

Deep Learning

A course about theory & practice



Generative Networks

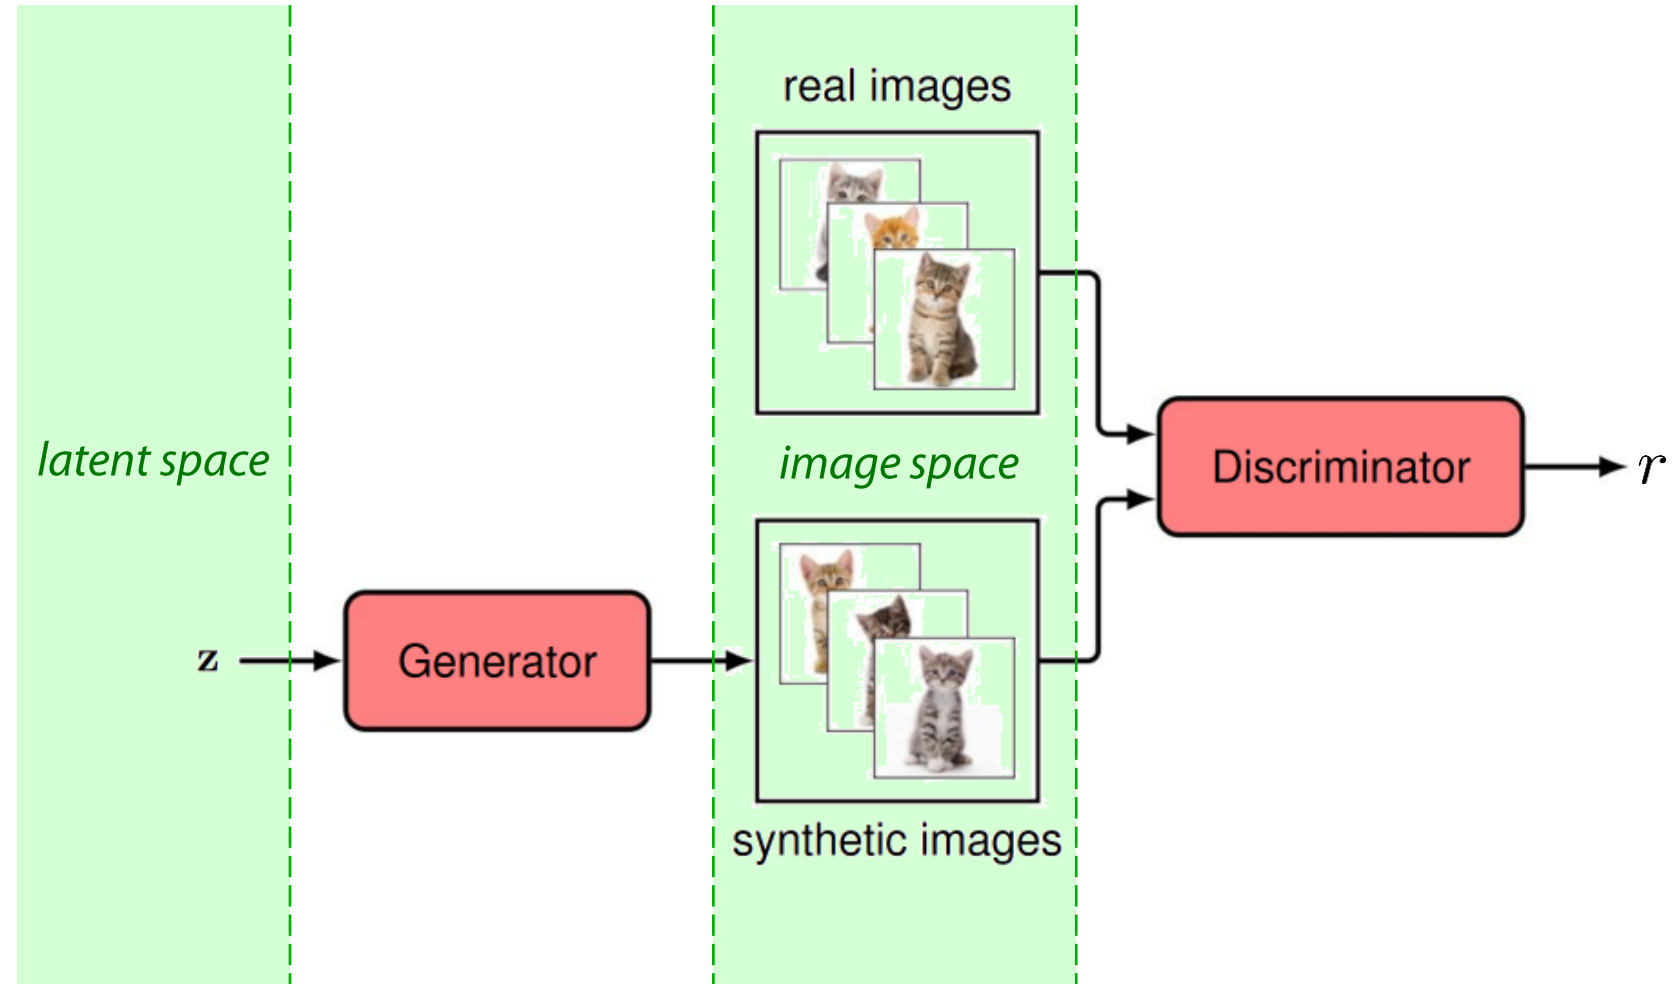
Marco Piastra

Generative Adversarial Networks

Basic idea

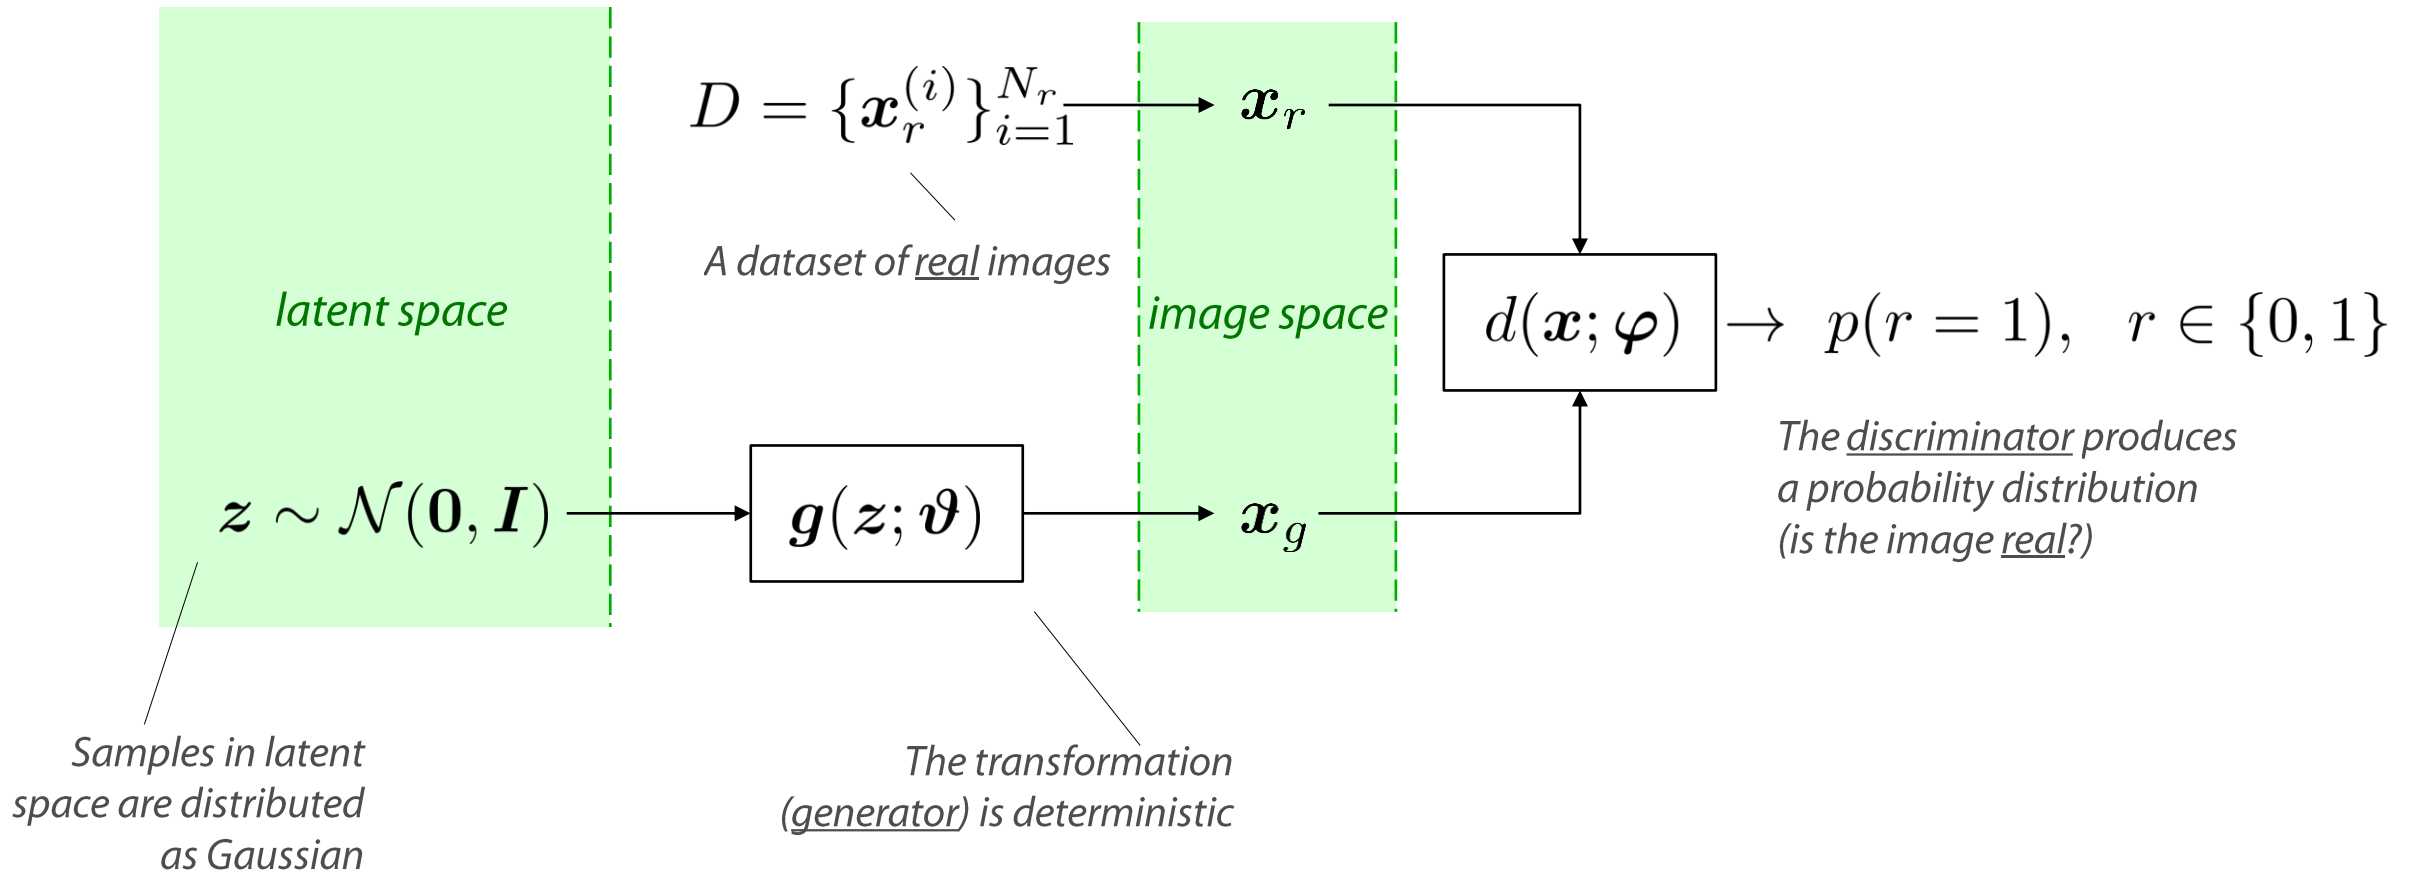
Objective:
creating a non-linear transformation from a latent space to a data space

Method:
training together a generator and a discriminator using a real dataset




[Image from <https://www.bishopbook.com/>]

Generative Adversarial Network (GAN)



Generative Adversarial Network (GAN)



