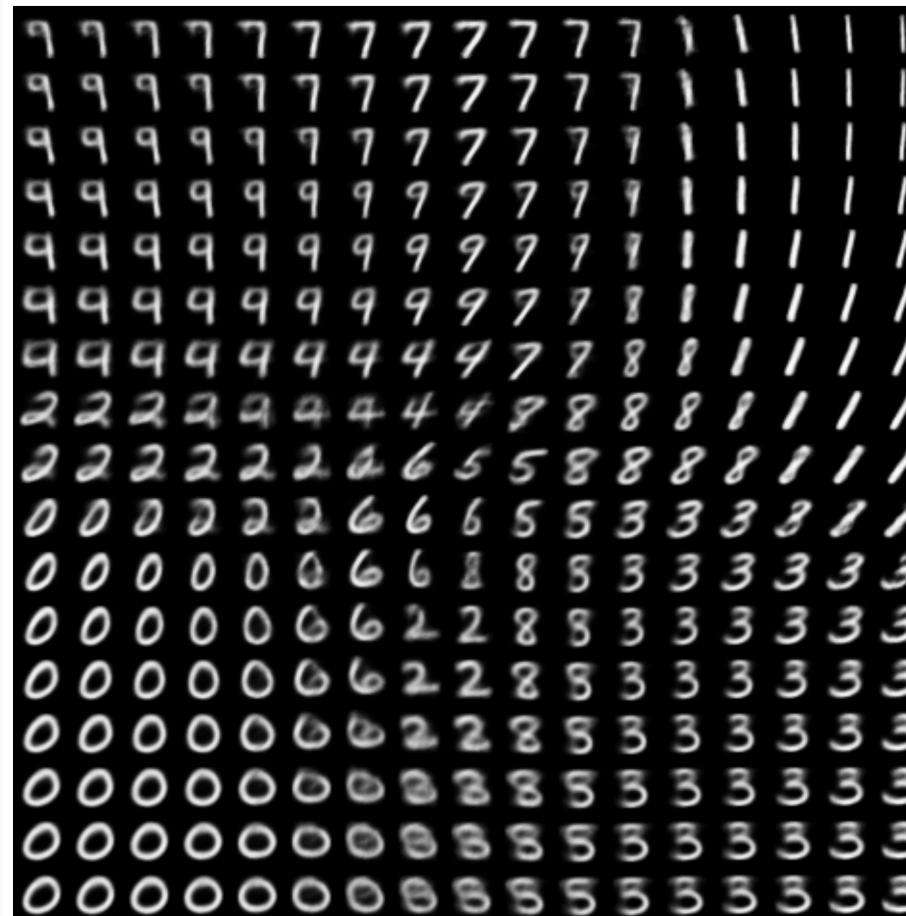


Big networks require big data, and getting high-quality, labeled data is difficult. If we're generating that data ourselves, we can make as much of it as we like.



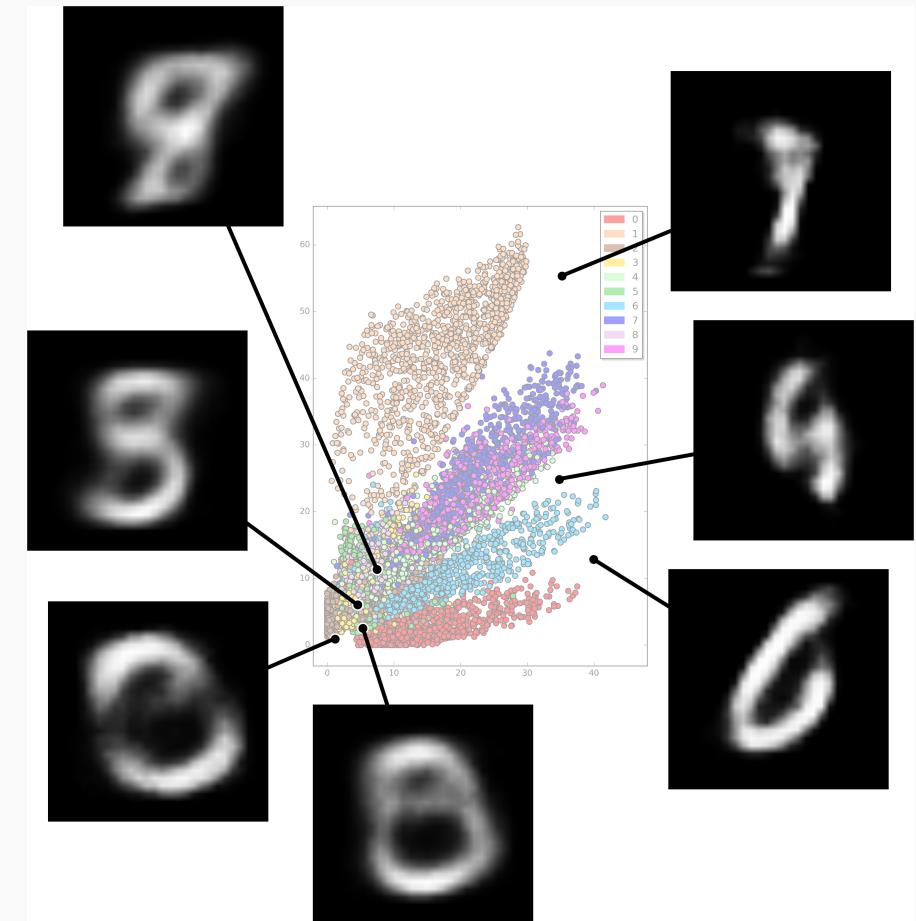
# Generating Data

We saw how to generate new data with a AE in Lecture 18.



# Problems with Autoencoders

- Gaps in the latent space
- Discrete latent space
- Separability in the latent space



# Outline

---

Motivation for Variational Autoencoders (VAE)

**Mechanics of VAE**

Seperability of VAE

The math behind everything

Generative models

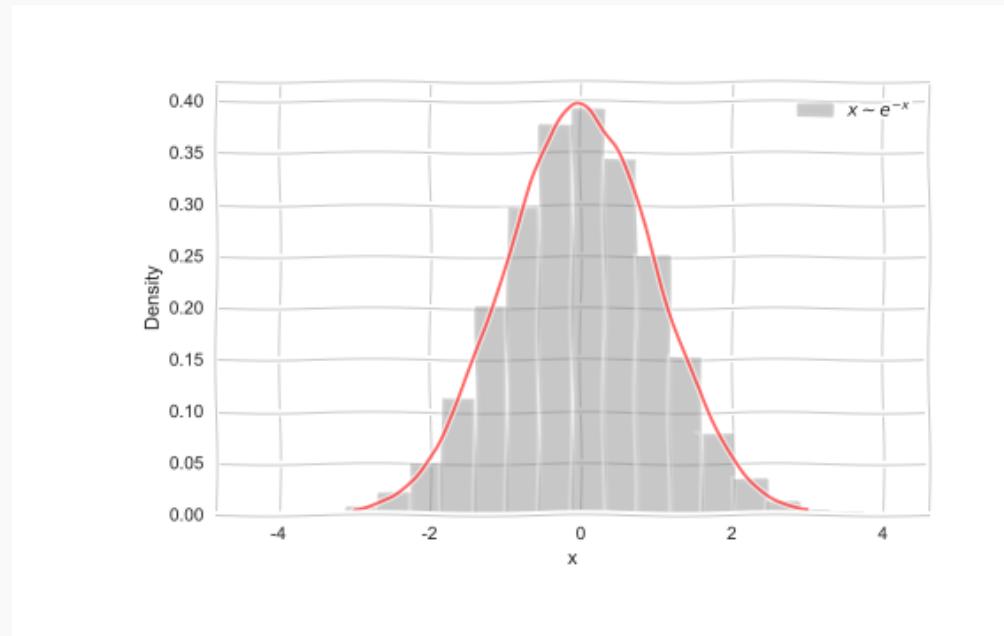
# Generative models

Imagine we want to generate data from a distribution,

e.g.

$$x \sim p(x)$$

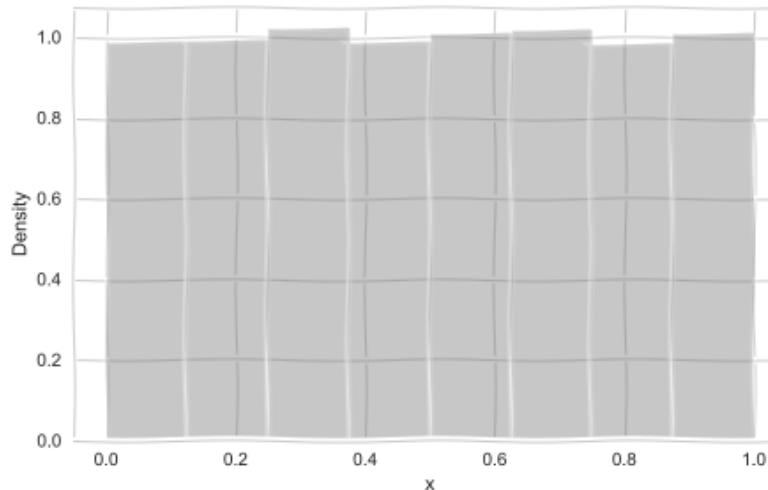
$$x \sim \mathcal{N}(\mu, \sigma)$$



# Generative models

But how do we generate such samples?

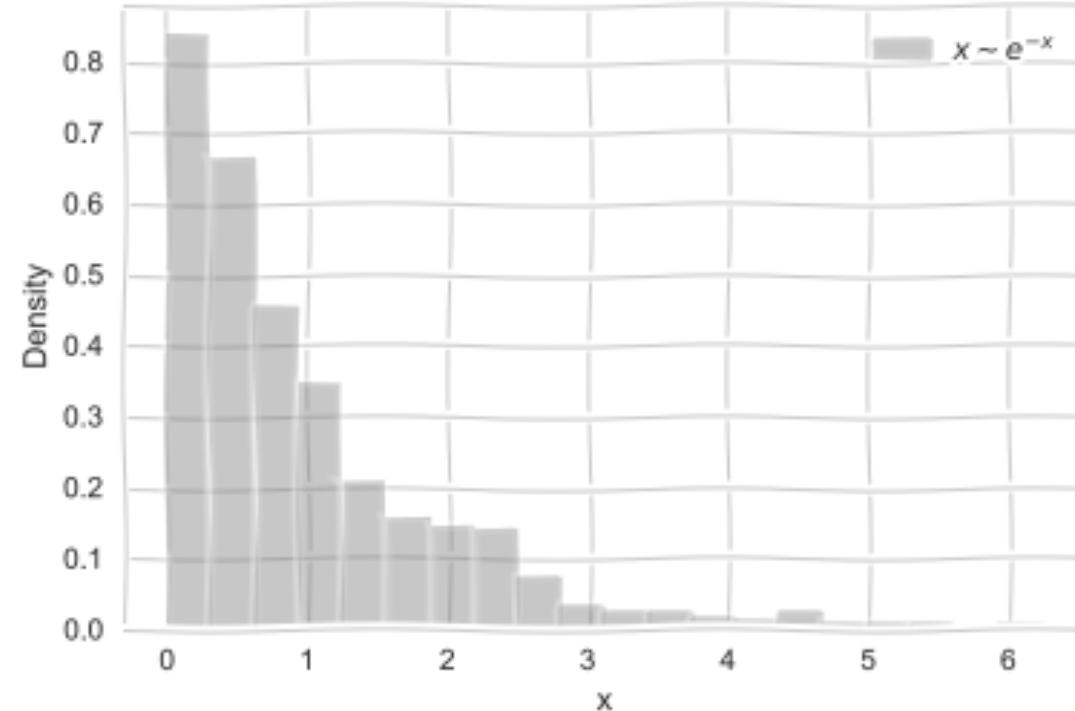
$$z \sim \text{Unif}(0, 1)$$



# Generative models

But how do we generate such samples?

$$z \sim \text{Unif}(0, 1) \quad x = \ln z$$



# Generative models

---

In other words we can think that if we choose  $z \sim \text{Uniform}$  then there is a mapping:

$$x = f(z)$$

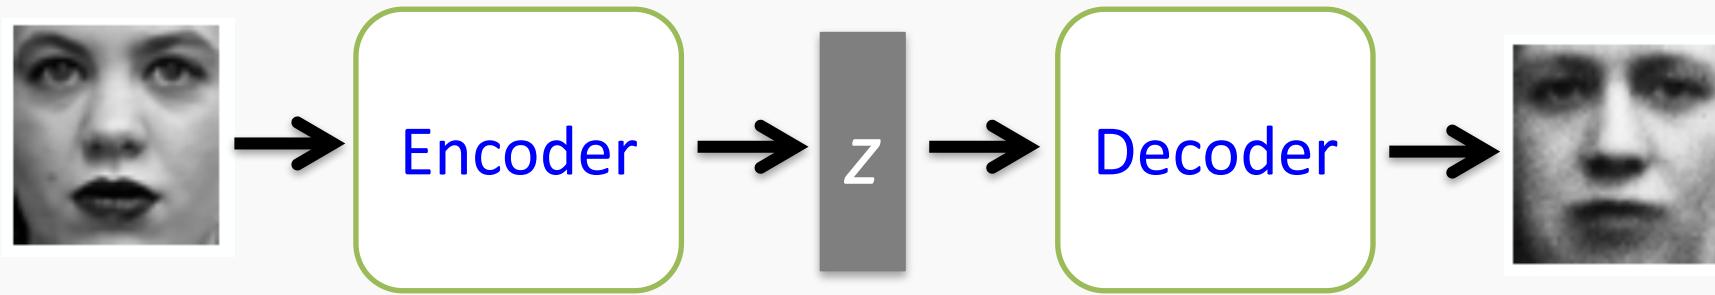
such as:

$$x \sim p(x)$$

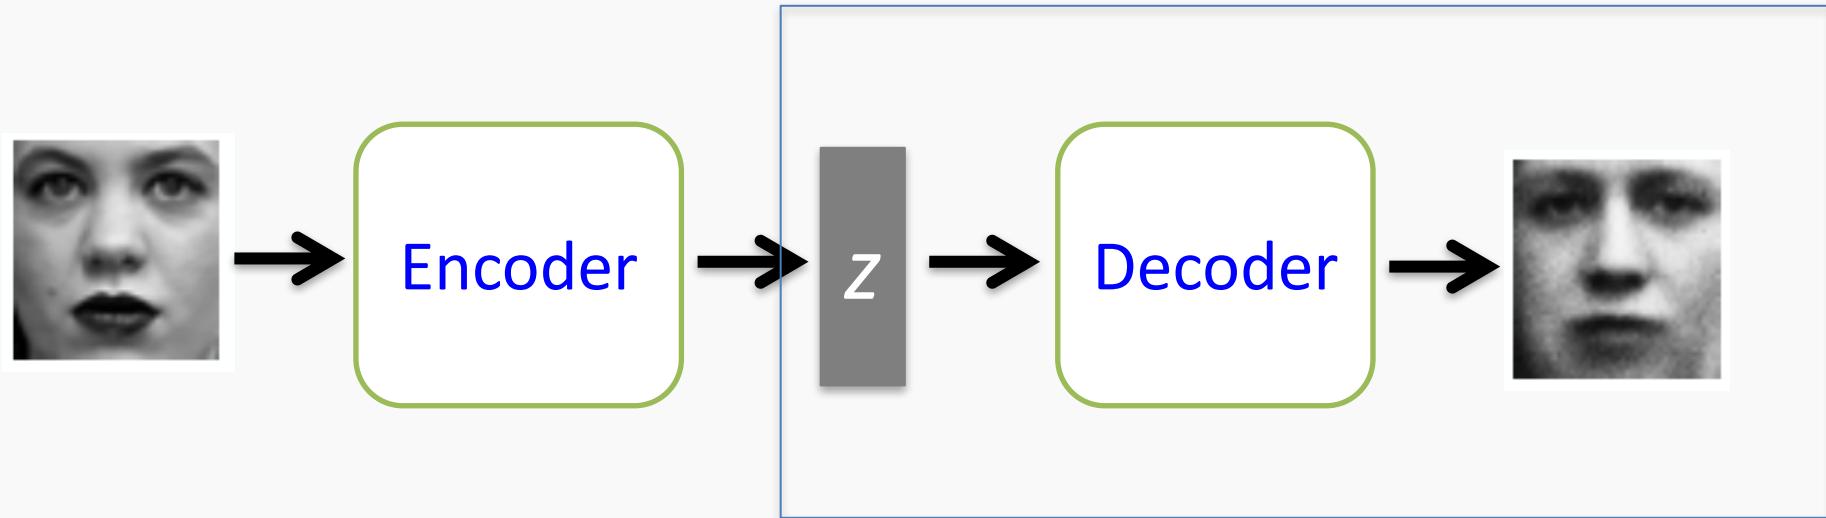
where in general  $f$  is some complicated function.

We already know that Neural Networks are great in learning complex functions.

# Variational Autoencoders

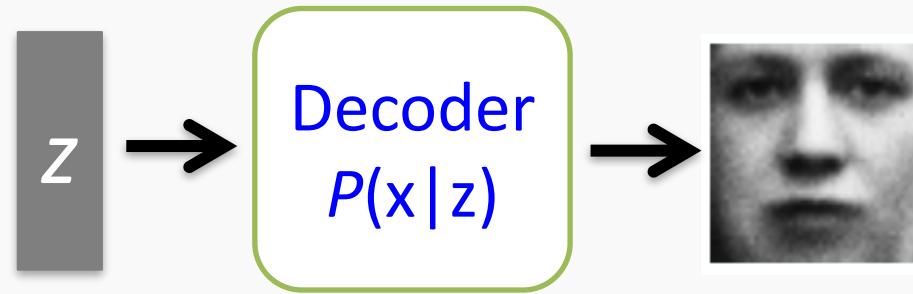


# Variational Autoencoders

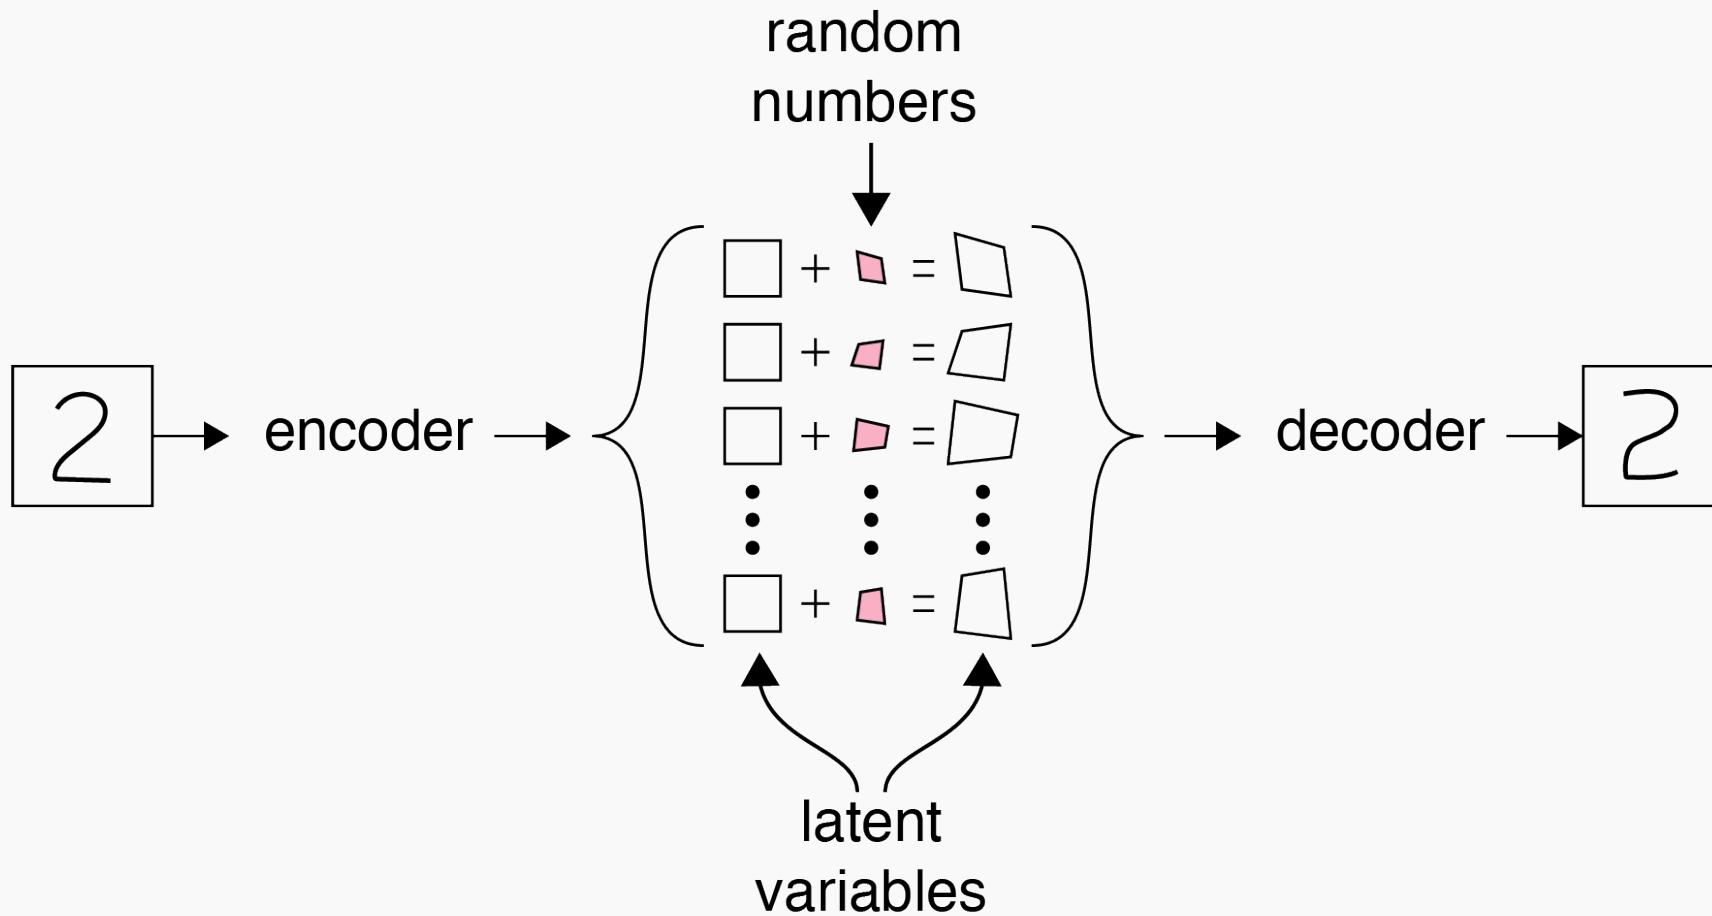


# Variational Autoencoders

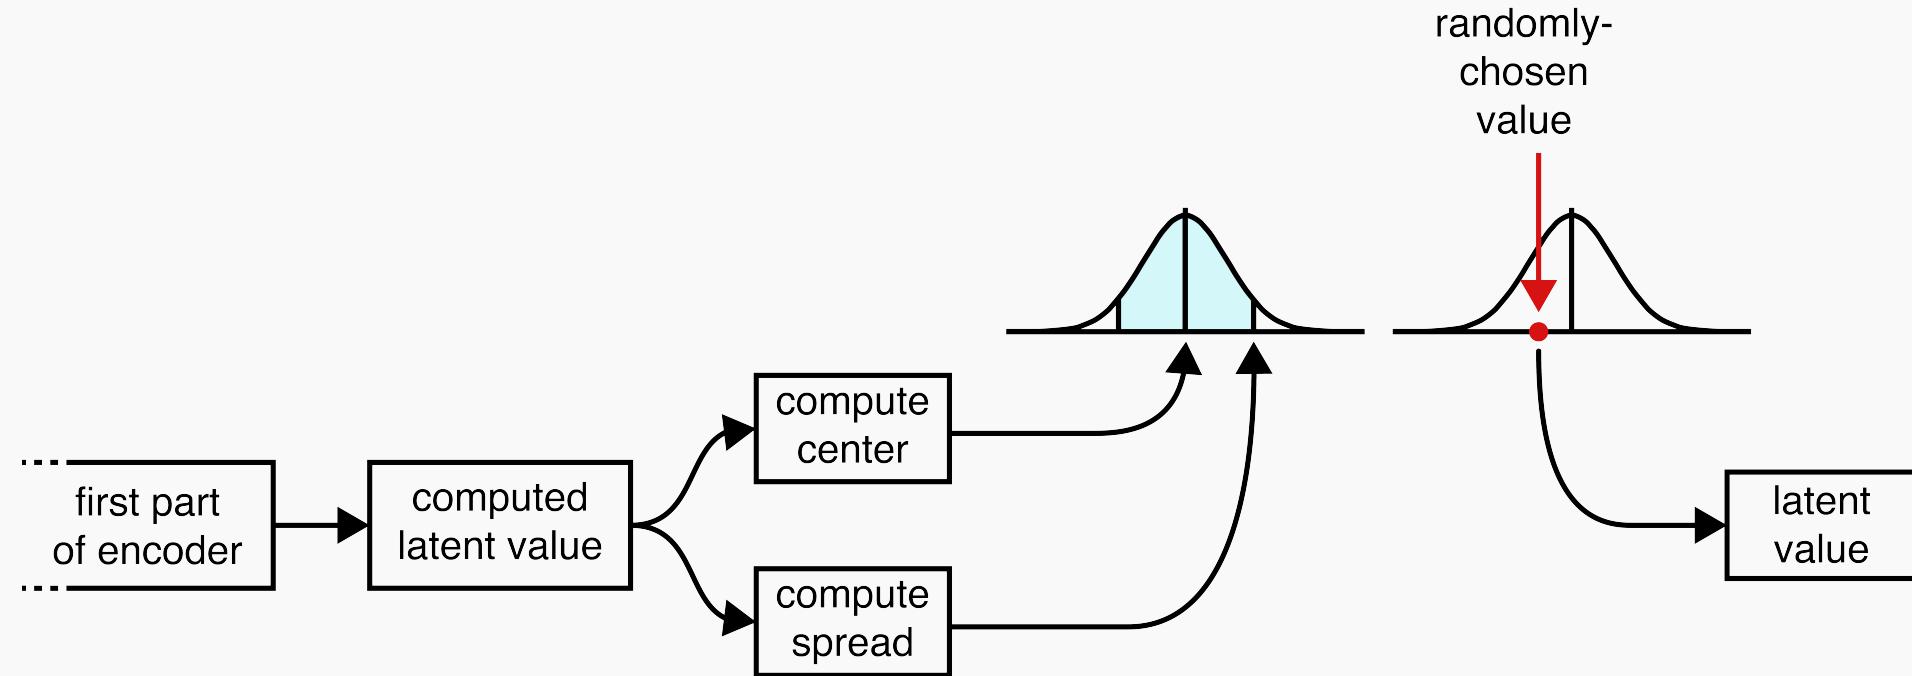
Sample from  $P(z)$   
*Standard Gaussian*



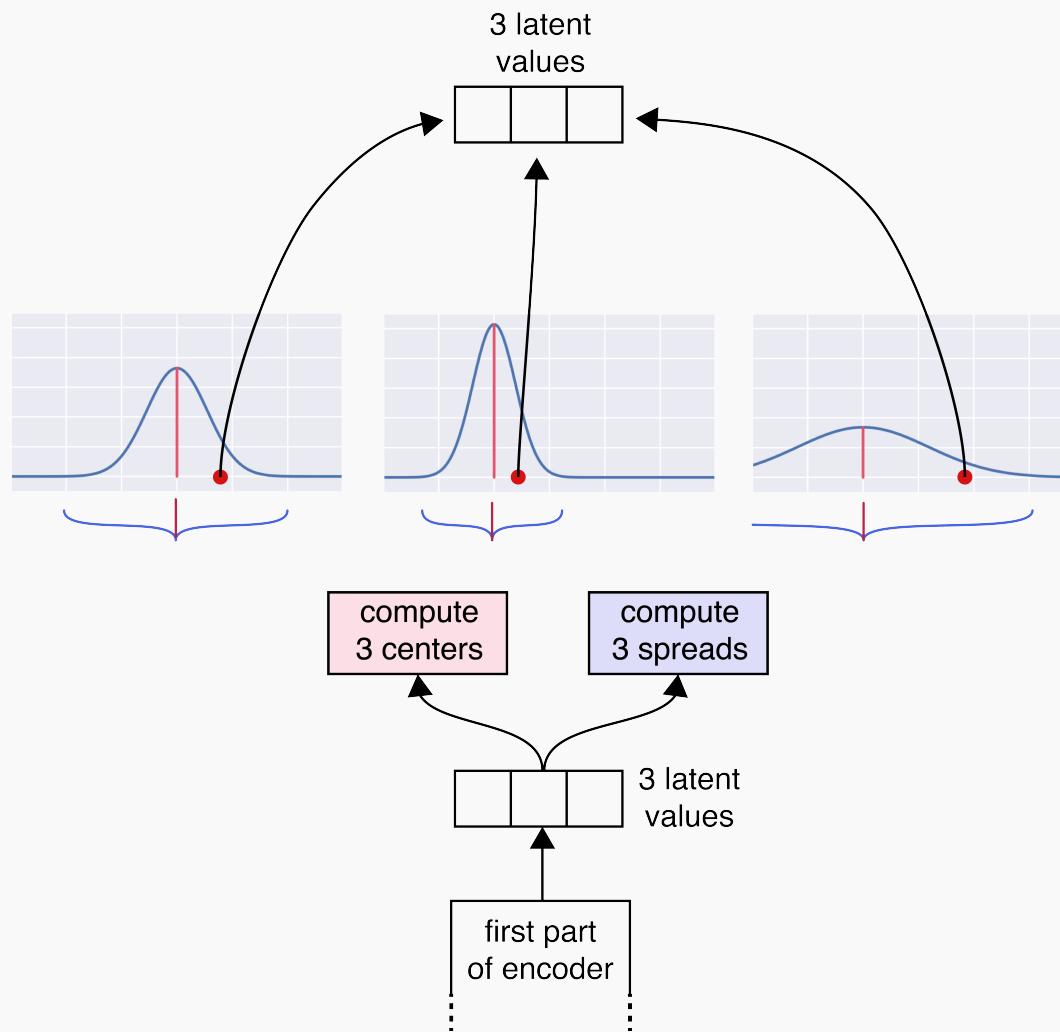
# Variational Autoencoders



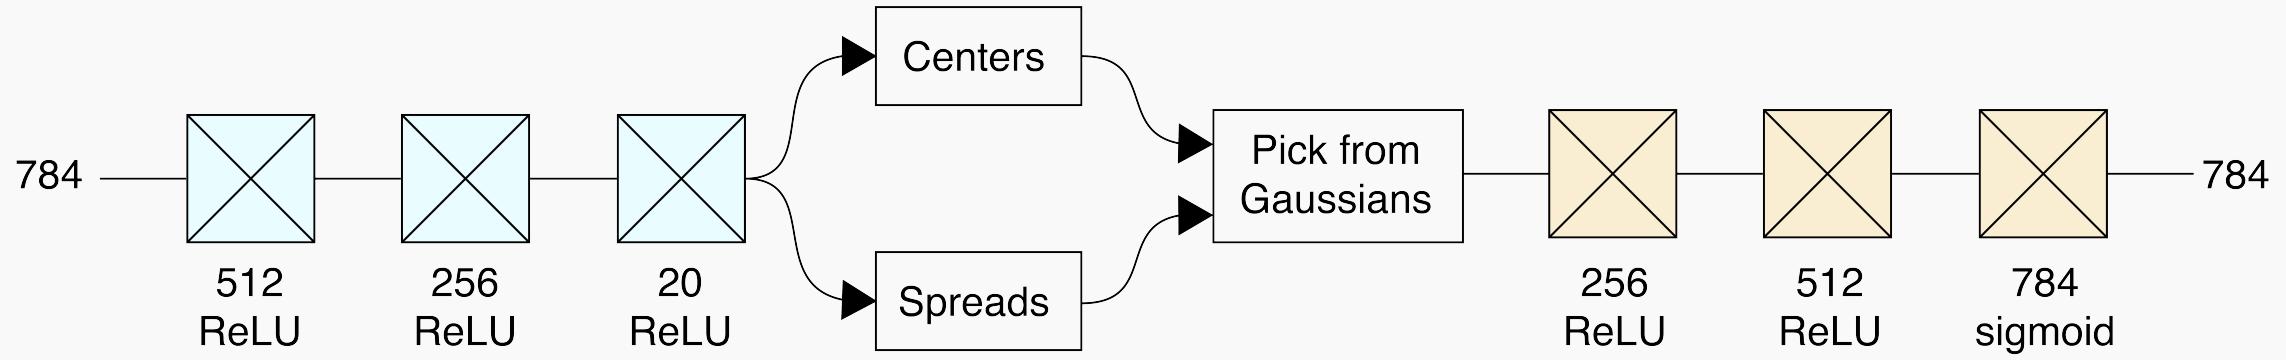
# Variational Autoencoders



# Variational Autoencoders



# Variational Autoencoders



# Outline

---

Motivation for Variational Autoencoders (VAE)