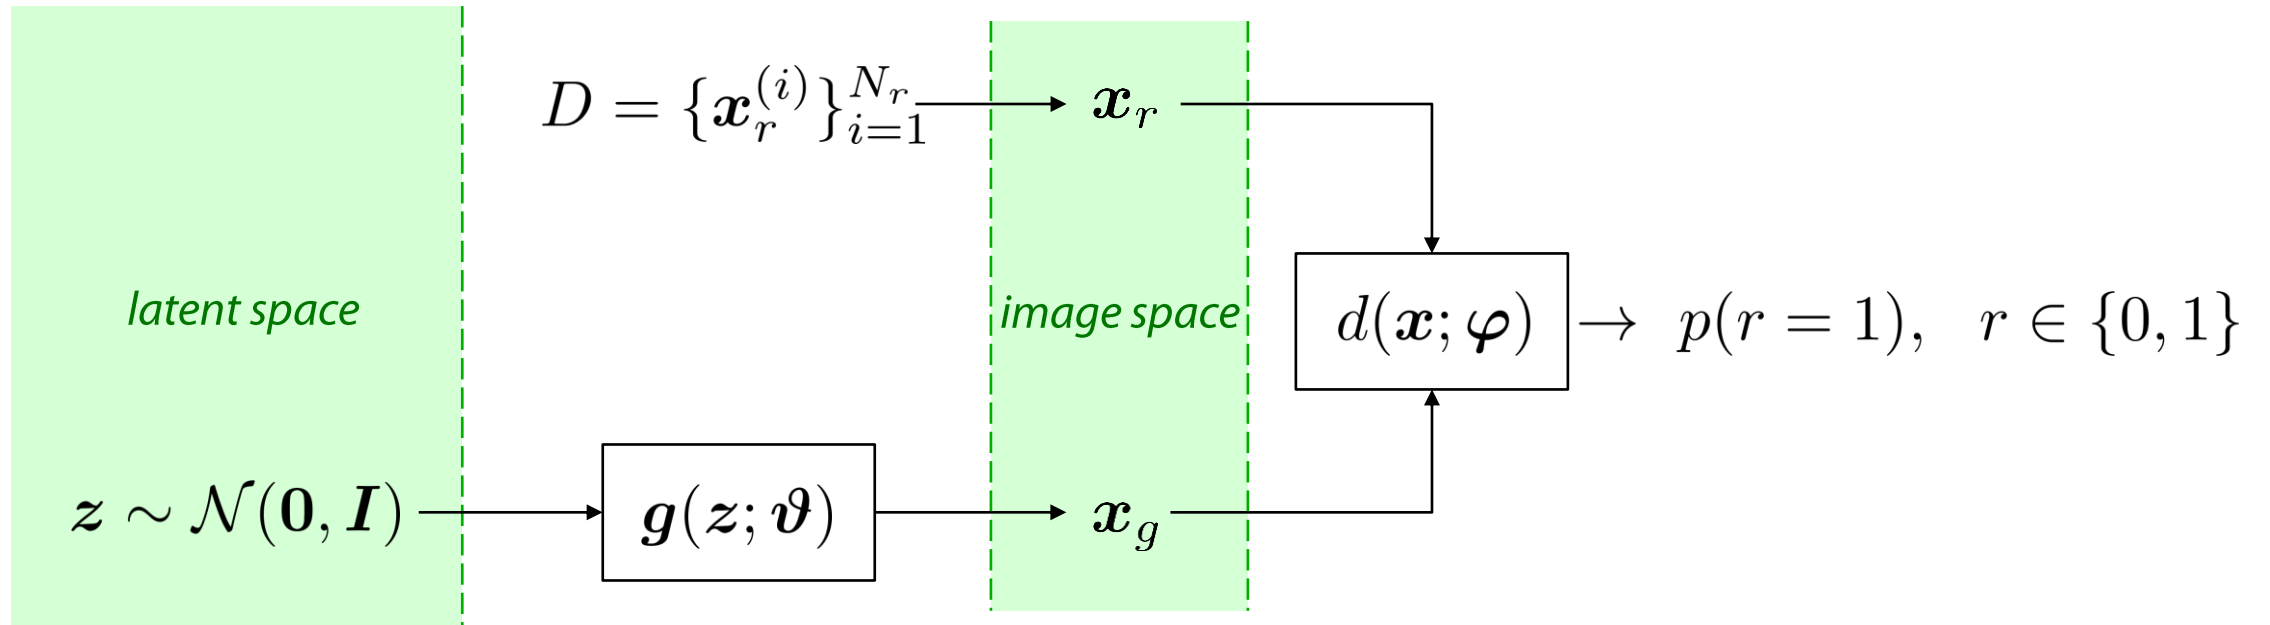
Loss function

$$L(\vartheta, \varphi) := -\frac{1}{N_r} \sum_{i \in \mathcal{R}} \ln(d(x_r^{(i)}; \varphi)) - \frac{1}{N_g} \sum_{j \in \mathcal{G}} \ln(1 - d(g(z^{(j)}; \vartheta); \varphi))$$

Cross-entropy
(d should recognize real images)

Cross-entropy
(d should recognize 'false' images)

Generative Adversarial Network (GAN)



Loss function

$$L(\vartheta, \varphi) := -\frac{1}{N_r} \sum_{i \in \mathcal{R}} \ln(d(x_r^{(i)}; \varphi)) - \frac{1}{N_g} \sum_{j \in \mathcal{G}} \ln(1 - d(g(z^{(j)}; \vartheta); \varphi))$$

Gradients

$$\Delta \varphi = -\eta \frac{\partial}{\partial \varphi} L(\vartheta, \varphi)$$

Make the discriminator smarter


$$\Delta \vartheta = +\eta \frac{\partial}{\partial \vartheta} L(\vartheta, \varphi)$$

Make the generator smarter: the discriminator should be fooled

Generative Adversarial Network (GAN)

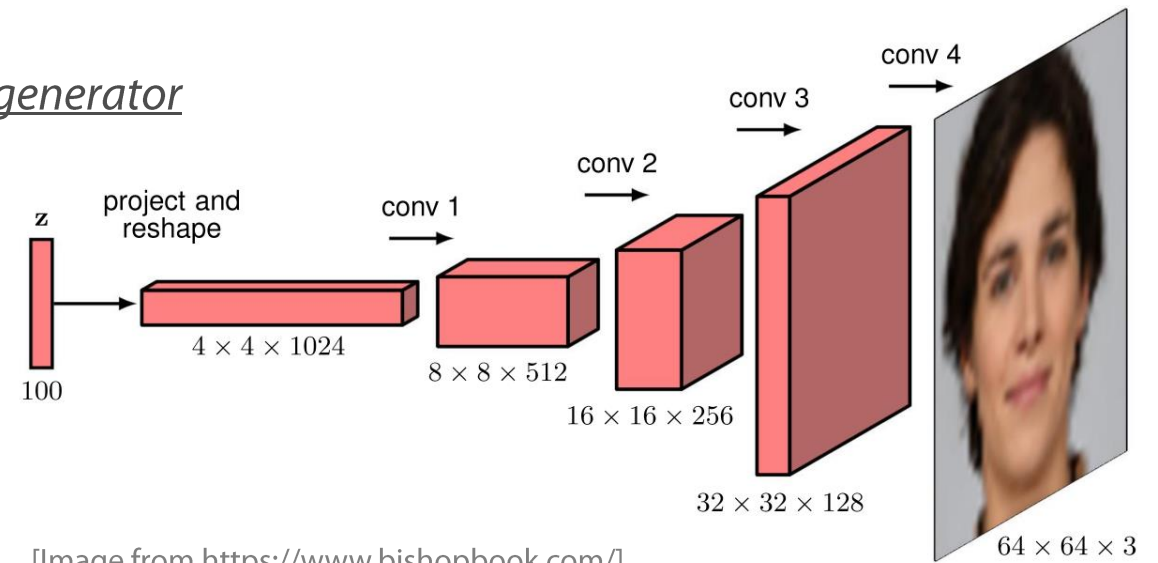
After training

The generator can be used to transform random samples in latent space into realistic data items



ImageGAN

Typically, a (de)convolutional network is used for the generator




[Image from <https://www.bishopbook.com/>]

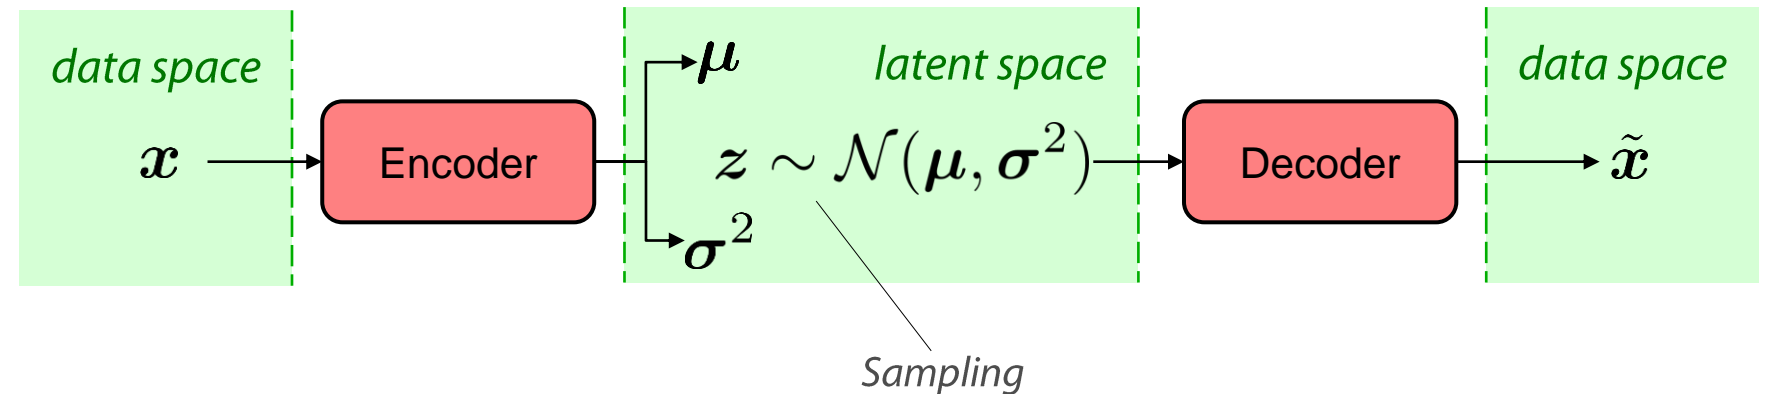
Variational Auto-Encoders

Basic idea

Auto-Encoder:
from data space
into latent space
then back





**Variational
Auto-Encoder:**
use a Gaussian spread
function to organize
the latent space





Basic idea

Auto-Encoder:
the correspondence
between data space
and latent space
is one to one





**Variational
Auto-Encoder:**
the correspondence
is one to many



Smooth generation

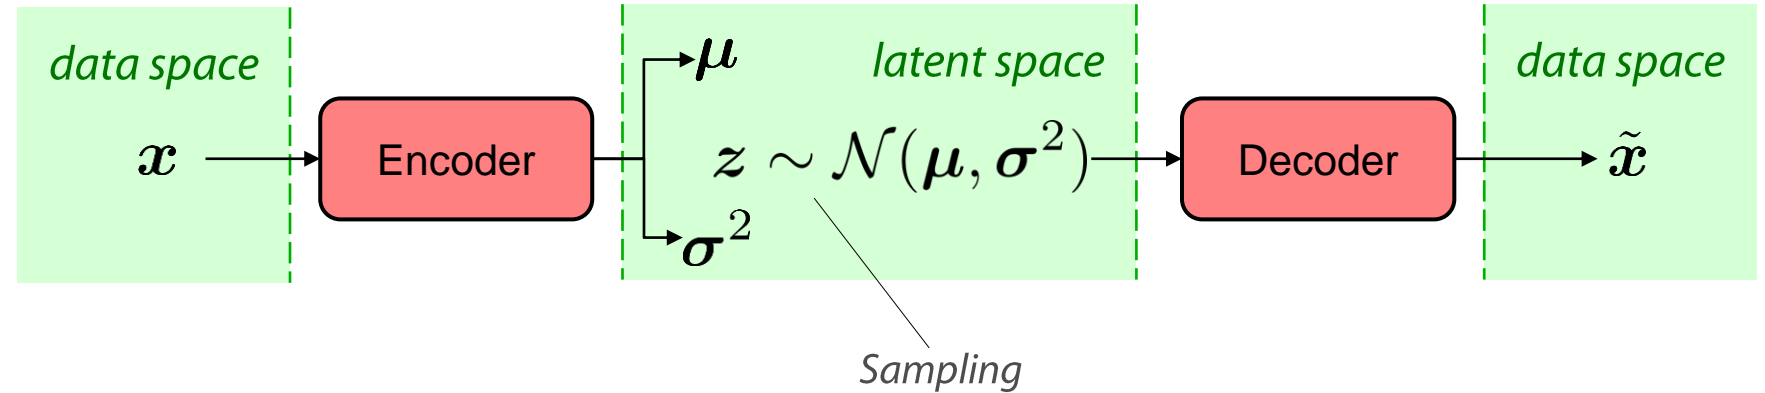
Variational Auto-Encoder:

after training, any convex combination of two points in latent space will generate a data item that changes smoothly from one extreme to the other



A problem: which loss function?

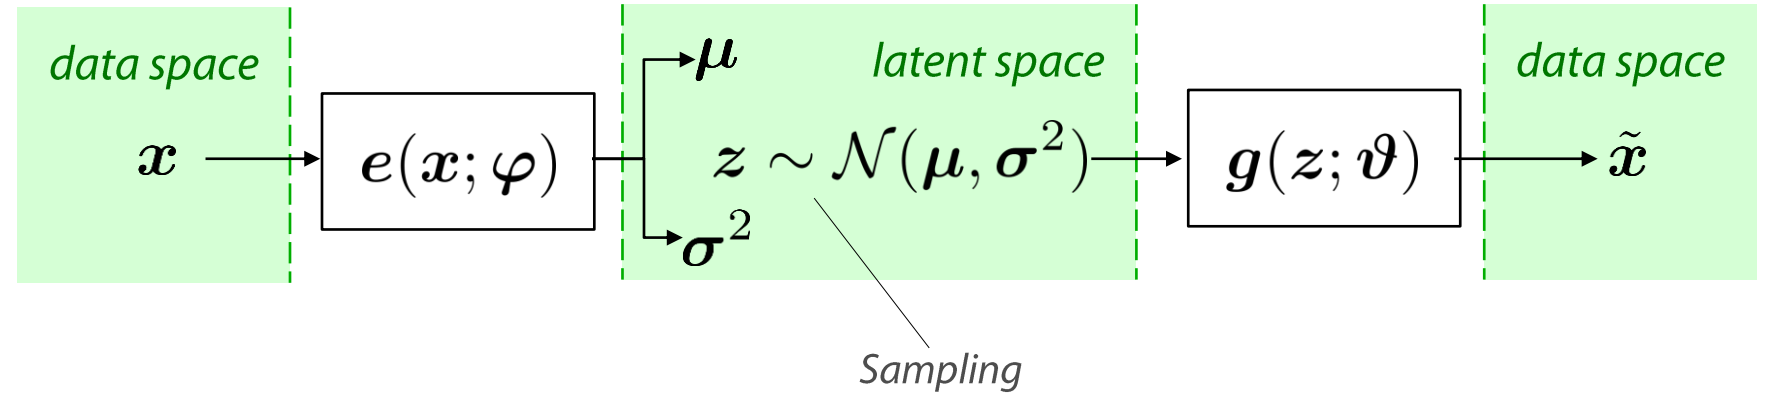
Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space



This is what we want to train from a real dataset $D = \{\mathbf{x}_r^{(i)}\}_{i=1}^{N_r}$

A problem: which loss function?