Mechanics of VAE

**Seperability of VAE**

The math behind everything

Generative models



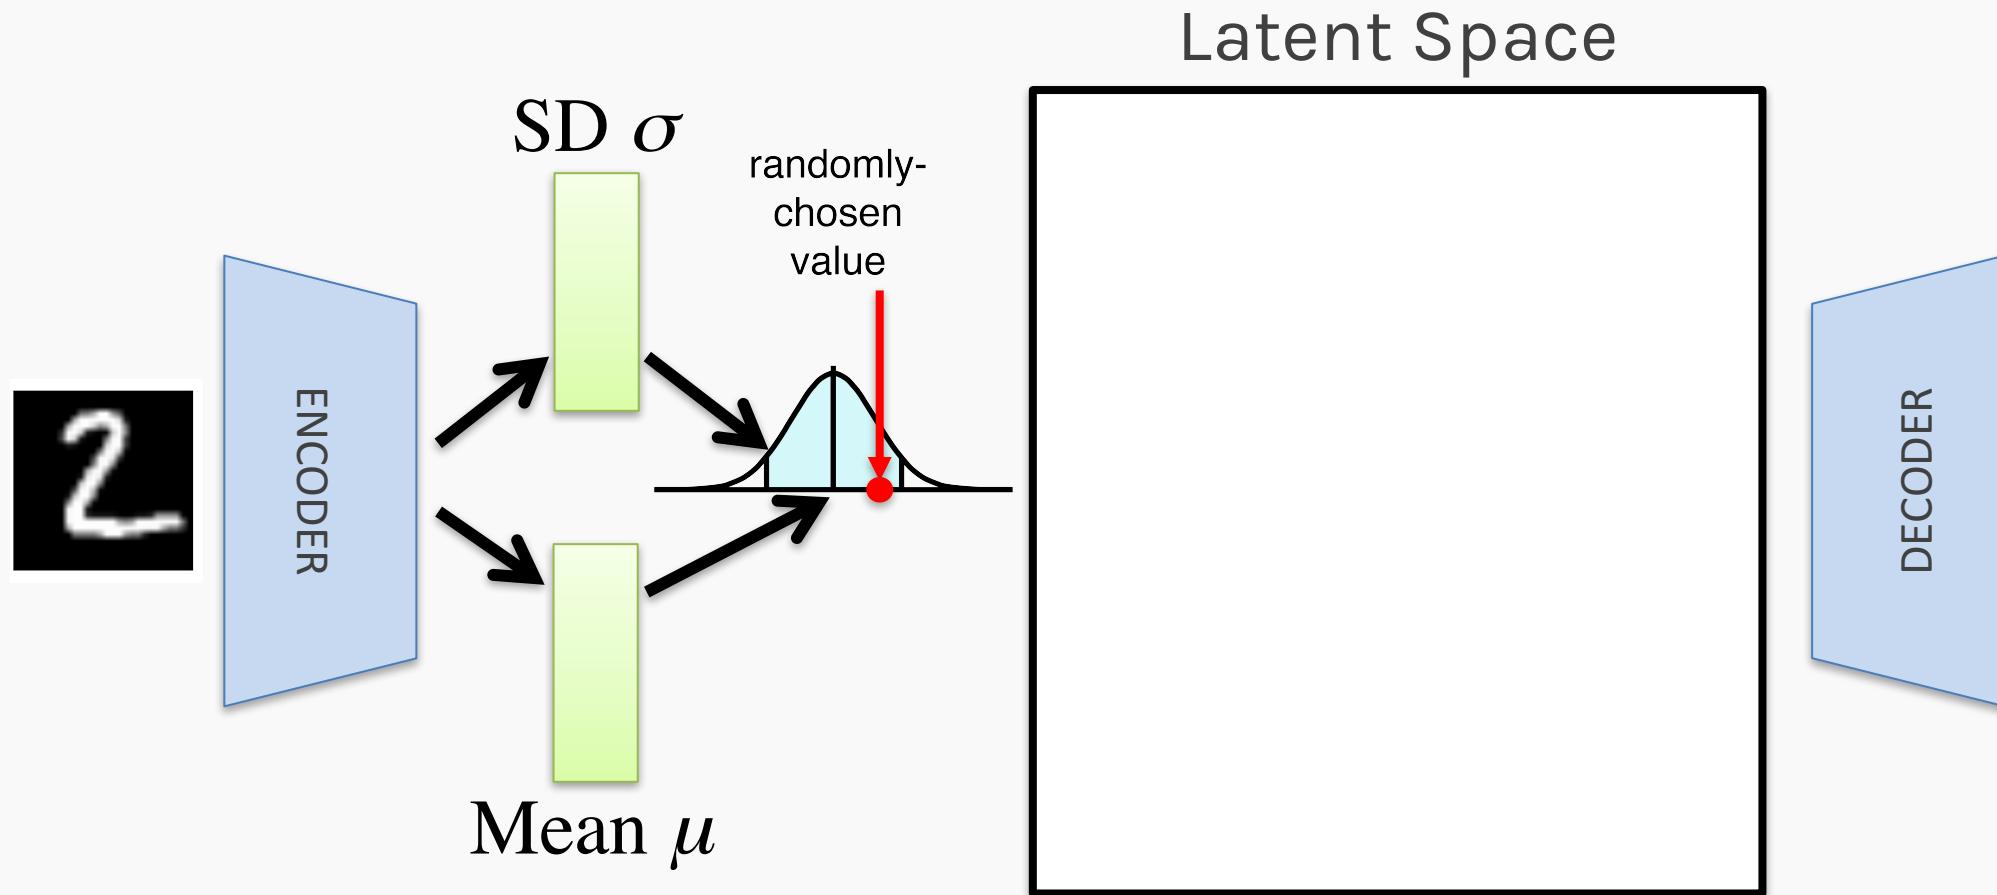
# Variational Autoencoders

---

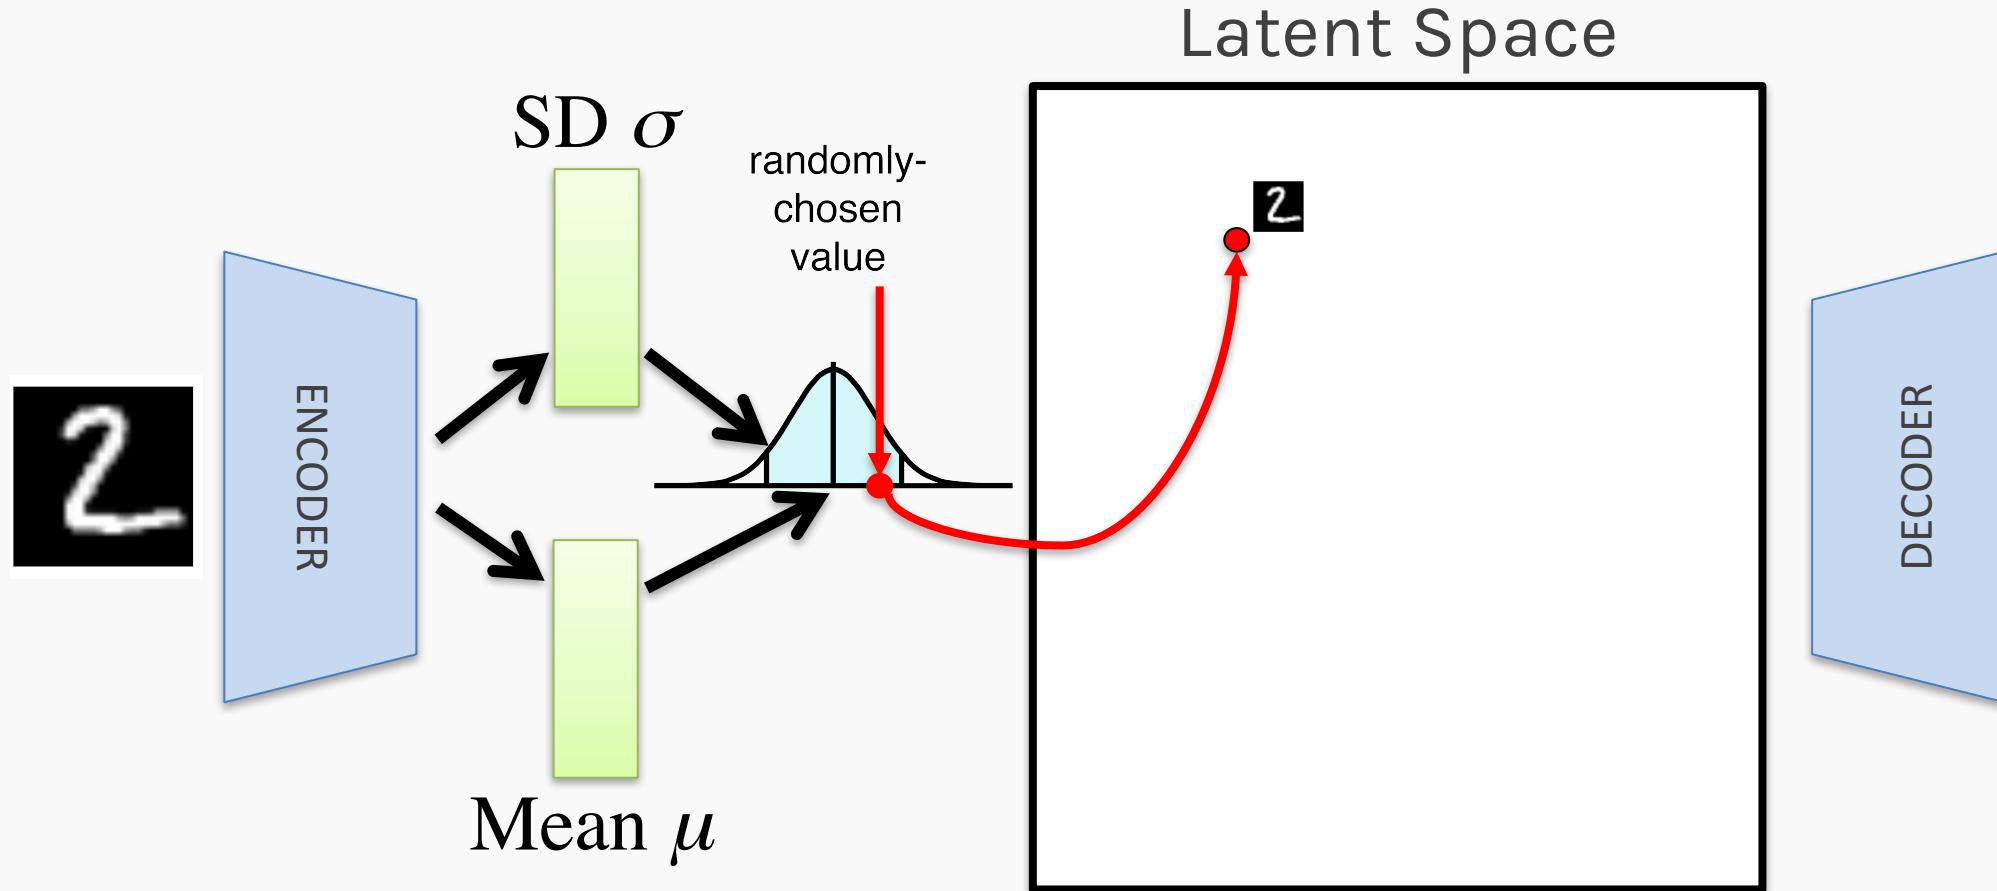
2 2 2 2 2 2 2  
2 2 2 2 2 2 2  
2 2 2 2 2 2 2  
2 2 2 2 2 2 2  
2 2 2 2 2 2 2  
2 2 2 2 2 2 2  
2 2 2 2 2 2 2  
2 2 2 2 2 2 2