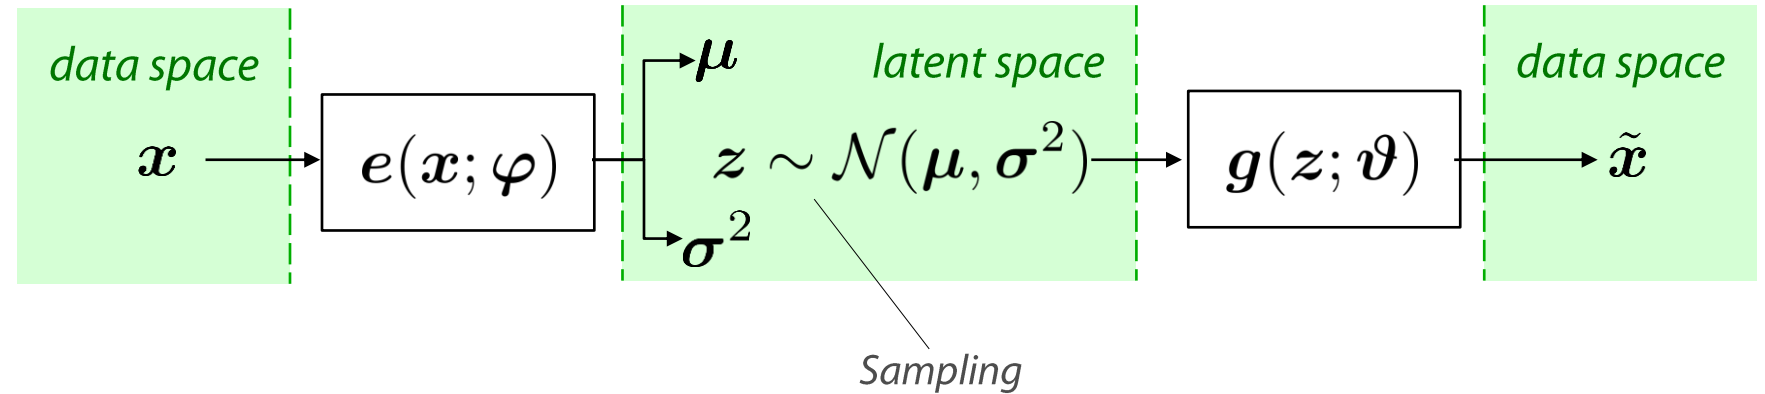
Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space



This is what we want to train from a real dataset $D = \{\mathbf{x}_r^{(i)}\}_{i=1}^{N_r}$

A problem: which loss function?

Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space



This is what we want to train from a real dataset $D = \{\mathbf{x}_r^{(i)}\}_{i=1}^{N_r}$

$$L(\mathbf{x}; \vartheta, \varphi) := \text{KL}(q(\mathbf{z} | \mathbf{x}, \varphi) \| p(\mathbf{z})) + \frac{1}{2} \frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}{c}$$


Kullback-Leibler divergence

This is similar to the loss of a standard autoencoder

This is an hyperparameter (see later)

A problem: which loss function?

Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space



This is what we want to train from a real dataset $D = \{\mathbf{x}_r^{(i)}\}_{i=1}^{N_r}$

$$L(\mathbf{x}; \vartheta, \varphi) := \text{KL}(q(\mathbf{z} | \mathbf{x}, \varphi) \| p(\mathbf{z})) + \frac{1}{2} \frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}{c}$$

Design choices


$$q(\mathbf{z} | \mathbf{x}, \varphi) := \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}; \varphi), \boldsymbol{\sigma}^2(\mathbf{x}; \varphi) \mathbf{I})$$

$$p(\mathbf{z}) := \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Normalization constraint: a soft limit against overspreading latent values

A problem: which loss function?

Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space




In general $\text{KL}(q(\mathbf{z}) \parallel p(\mathbf{z})) := \int q(\mathbf{z}) \ln \frac{q(\mathbf{z})}{p(\mathbf{z})} d\mathbf{z}$ ——— Kullback-Leibler divergence;
always positive, zero when the two distributions are identical

Under the conditions adopted

$$\text{KL}(q(\mathbf{z} \mid \mathbf{x}, \varphi) \parallel p(\mathbf{z})) = -\frac{1}{2} \sum_{j=1}^{\dim(\mathbf{z})} (1 + \ln \sigma_j^2(\mathbf{x}; \varphi) - \mu_j^2(\mathbf{x}; \varphi) - \sigma_j^2(\mathbf{x}; \varphi))$$

A problem: which loss function?

Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space



With a bit more of mathematics (omitted :)) it can be shown that the second term in the loss function

$$\frac{1}{2} \frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}{c}$$

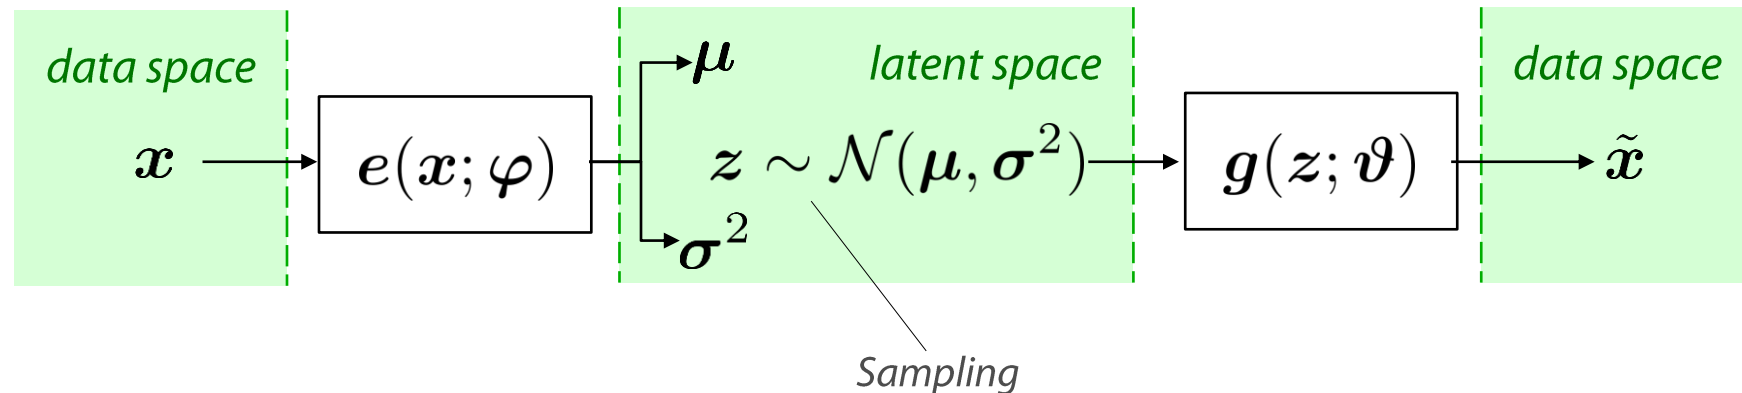
relates to an assumption of:

$$p(\tilde{\mathbf{x}}) := \mathcal{N}(\mathbf{x}, c\mathbf{I}), \quad c > 0$$

Design choice (points to c)
hyperparameter (points to \mathbf{I})
(hyper) spherical normal (points to \mathcal{N})

Reparametrization Trick

Variational Auto-Encoder:
use a Gaussian spread function to organize the latent space



$$L(\mathbf{x}; \vartheta, \varphi) := -\frac{1}{2} \sum_{j=1}^{\dim(\mathbf{z})} (1 + \ln \sigma_j^2(\mathbf{x}; \varphi) - \mu_j^2(\mathbf{x}; \varphi) - \sigma_j^2(\mathbf{x}; \varphi)) + \frac{1}{2} \frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}{c}$$


$$\Delta \varphi = -\eta \frac{\partial}{\partial \varphi} L(\mathbf{x}; \vartheta, \varphi)$$

$$\Delta \vartheta = -\eta \frac{\partial}{\partial \vartheta} L(\mathbf{x}; \vartheta, \varphi)$$

Reparametrization Trick

\tilde{x} depends on both ϑ and φ via z
 yet, when z is **sampled**, the derivative in φ is blocked

Variational Auto-Encoder:
 use a Gaussian spread function to organize the latent space



$$L(\mathbf{x}; \vartheta, \varphi) := -\frac{1}{2} \sum_{j=1}^{\dim(\mathbf{z})} (1 + \ln \sigma_j^2(\mathbf{x}; \varphi) - \mu_j^2(\mathbf{x}; \varphi) - \sigma_j^2(\mathbf{x}; \varphi)) + \frac{1}{2} \frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}{c}$$

The trick is assuming:

$$\mathbf{z} = \boldsymbol{\mu}(\mathbf{x}; \varphi) + \varepsilon \boldsymbol{\sigma}^2(\mathbf{x}; \varphi)$$

where:

$$\varepsilon \sim \mathcal{N}(0, 1)$$

is a parameter, therefore it is constant to the derivative.

In plain words, during training and per each data item $\mathbf{x}^{(i)}$ the system draws one random value ε and computes the derivatives

Links

<https://johfischer.com/2022/09/18/denoising-score-matching/>

<https://towardsdatascience.com/understanding-variational-autoencoders-vaes-f70510919f73>

https://en.wikipedia.org/wiki/Variational_autoencoder


<https://mbernste.github.io/posts/vae/>

Diffusion Models

Basic idea

Forward Diffusion


Assume that images are corrupted by Gaussian noise with known parameters



The idea behind

Denoising Diffusion Probabilistic Models

is learning how to reverse the process



Starting from the end: training algorithm

Forward Diffusion

Assume that images
are corrupted by
Gaussian noise
with known parameters

The idea behind

Denoising Diffusion Probabilistic Models
is learning how to reverse the process

Algorithm 20.1: Training a denoising diffusion probabilistic model

```
Input: Training data  $\mathcal{D} = \{\mathbf{x}_n\}$   
Noise schedule  $\{\beta_1, \dots, \beta_T\}$   
Output: Network parameters  $\mathbf{w}$   


---

for  $t \in \{1, \dots, T\}$  do  
|  $\alpha_t \leftarrow \prod_{\tau=1}^t (1 - \beta_\tau)$  // Calculate alphas from betas  
end for  
repeat  
|  $\mathbf{x} \sim \mathcal{D}$  // Sample a data point  
|  $t \sim \{1, \dots, T\}$  // Sample a point along the Markov chain  
|  $\epsilon \sim \mathcal{N}(\epsilon | \mathbf{0}, \mathbf{I})$  // Sample a noise vector  
|  $\mathbf{z}_t \leftarrow \sqrt{\alpha_t} \mathbf{x} + \sqrt{1 - \alpha_t} \epsilon$  // Evaluate noisy latent variable  
|  $\mathcal{L}(\mathbf{w}) \leftarrow \|\mathbf{g}(\mathbf{z}_t, \vartheta, t) - \epsilon\|^2$  // Compute loss term  
| Take optimizer step  
until converged  
return  $\mathbf{w}$ 
```


Neural network
with suitable architecture

[Image from <https://www.bishopbook.com/>]

Forward diffusion

Forward Diffusion

Assume that images are corrupted by Gaussian noise with known parameters



$$\mathbf{x}_t \sim \mathcal{N} \left(\sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I} \right)$$

$$\beta_t \in (0, 1), \forall t$$

$$\beta_1 < \beta_2 < \dots < \beta_T$$

Hyperparameters

Could be rewritten as


$$\mathbf{x}_t = \sqrt{1 - \beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \boldsymbol{\varepsilon}$$

$$\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Forward diffusion

Forward Diffusion

Assume that images are corrupted by Gaussian noise with known parameters



At any forward step t , the diffusion sequence can be compacted as

$$x_t \sim \mathcal{N}(\sqrt{\alpha_t} x_0, (1 - \alpha_t) \mathbf{I})$$

where:


$$\alpha_t = \prod_{\tau=1}^t (1 - \beta_\tau)$$



Going backward: denoising

Backward Denoising

A neural network is at the core of the backward process



We assume that:

$$\mathbf{x}_{t-1} = \boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) + \sqrt{\beta_t} \boldsymbol{\varepsilon}$$


$$\boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) = \frac{1}{\sqrt{1 - \beta_t}} \left\{ \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \overbrace{g(\mathbf{x}_t, t; \boldsymbol{\vartheta})}^{\text{Neural Network}} \right\}$$

$$\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Going backward: denoising

Backward Denoising

A neural network is at the core of the backward process



We assume that:

$$\mathbf{x}_{t-1} = \boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) + \sqrt{\beta_t} \boldsymbol{\varepsilon}$$

$$\boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) = \frac{1}{\sqrt{1 - \beta_t}} \left\{ \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \overbrace{g(\mathbf{x}_t, t; \boldsymbol{\vartheta})}^{\text{Neural Network}} \right\}$$

$$\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

How can the neural network be trained?
(A suitable loss function is needed)

Going backward: denoising

$$\tilde{q}(\mathbf{x}_{t-1}|\mathbf{x}_t)$$

An approximation to $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$

We assume that:

$$\mathbf{x}_{t-1} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}), \beta_t \mathbf{I})$$

During training, \mathbf{x}_0 is known. Then we can sample $\boldsymbol{\varepsilon}_t$

$$\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \boldsymbol{\varepsilon}_t$$

Noise added at step t

Therefore, it can be proven that:

$$\mathbf{m}(\mathbf{x}_{t-1}) = \frac{1}{\sqrt{1 - \beta_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \boldsymbol{\varepsilon}_t \right)$$

is the true mean of:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t)$$

Then the Kullback-Leibler divergence is:

$$\begin{aligned} \text{KL} (q(\mathbf{x}_{t-1}|\mathbf{x}_t) \parallel \tilde{q}(\mathbf{x}_{t-1}|\mathbf{x}_t)) \\ = \frac{1}{2\beta_t} \|\mathbf{m}(\mathbf{x}_{t-1}) - \boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta})\|^2 + \text{const} \end{aligned}$$

Going backward: denoising

$$\begin{aligned} & \text{KL} (q(\mathbf{x}_{t-1}|\mathbf{x}_t) \parallel \tilde{q}(\mathbf{x}_{t-1}|\mathbf{x}_t)) \\ &= \frac{1}{2\beta_t} \|\boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) - \mathbf{m}(\mathbf{x}_{t-1})\|^2 + \text{const} \end{aligned}$$

Therefore, given (see before):



$$\begin{aligned} \mathbf{m}(\mathbf{x}_{t-1}) &= \frac{1}{\sqrt{1-\beta_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \boldsymbol{\varepsilon}_t \right) \\ \boldsymbol{\mu}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) &= \frac{1}{\sqrt{1-\beta_t}} \left\{ \mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \mathbf{g}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) \right\} \end{aligned}$$

$$\text{KL} (\tilde{q}(\mathbf{x}_{t-1}|\mathbf{x}_t) \parallel q(\mathbf{x}_{t-1}|\mathbf{x}_t)) \propto \|\mathbf{g}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) - \boldsymbol{\varepsilon}_t\|^2$$


$$L(\boldsymbol{\vartheta}) := \|\mathbf{g}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) - \boldsymbol{\varepsilon}_t\|^2 \quad \text{To be minimized}$$

*Diffusion Models:
Why so many steps?*



Going forward: adding noise



Marginal distribution at step $t - 1$






Forward probability distribution (Gaussian)



$$\mathbf{x}_t = \sqrt{1 - \beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \varepsilon$$

Going backward: denoising

The backward probability distribution can be computed from forward and marginals using Bayes' theorem




This is what we want to learn


$$q(\mathbf{x}_{t-1} | \mathbf{x}_t) = \frac{q(\mathbf{x}_t | \mathbf{x}_{t-1}) q(\mathbf{x}_{t-1})}{q(\mathbf{x}_t)}$$