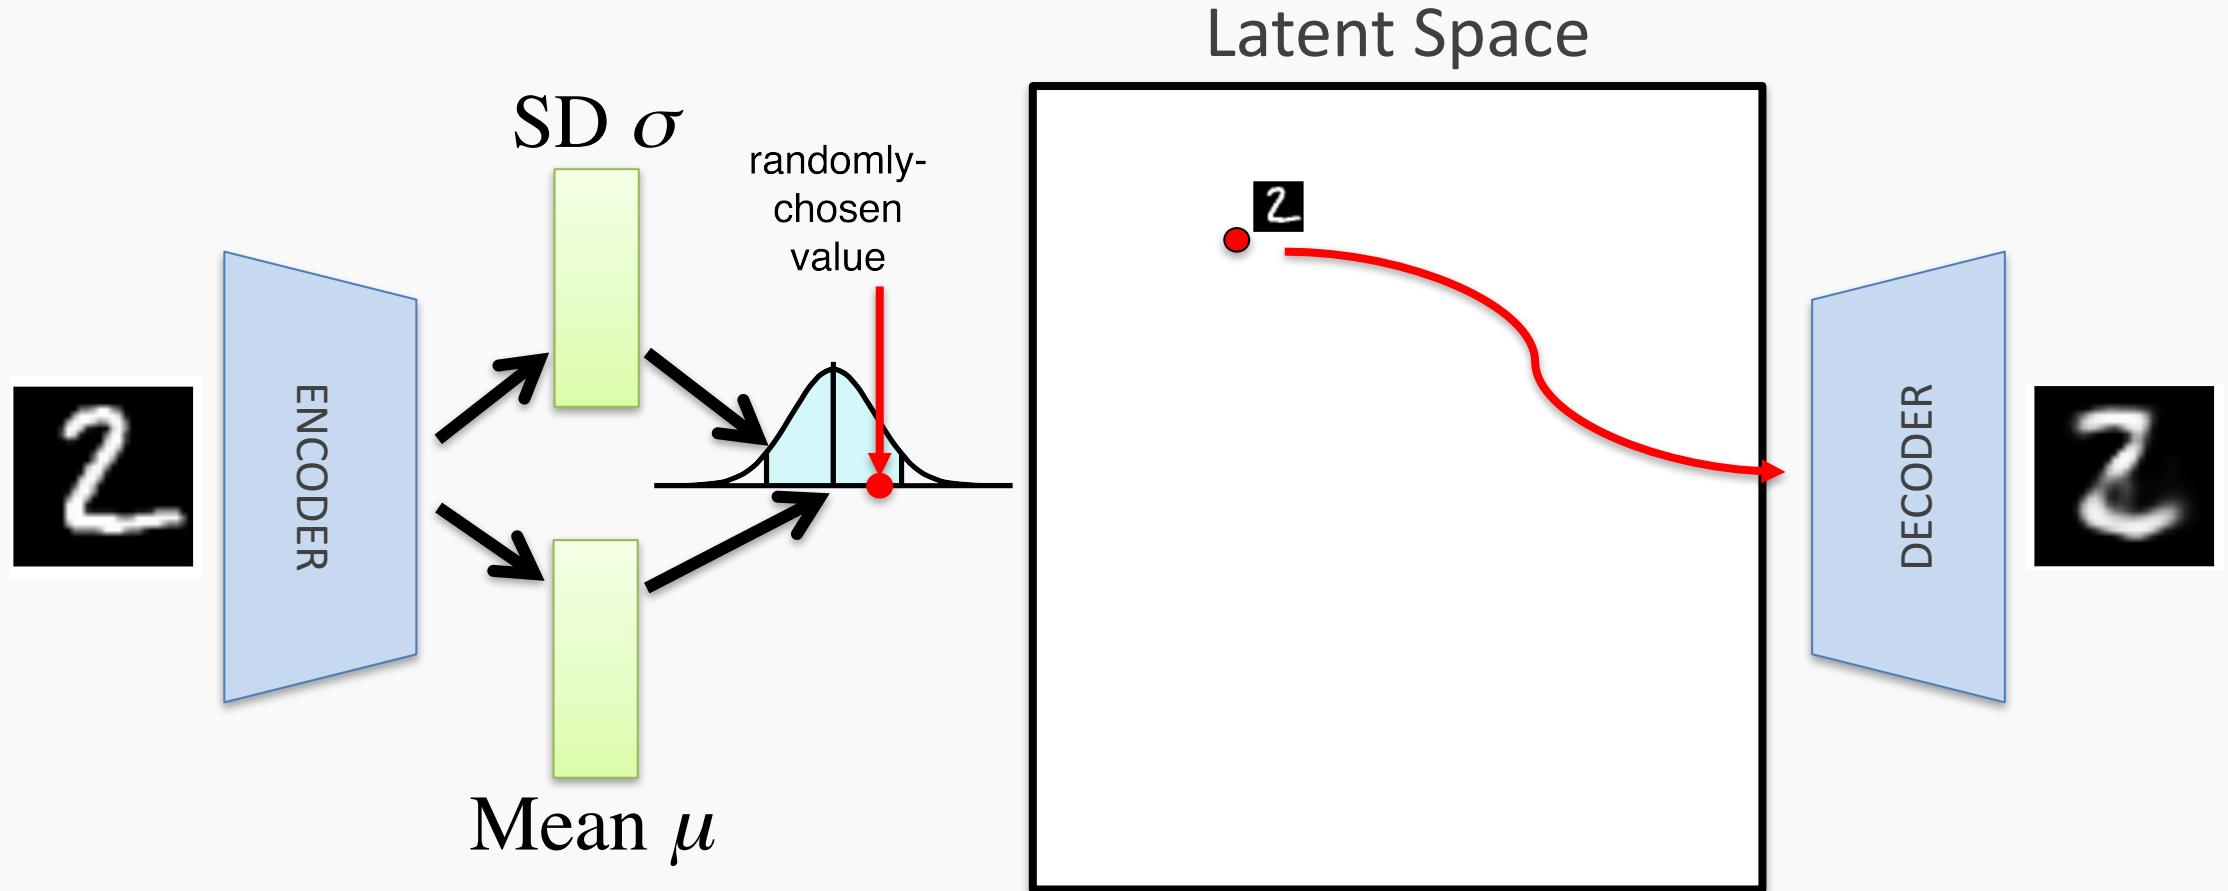
# Blending Latent Variables



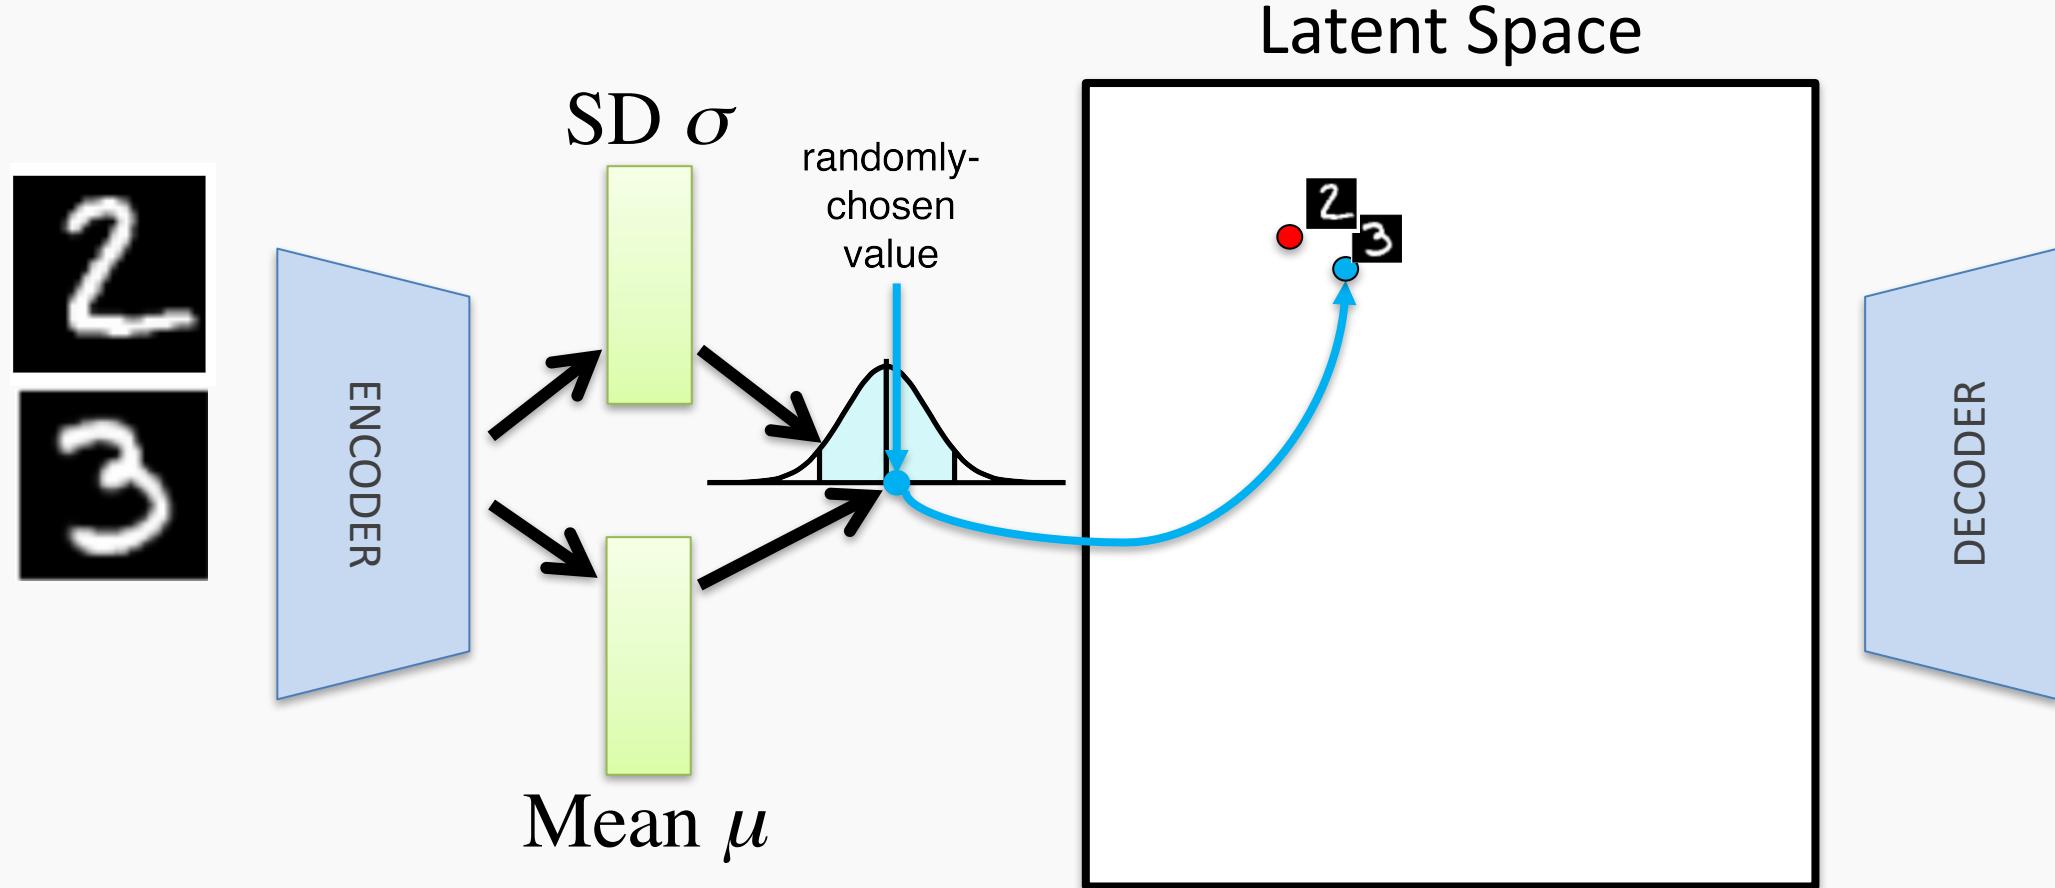
# Blending Latent Variables



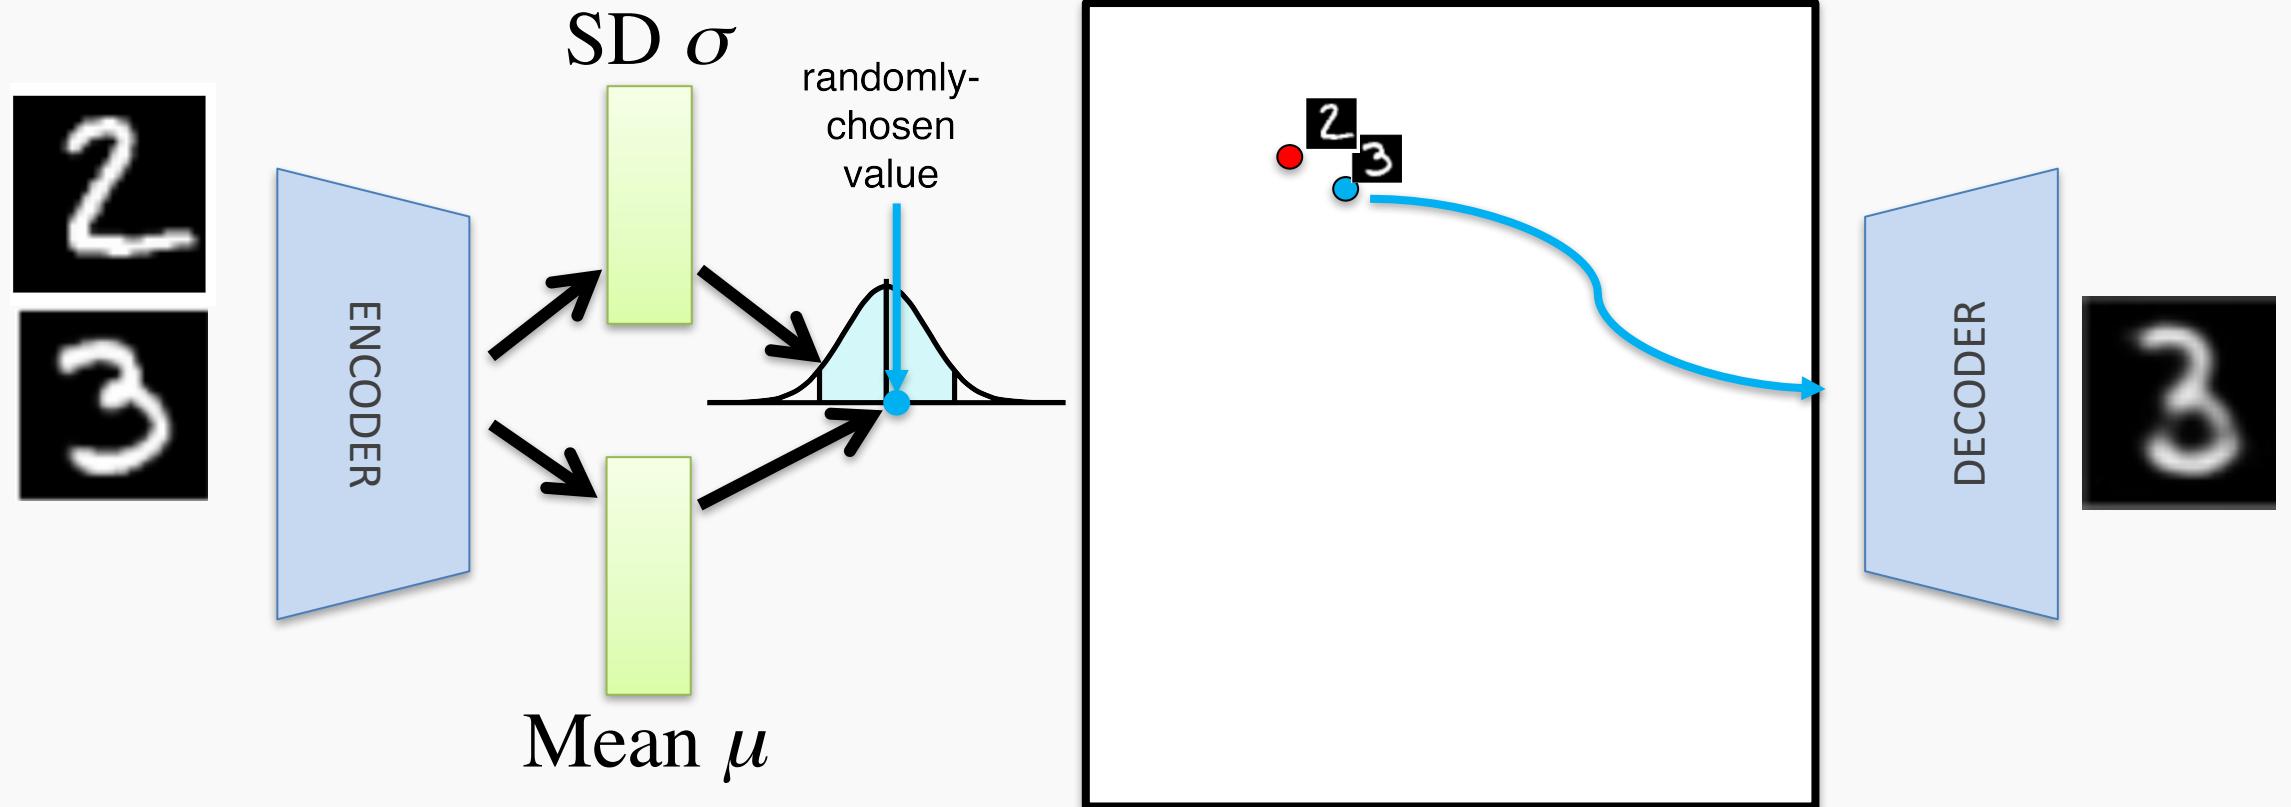
# Blending Latent Variables



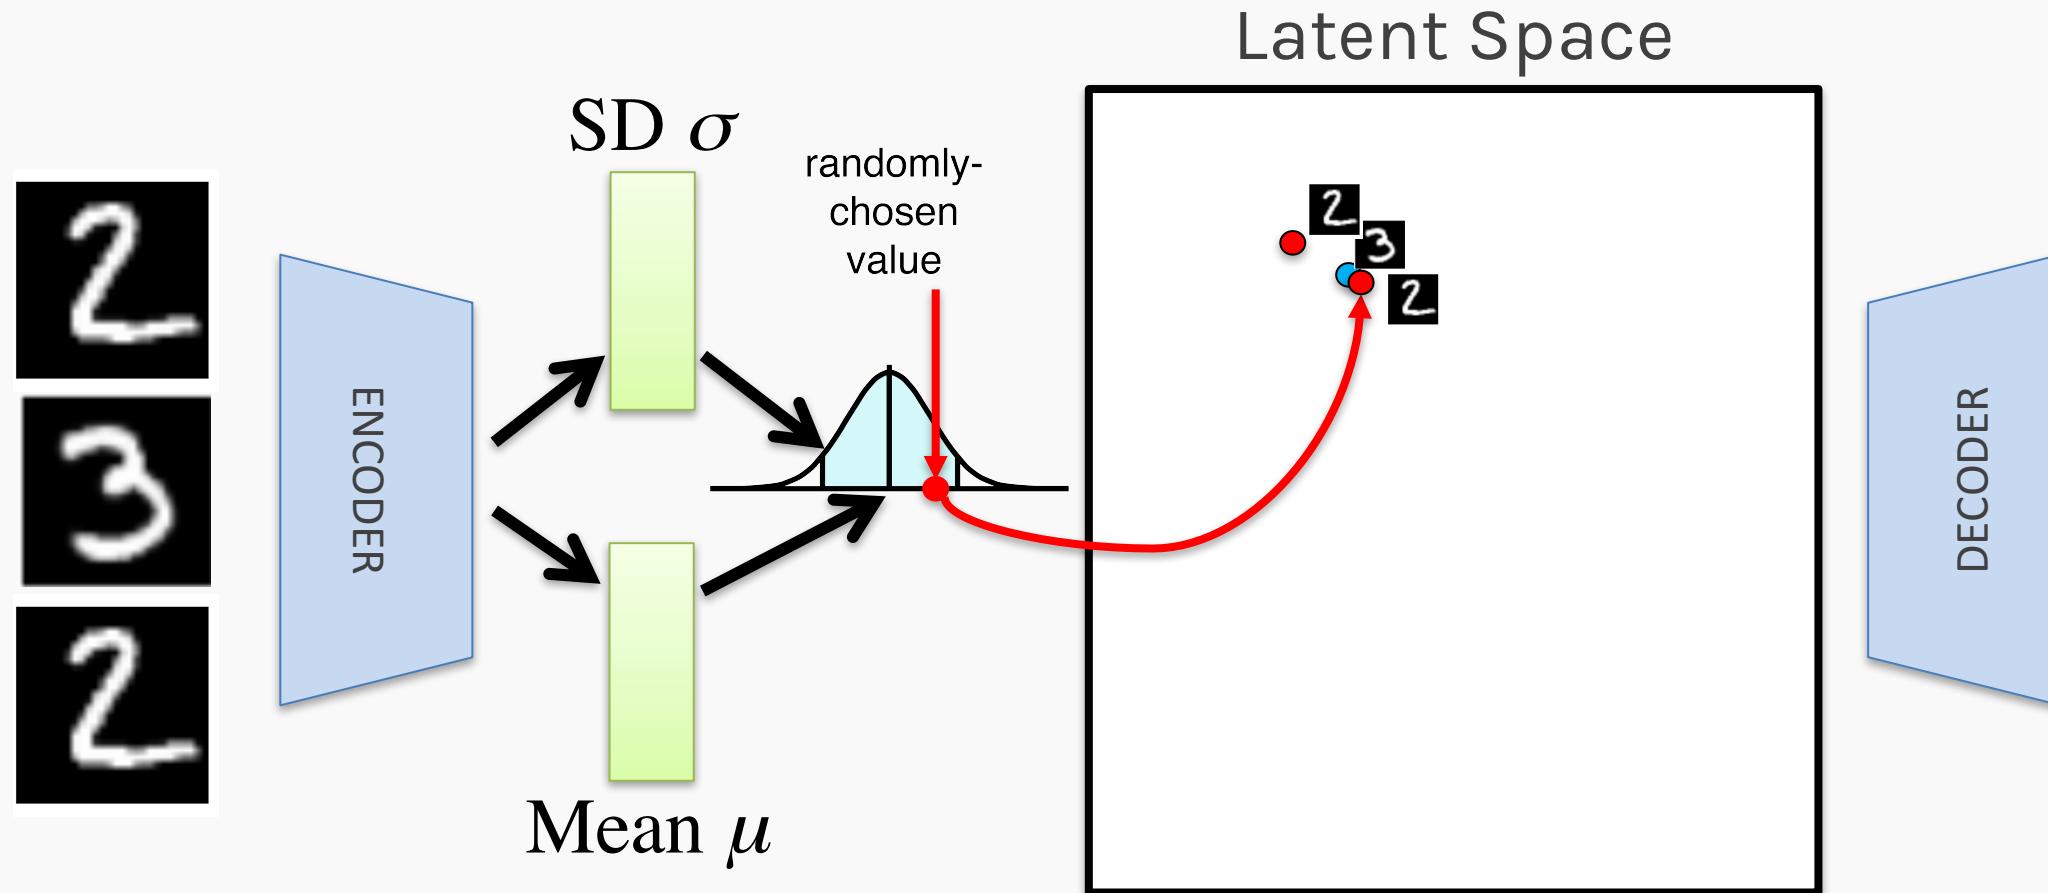
# Blending Latent Variables



# Blending Latent Variables

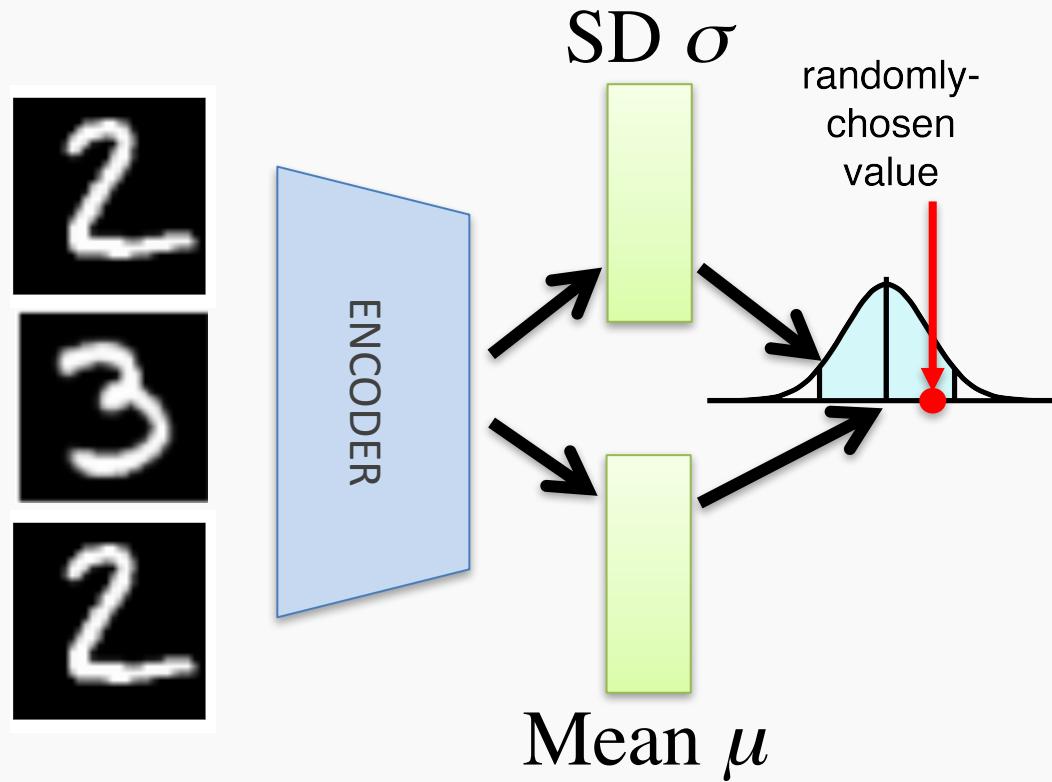


# Blending Latent Variables

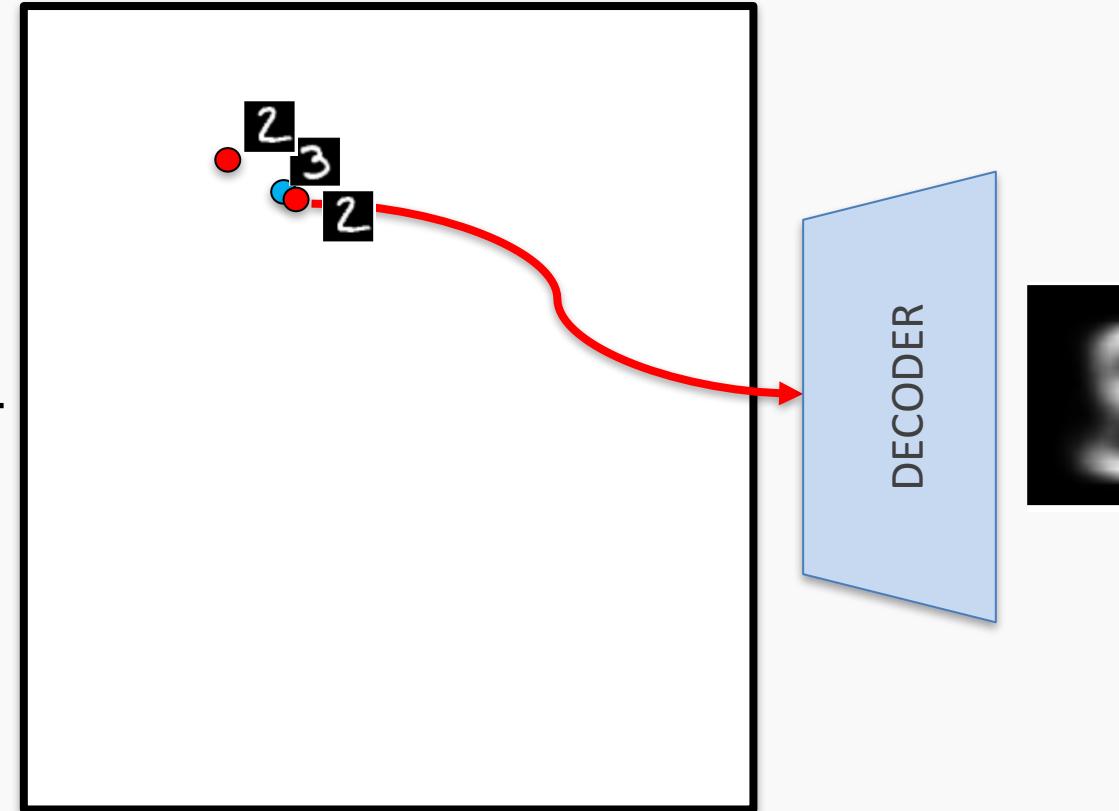