Bayes' Theorem

Going backward: denoising



At training time, \mathbf{x}_t is known
(dataset + forward diffusion)




$$q(\mathbf{x}_{t-1} | \mathbf{x}_t) = \frac{q(\mathbf{x}_t | \mathbf{x}_{t-1}) q(\mathbf{x}_{t-1})}{q(\mathbf{x}_t)}$$

Bayes' Theorem


[Image from <https://www.bishopbook.com/>]

Going backward: denoising



At training time, \mathbf{x}_t is known
(dataset + forward diffusion)

The reverse probability
is the blue curve






$$q(\mathbf{x}_{t-1} | \mathbf{x}_t) = \frac{q(\mathbf{x}_t | \mathbf{x}_{t-1})q(\mathbf{x}_{t-1})}{q(\mathbf{x}_t)}$$

Bayes' Theorem

[Image from <https://www.bishopbook.com/>]

Going backward: denoising

When β_t is large
 $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$
becomes very different
from a Gaussian, hence
unsuitable for training






$$q(\mathbf{x}_{t-1}|\mathbf{x}_t) = \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1})q(\mathbf{x}_{t-1})}{q(\mathbf{x}_t)}$$

Bayes' Theorem

[Image from <https://www.bishopbook.com/>]

Going backward: denoising

When β_t is small
 $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$
is approximately Gaussian



$$q(\mathbf{x}_{t-1}|\mathbf{x}_t) = \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1})q(\mathbf{x}_{t-1})}{q(\mathbf{x}_t)}$$

Bayes' Theorem

Links

<https://www.assemblyai.com/blog/diffusion-models-for-machine-learning-introduction/>

<https://lilianweng.github.io/posts/2021-07-11-diffusion-models/>

<https://www.superannotate.com/blog/diffusion-models>


<https://encord.com/blog/diffusion-models/>

Practical Implementation

Conditional U-Net as basic denoising block

Loss function: $L(\vartheta) := \|\mathbf{g}(\mathbf{x}_t, t; \vartheta) - \epsilon_t\|^2$

The network architecture for $\mathbf{g}(\mathbf{x}_t, t; \vartheta)$ is a U-Net



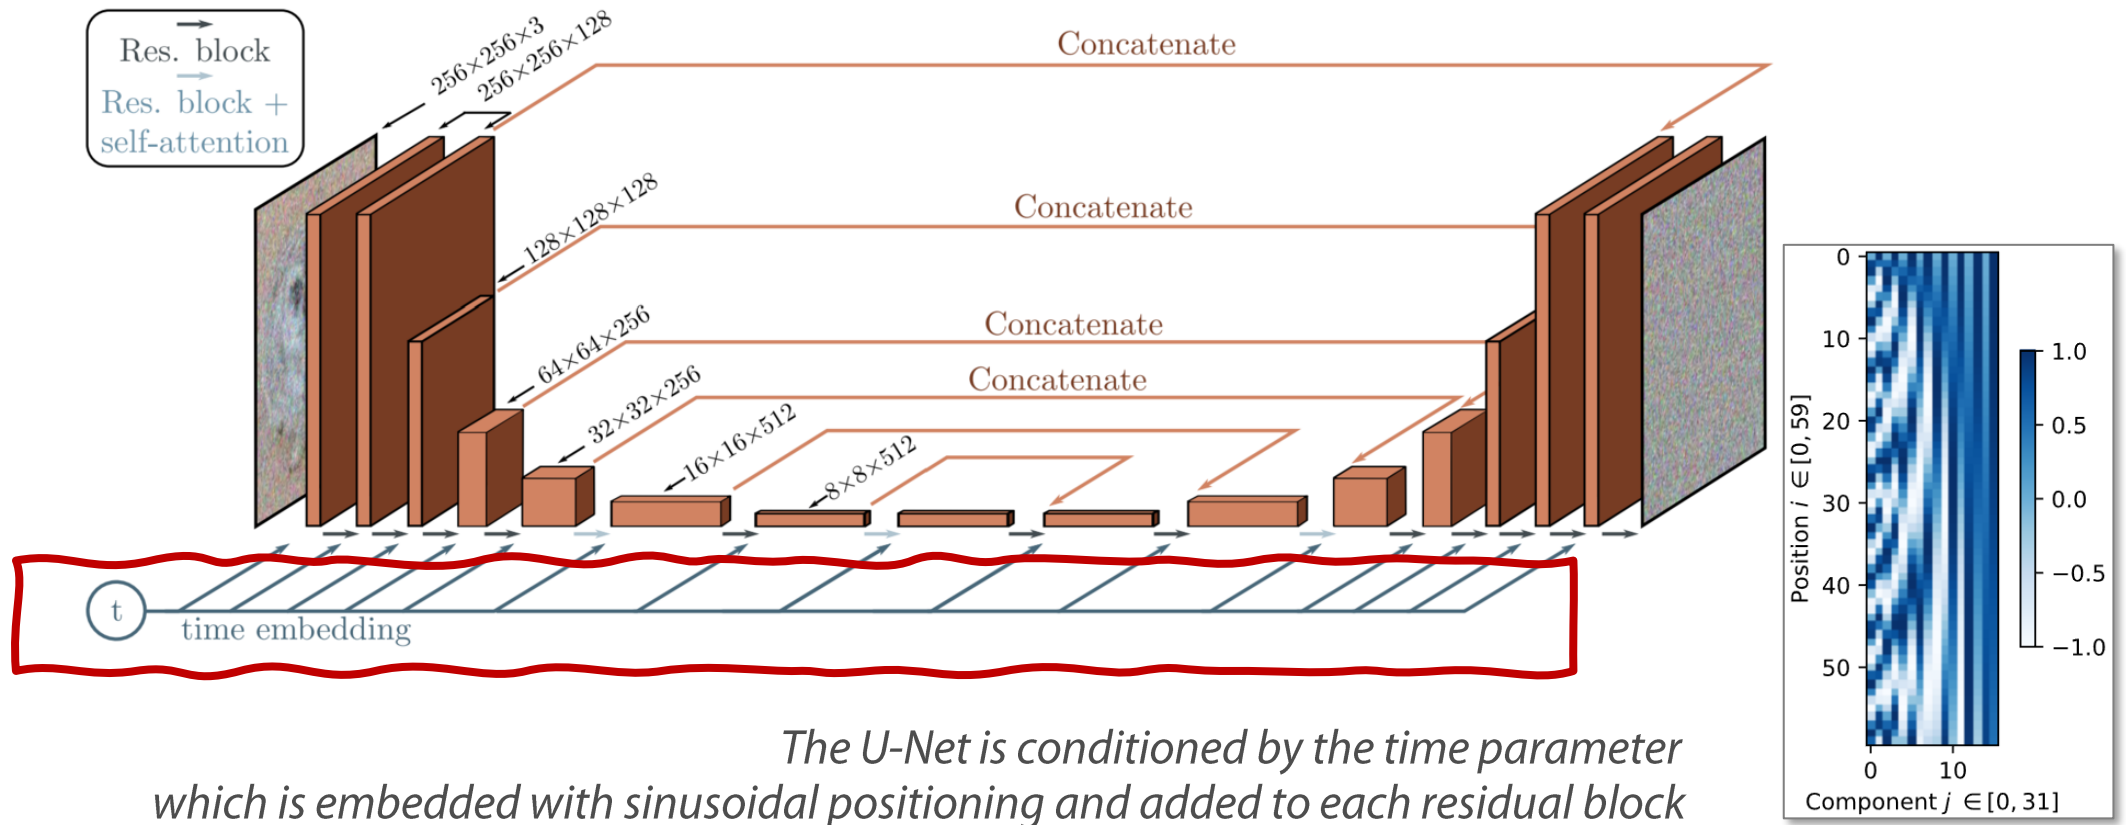
[Image from <https://learnopencv.com/denoising-diffusion-probabilistic-models/>]