# Blending Latent Variables

Train with 1<sup>st</sup> sample again

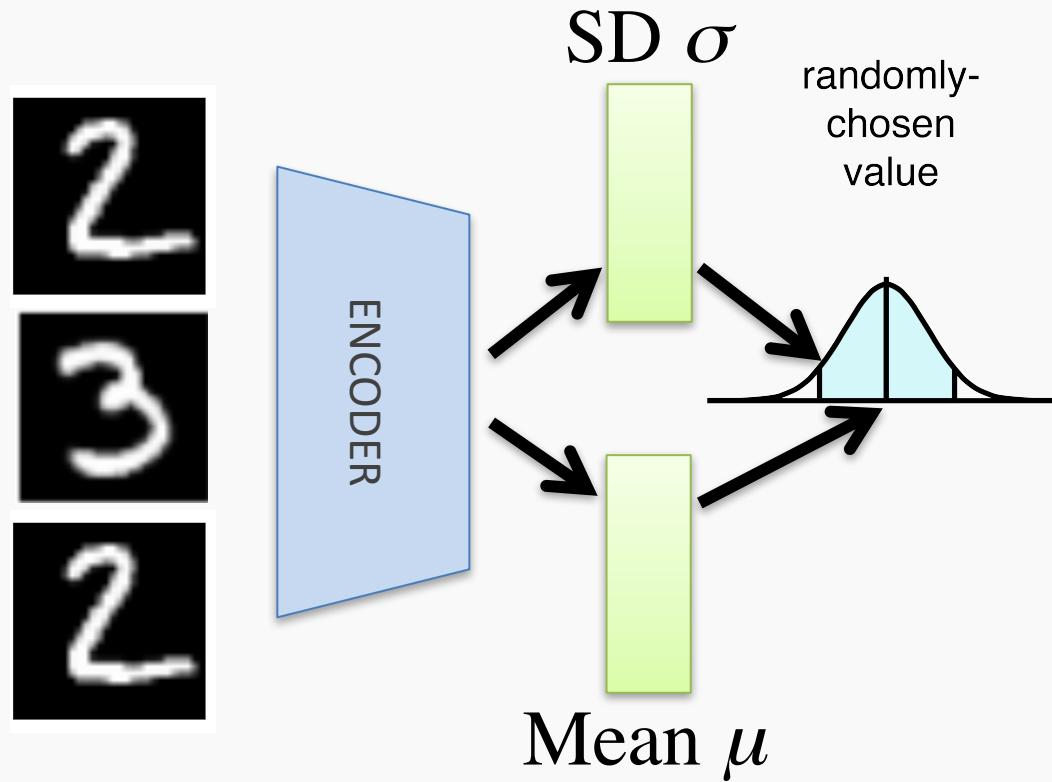


Latent Space

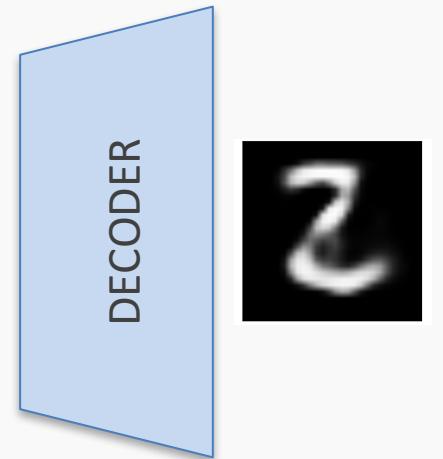
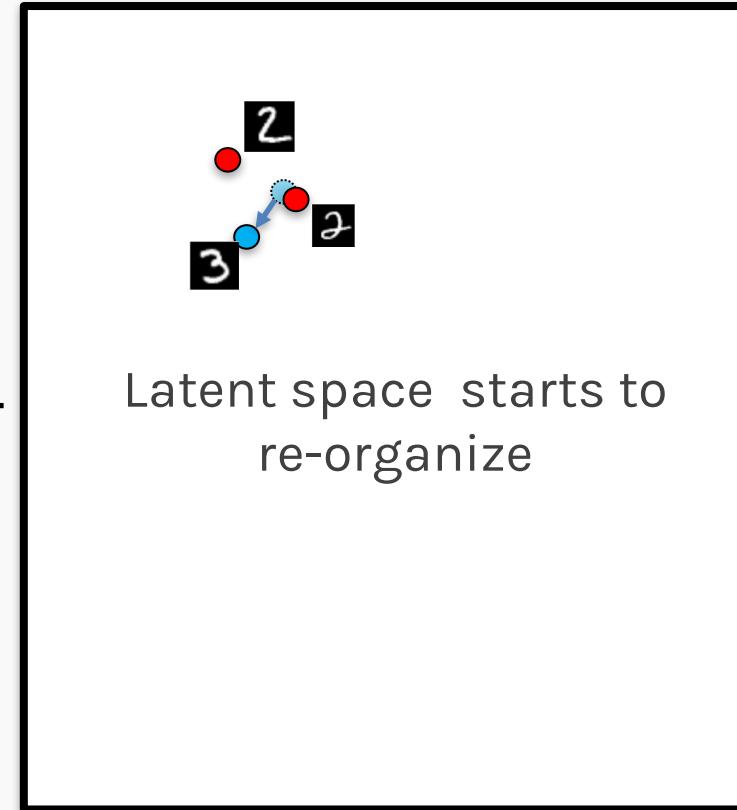


# Blending Latent Variables

Train with 1<sup>st</sup> sample again

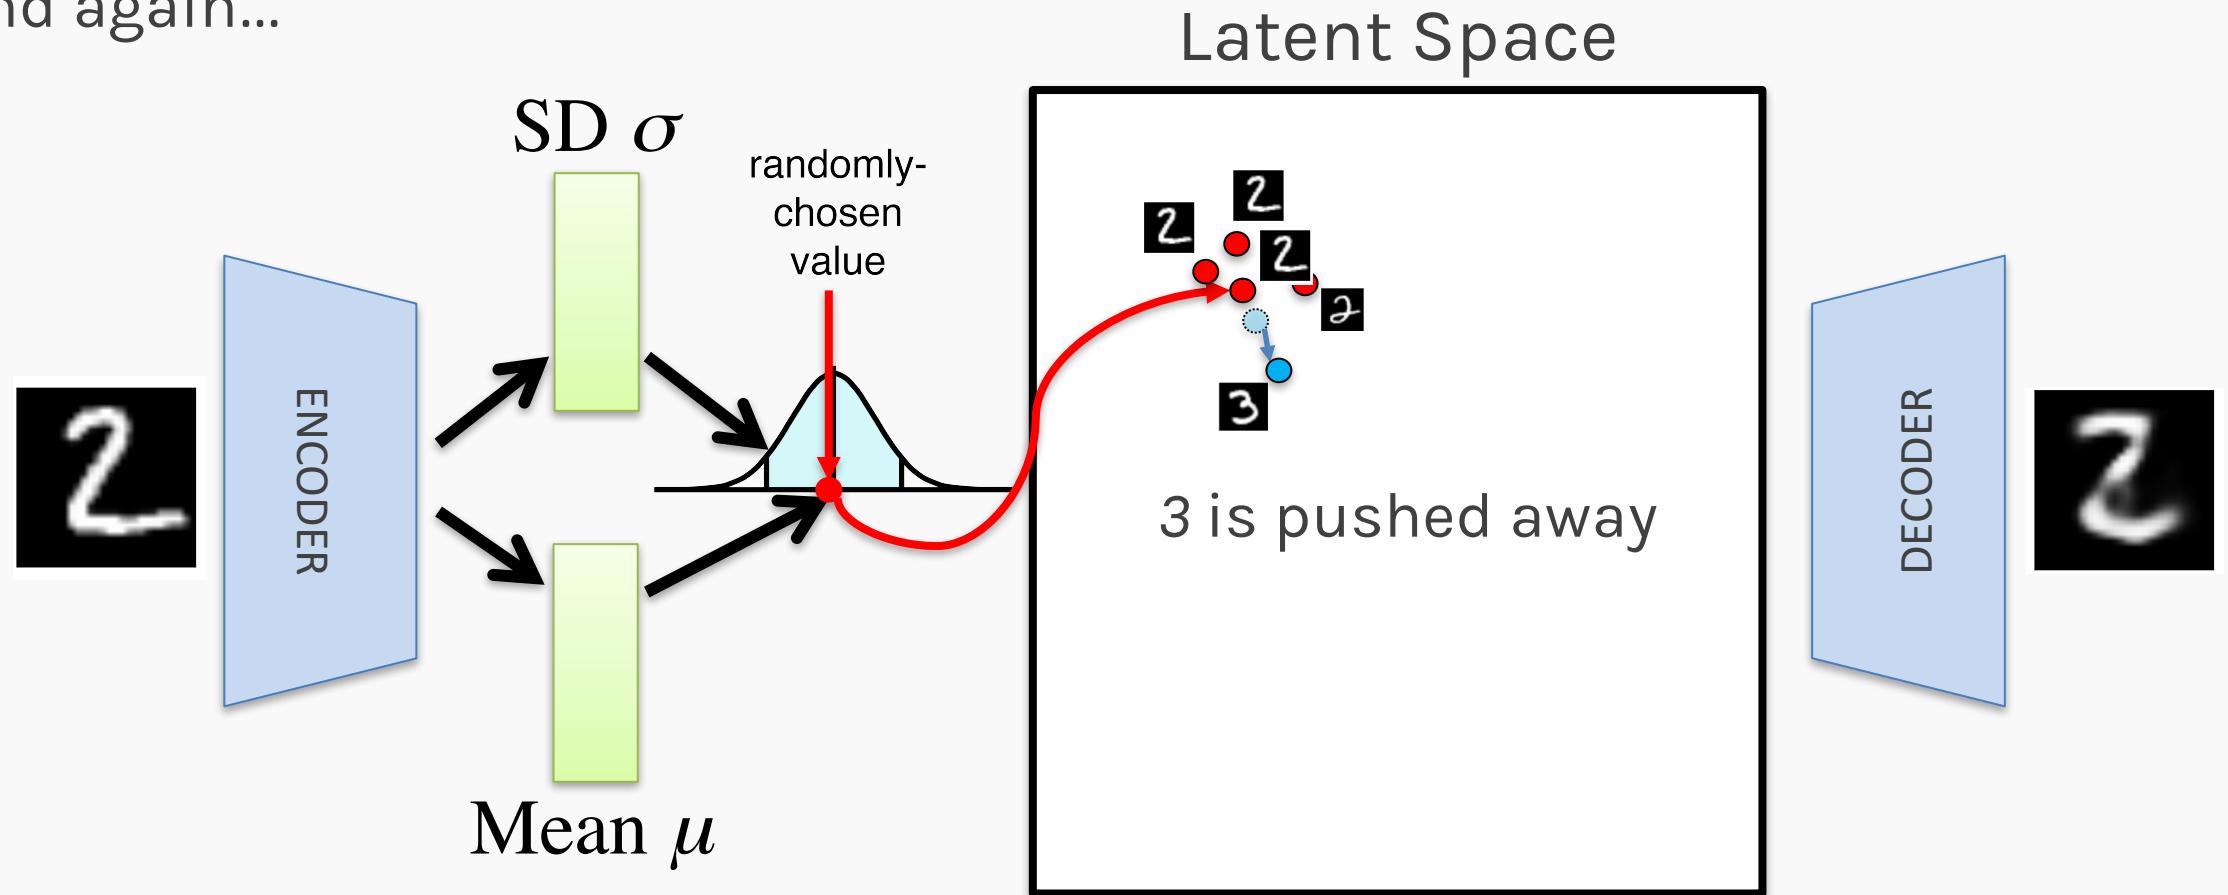


Latent Space



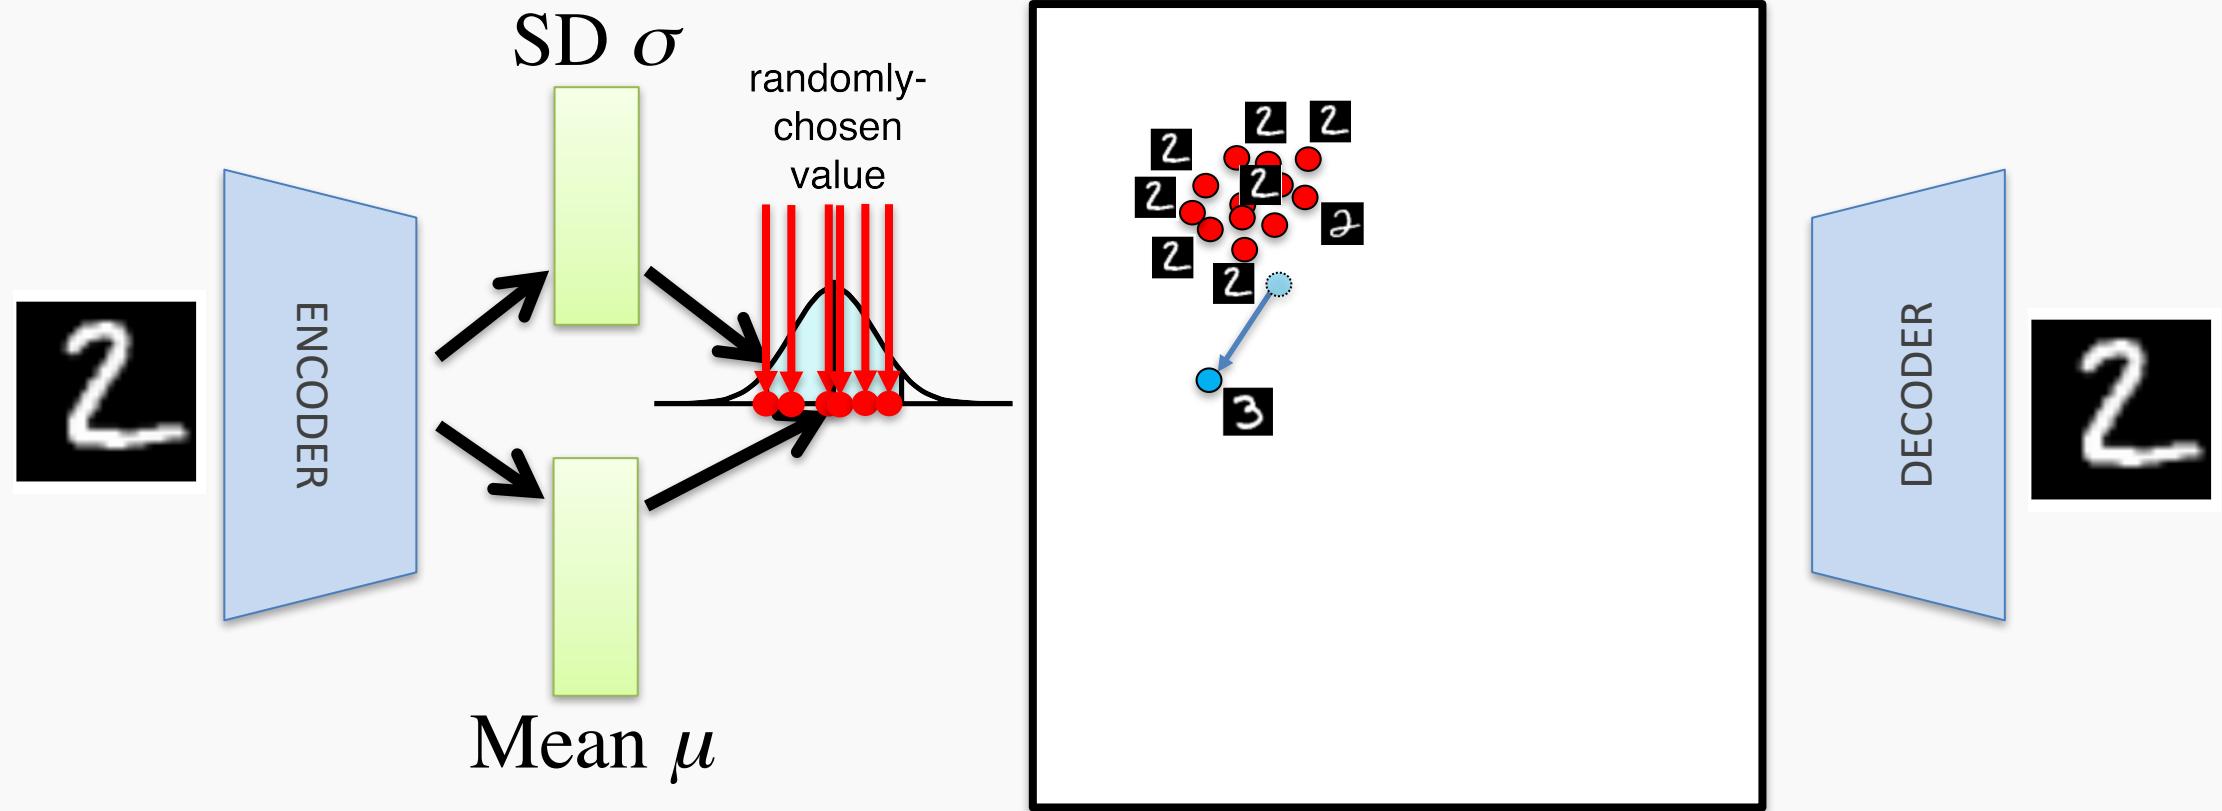
# Blending Latent Variables

And again...



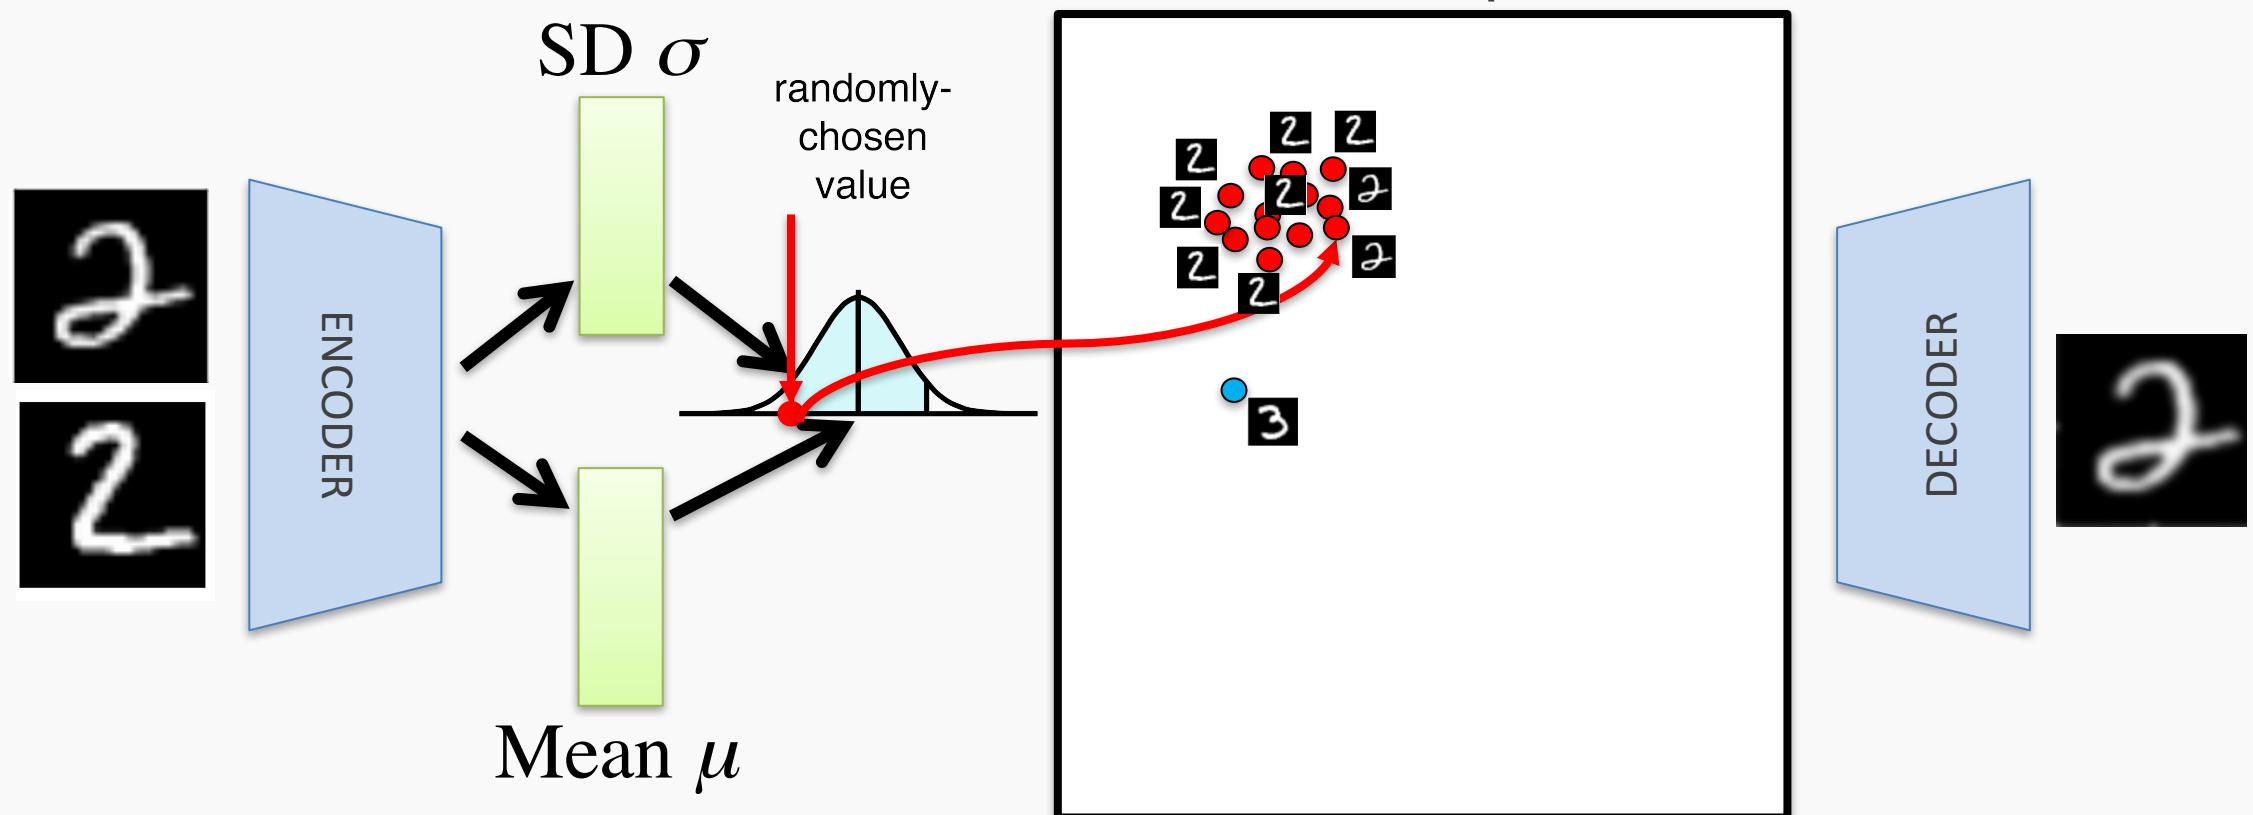
# Blending Latent Variables

Many times...



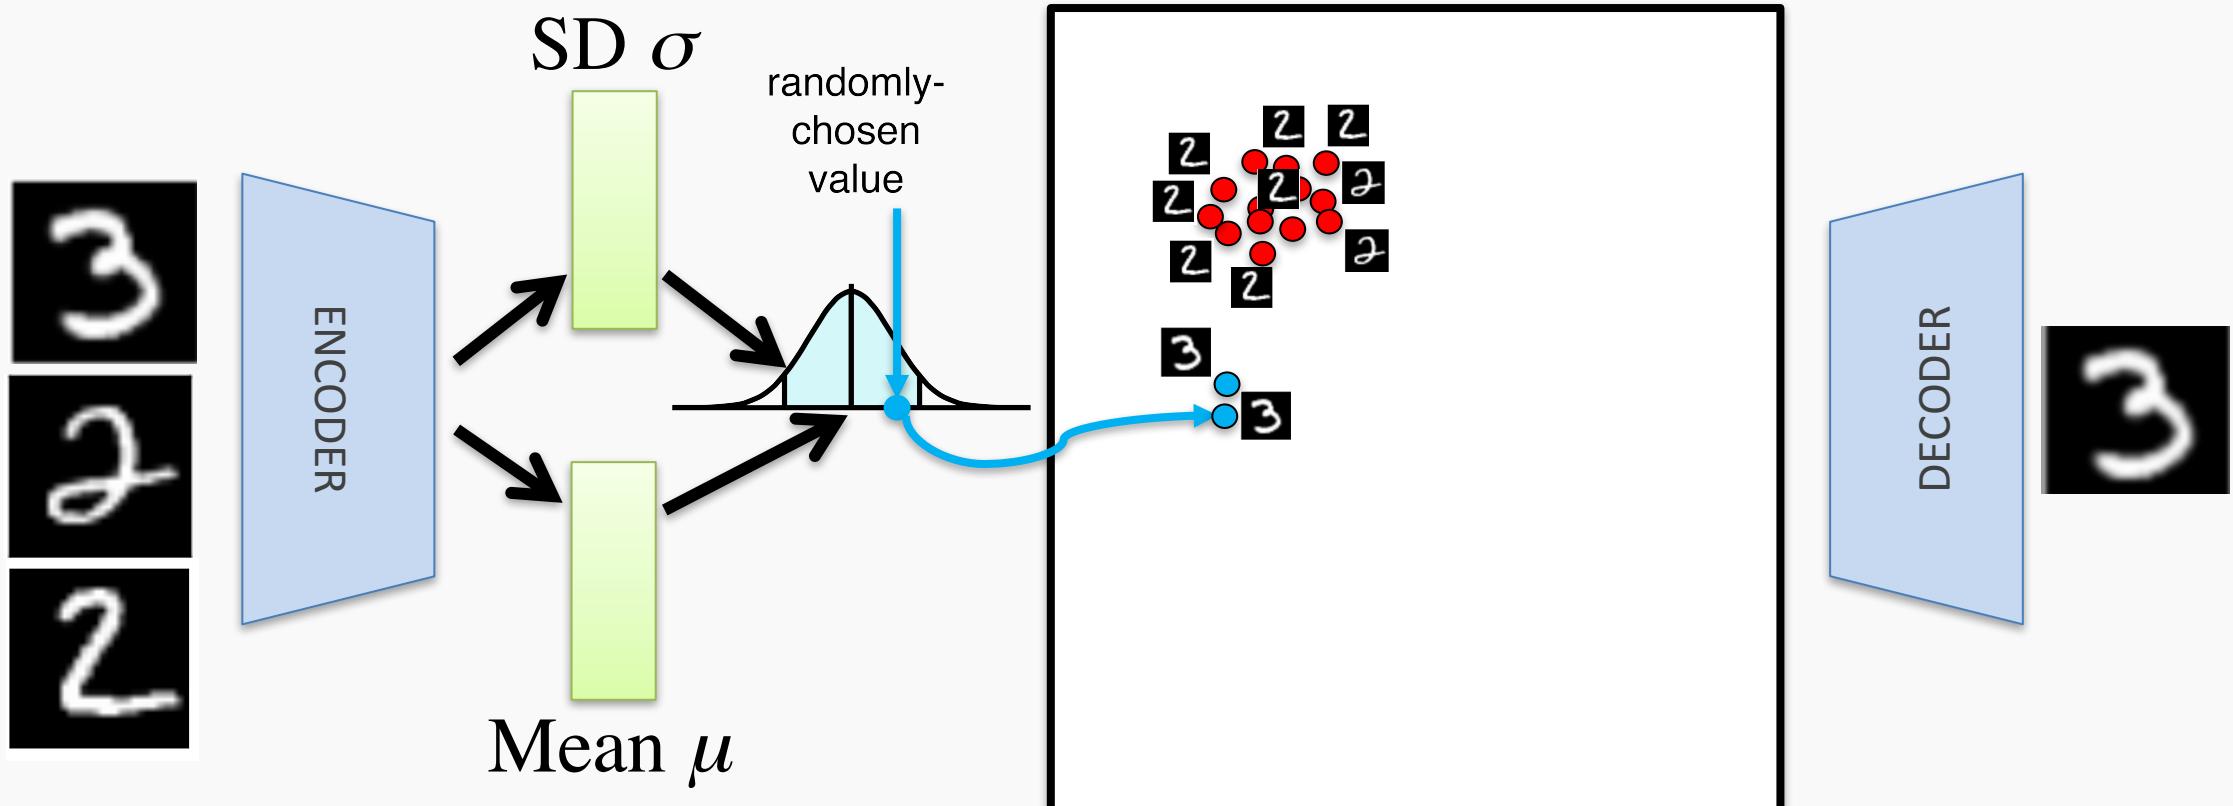
# Blending Latent Variables

Now lets test again



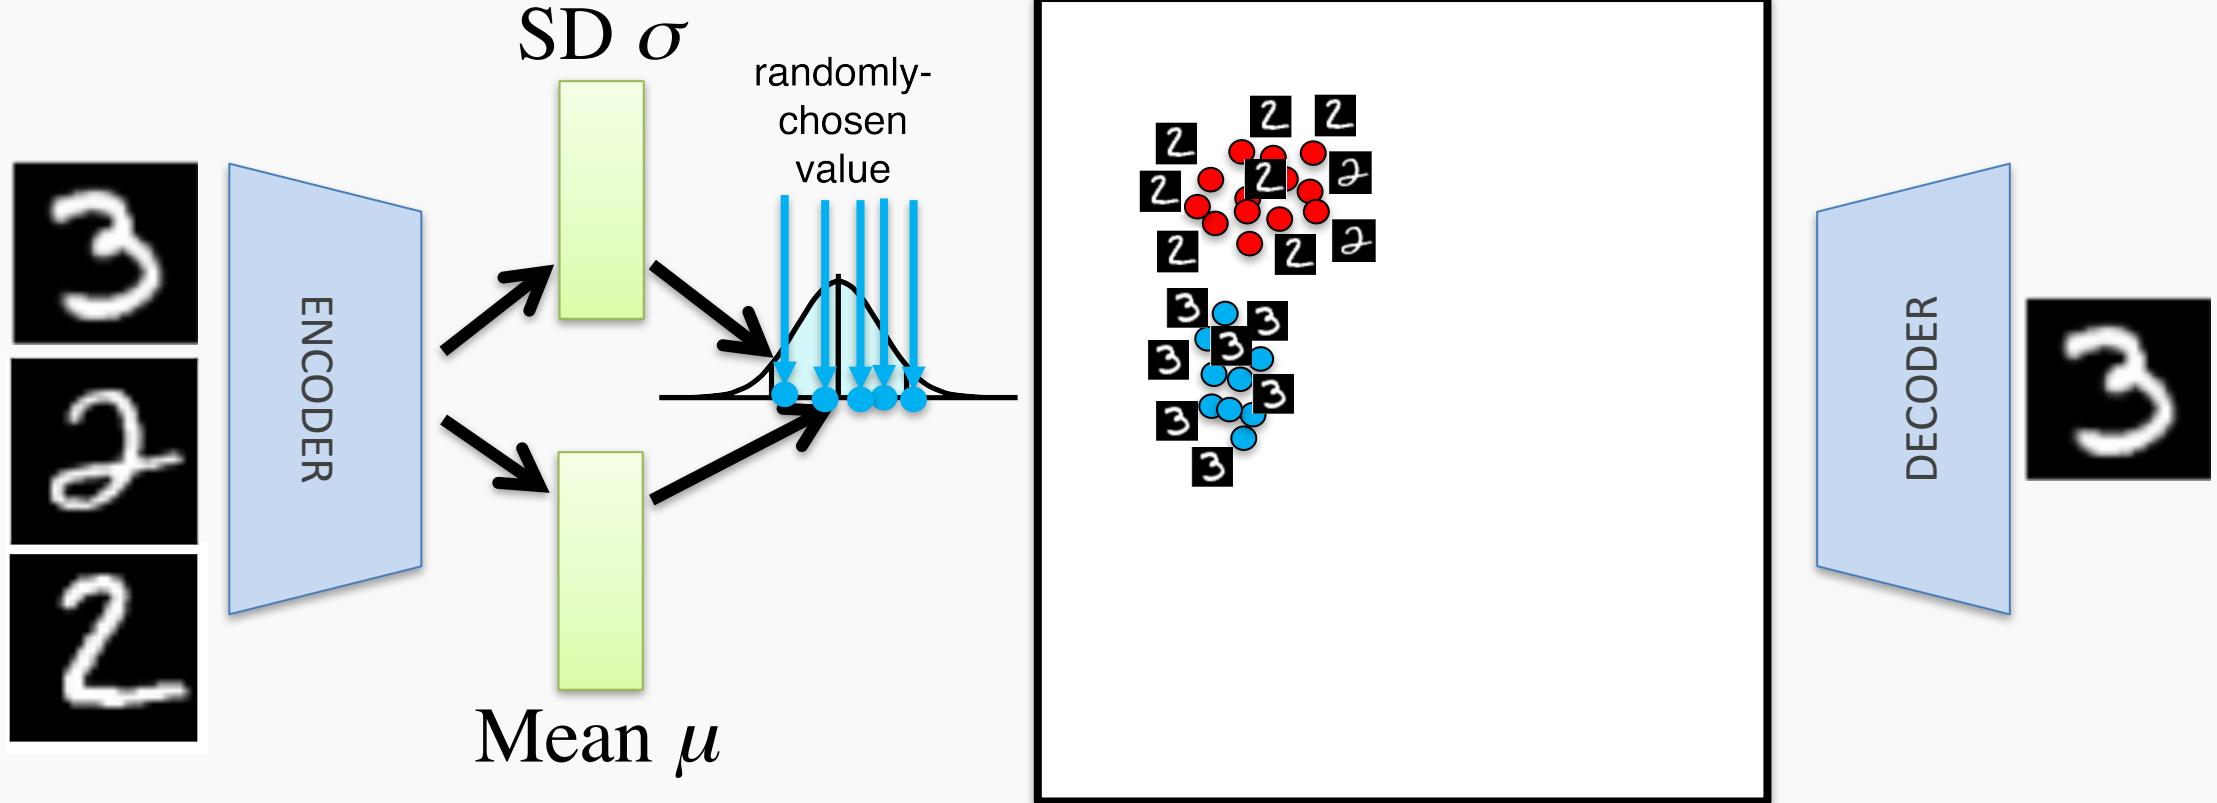
# Blending Latent Variables

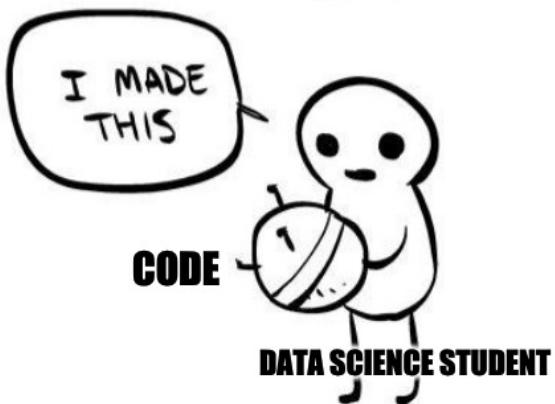
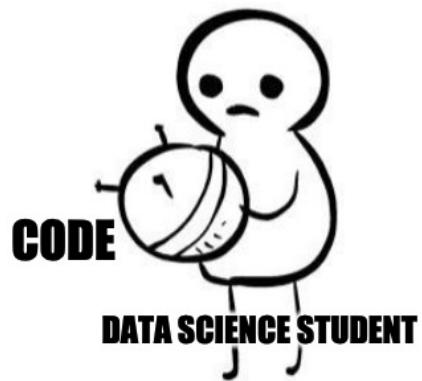
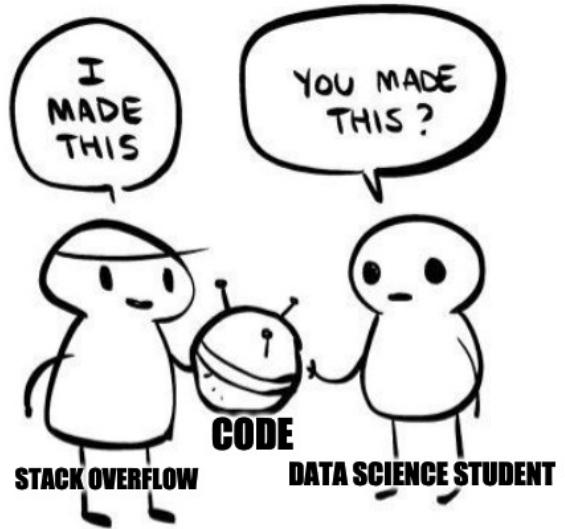
Training on 3's again



# Blending Latent Variables

Many times...





# Outline

---

Motivation for Variational Autoencoders (VAE)

Mechanics of VAE

Separability of VAE

**The math behind everything**

Generative models



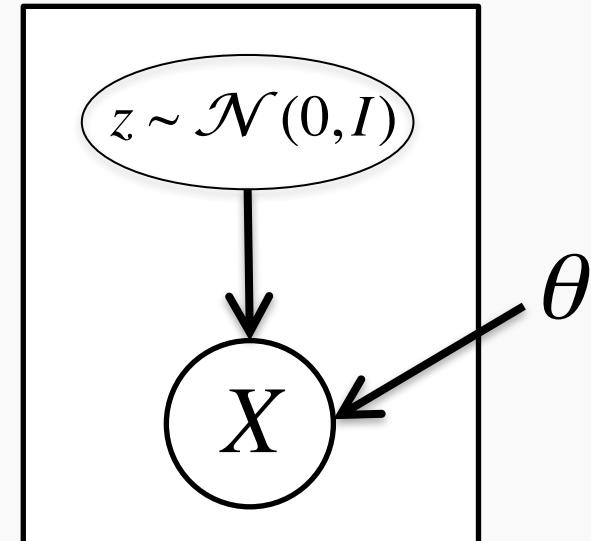
# VAE Likelihood

Neural network

$$p_{\theta}(x) = \int_z p_{\theta}(x|z)p_{\theta}(z)dz$$

Difficult to approximate in high  
dim through sampling

For most  $z$  values  $p(x|z)$  close to 0

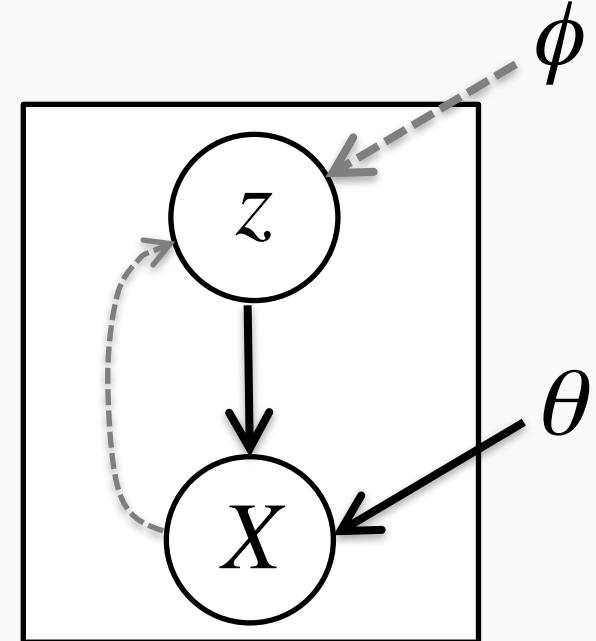


# VAE Likelihood

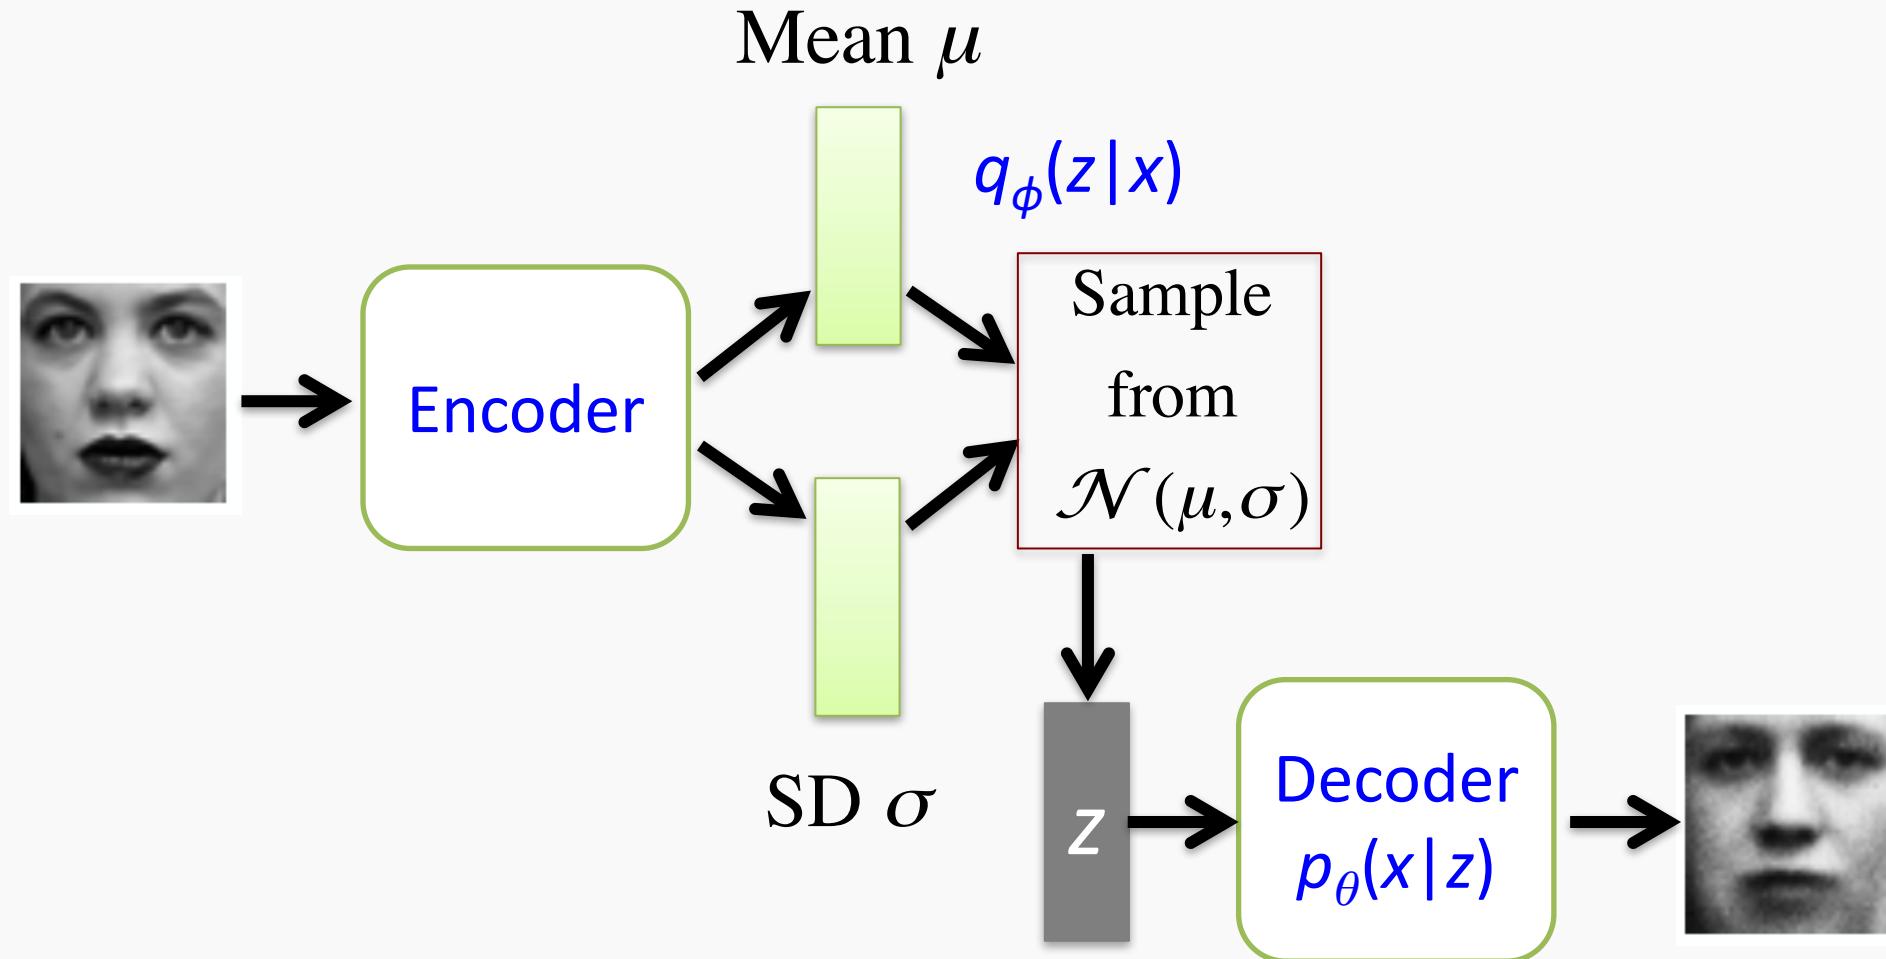
Another neural net

$$p_{\theta}(x) = \int_z p_{\theta}(x|z) q_{\phi}(z|x) dz$$

Proposal distribution:  
Likely to produce values of  $x$   
for which  $p(x|z)$  is non-zero



# VAE Architecture



# VAE Loss

Reconstruction Loss

$$-\mathbf{E}_{z \sim q_\phi(z|x)} \log(p_\theta(x|z))$$



# VAE Loss

Reconstruction Loss

Proposal distribution should  
resemble a Gaussian

$$-\mathbf{E}_{z \sim q_{\phi}(z|x)} \log(p_{\theta}(x|z)) + KL(q_{\phi}(z|x) \parallel p_{\theta}(z))$$



# VAE Loss

Reconstruction Loss

Proposal distribution should  
resemble a Gaussian

$$-\mathbf{E}_{z \sim q_{\phi}(z|x)} \log(p_{\theta}(x|z)) + KL(q_{\phi}(z|x) \parallel p_{\theta}(z))$$

$$\geq -\log p_{\theta}(x)$$

Variational upper bound  
on loss we care about!



# Training VAE

---

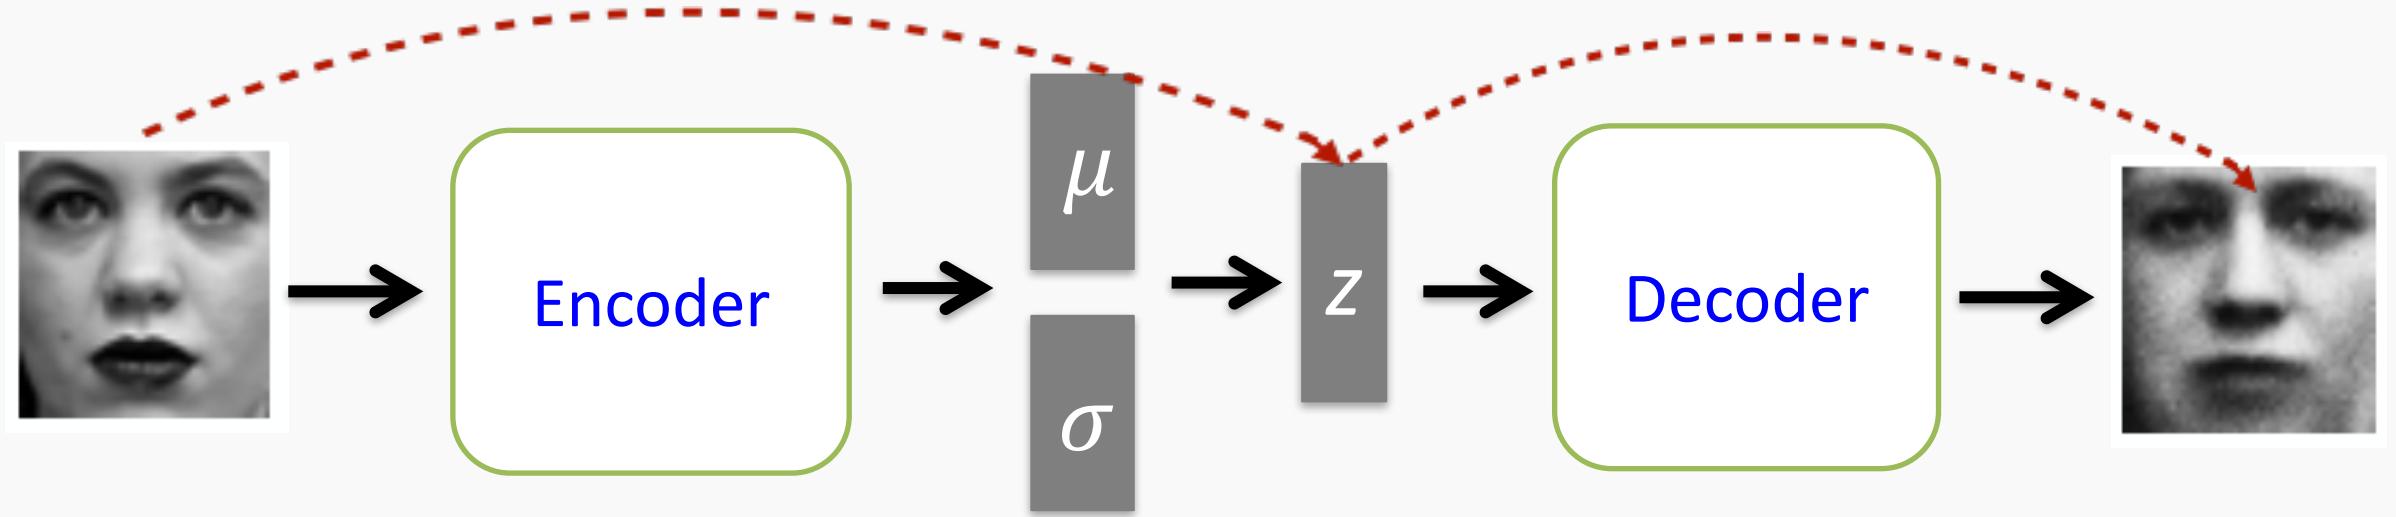
Apply stochastic gradient descent

Sampling step not differentiable

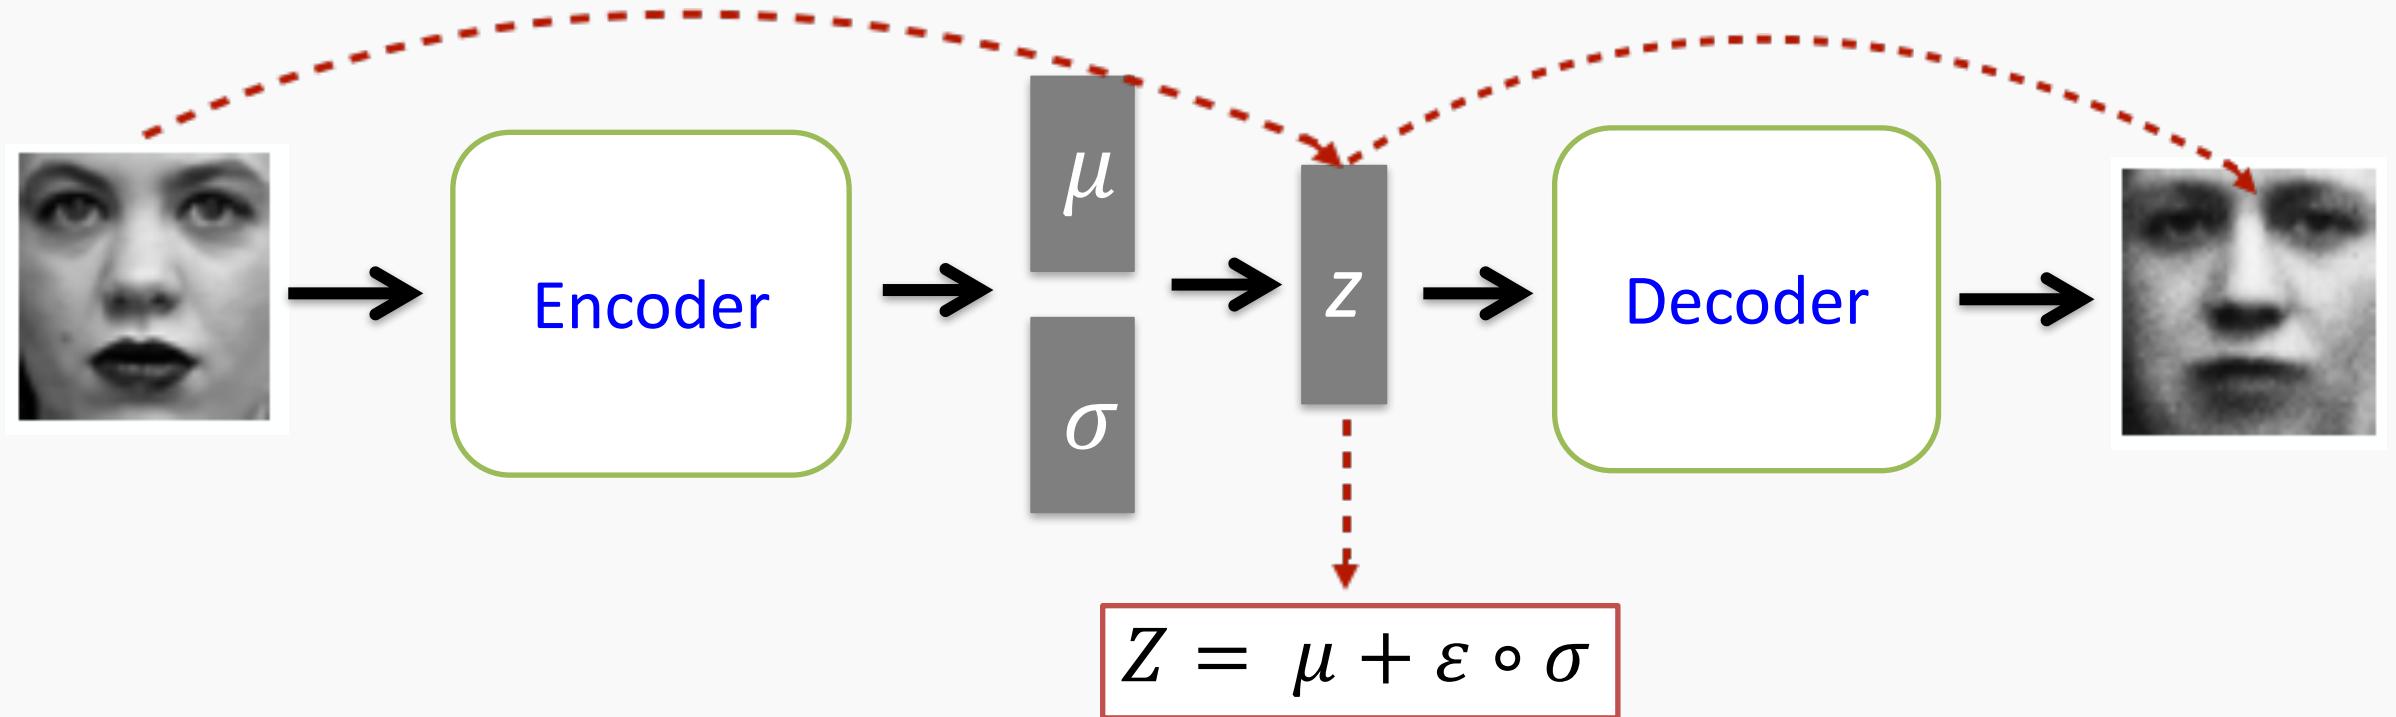
Use a re-parameterization trick

- Move sampling to input layer, so that the sampling step is independent of the model

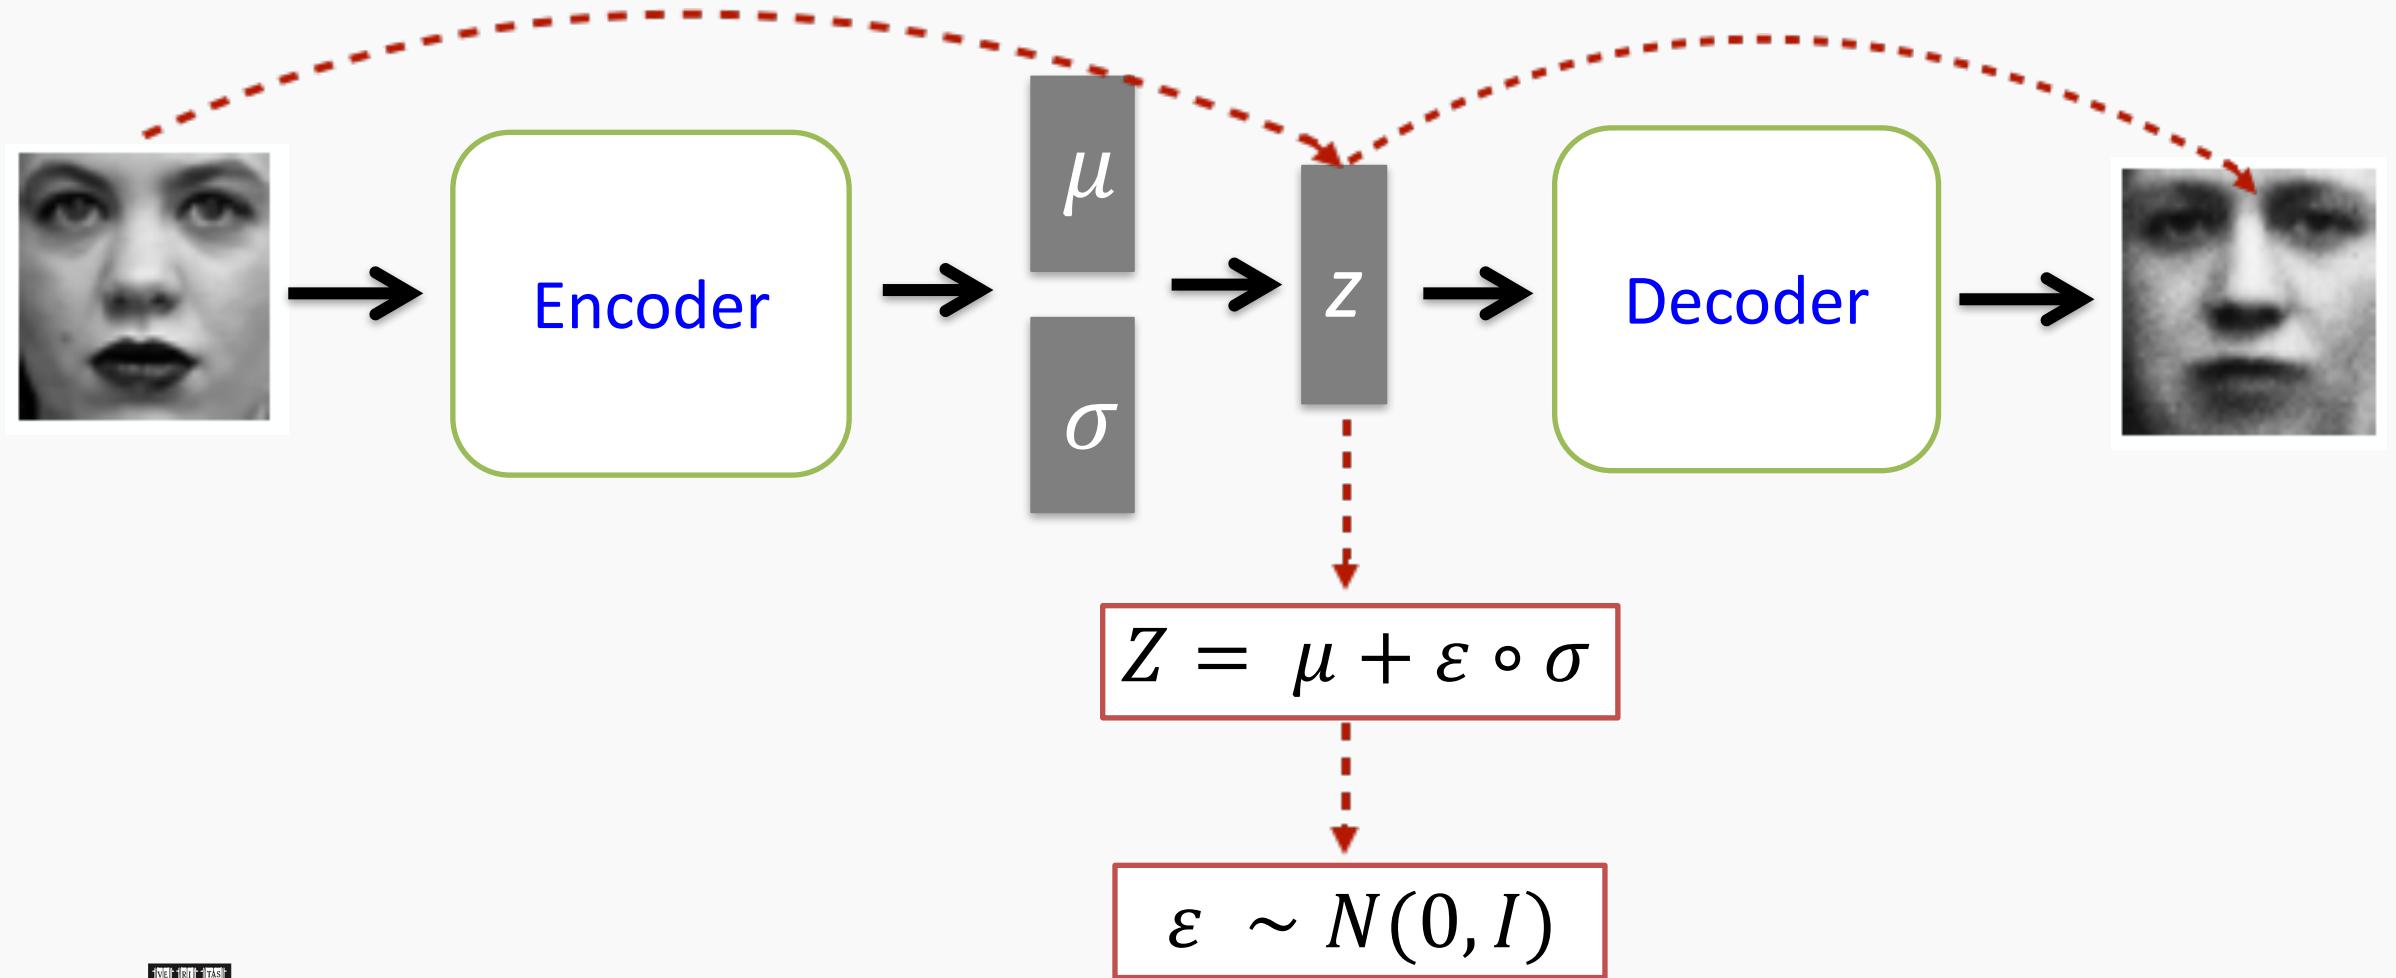
# Reparametrization Trick



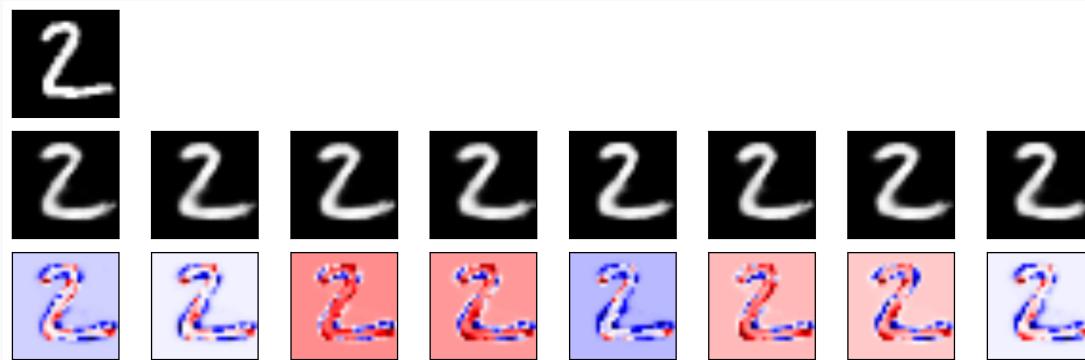
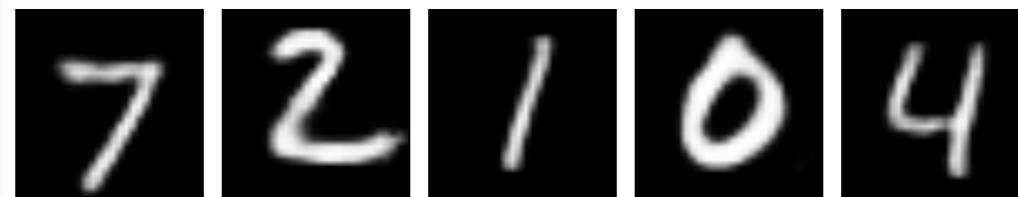
# Reparametrization Trick



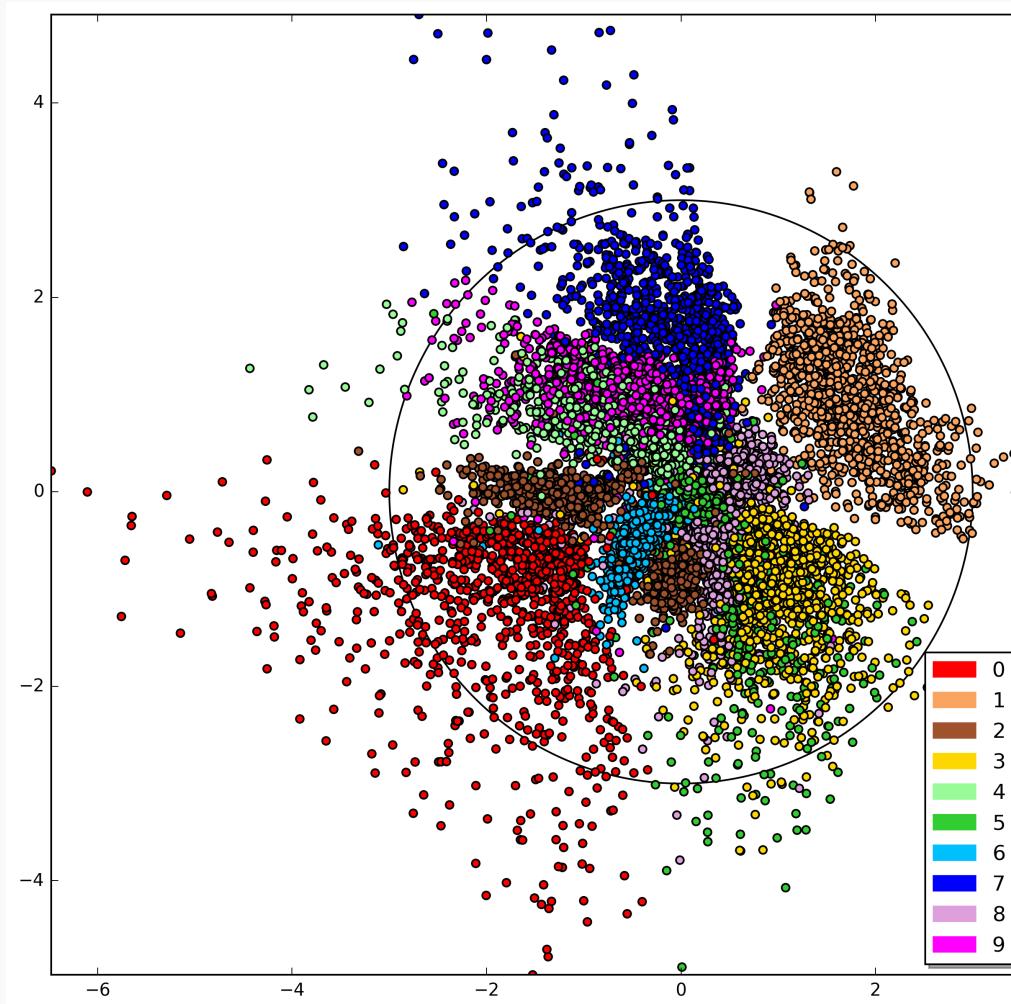
# Reparametrization Trick



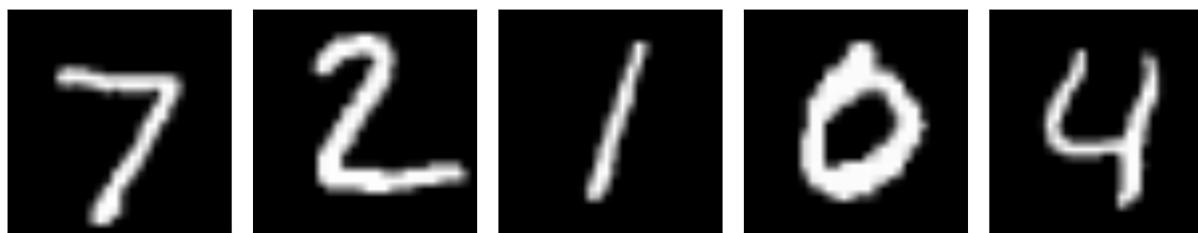
# Training VAE



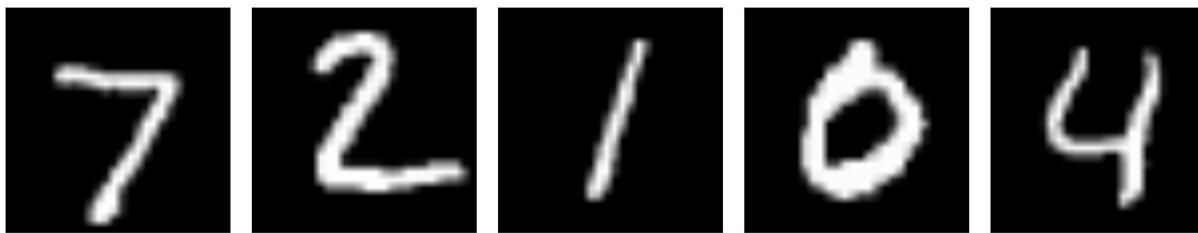
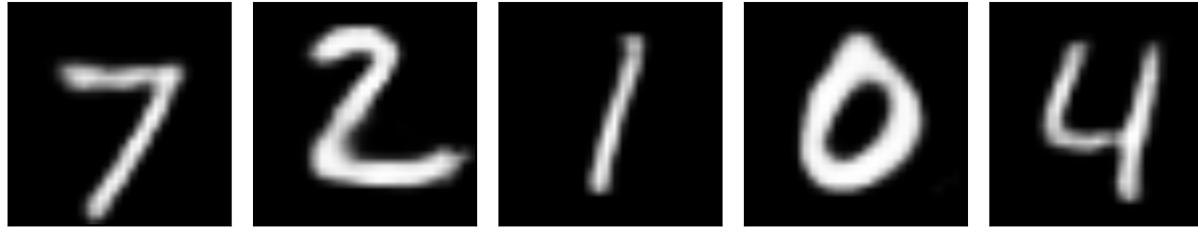
# Parameter space VAE



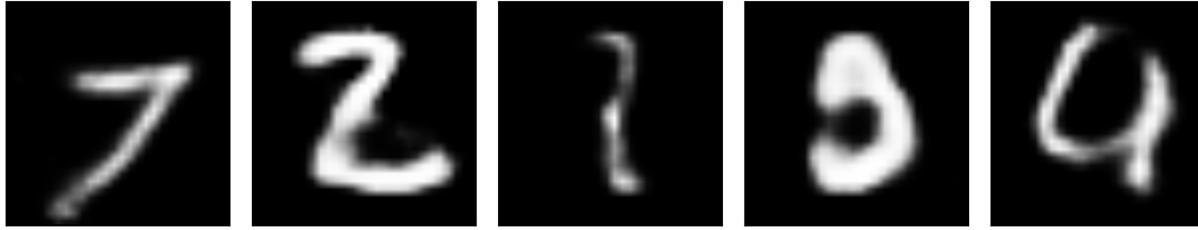
# Training VAE



10% error



30% error



# Parameter space VAE