[Ho, Jain & Abbeel, 2020 - <https://arxiv.org/pdf/2006.11239>]

Conditional U-Net as basic denoising block

Loss function: $L(\vartheta) := \|g(\mathbf{x}_t, t; \vartheta) - \epsilon_t\|^2$

The network architecture for $g(\mathbf{x}_t, t; \vartheta)$ is a U-Net



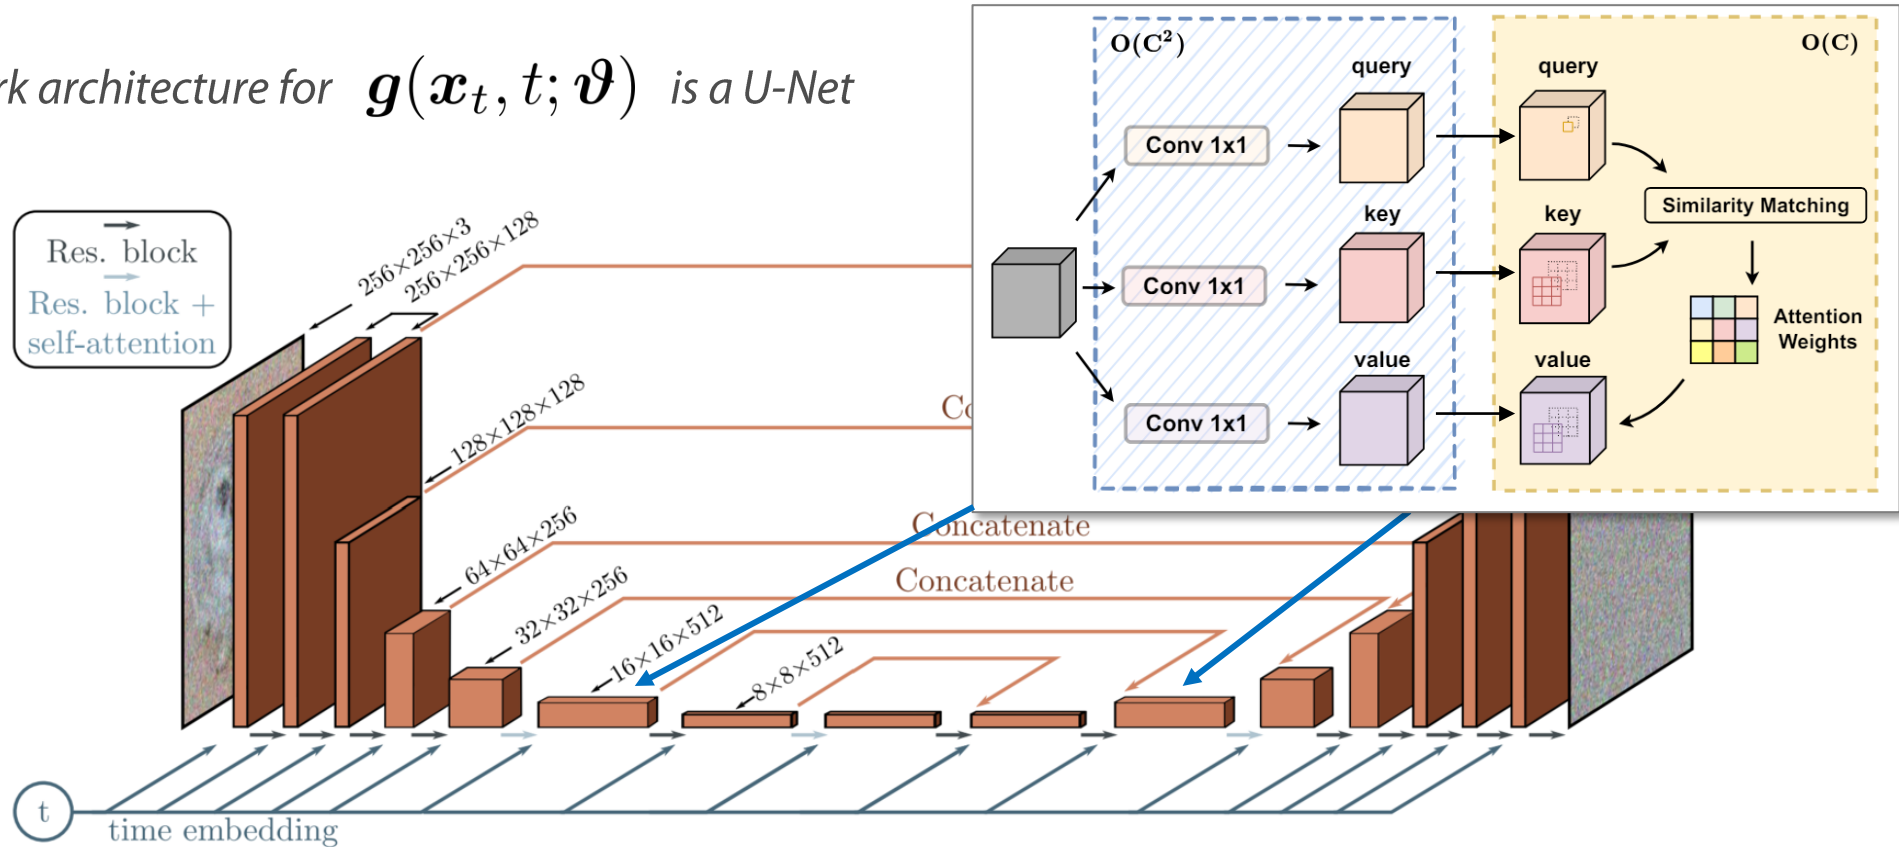
[Image from <https://learnopencv.com/denoising-diffusion-probabilistic-models/>]

[Ho, Jain & Abbeel, 2020 - <https://arxiv.org/pdf/2006.11239>]

Conditional U-Net as basic denoising block

$$\text{Loss function: } L(\boldsymbol{\vartheta}) := \|\mathbf{g}(\mathbf{x}_t, t; \boldsymbol{\vartheta}) - \boldsymbol{\varepsilon}_t\|^2$$

The network architecture for $\mathbf{g}(\mathbf{x}_t, t; \boldsymbol{\vartheta})$ is a U-Net



Self-Attention modules
are interspersed with convolutional blocks in the pipeline

[Ho, Jain & Abbeel, 2020 - <https://arxiv.org/pdf/2006.11239>]

Latent Diffusion Models


Forward Diffusion

It is relatively easy and inexpensive
(It can be performed in one step)



Backward Denoising


Must be performed in small steps
and is quite expensive,
in particular with high-resolution images



Latent Diffusion Models

Latent Diffusion Model

The intuitive idea is to perform diffusion in the *latent space*




[Rombach et al., 2022 - <https://arxiv.org/pdf/2112.10752>]

Latent Diffusion Models

Latent Diffusion Model

The intuitive idea is to perform diffusion in the *latent space*

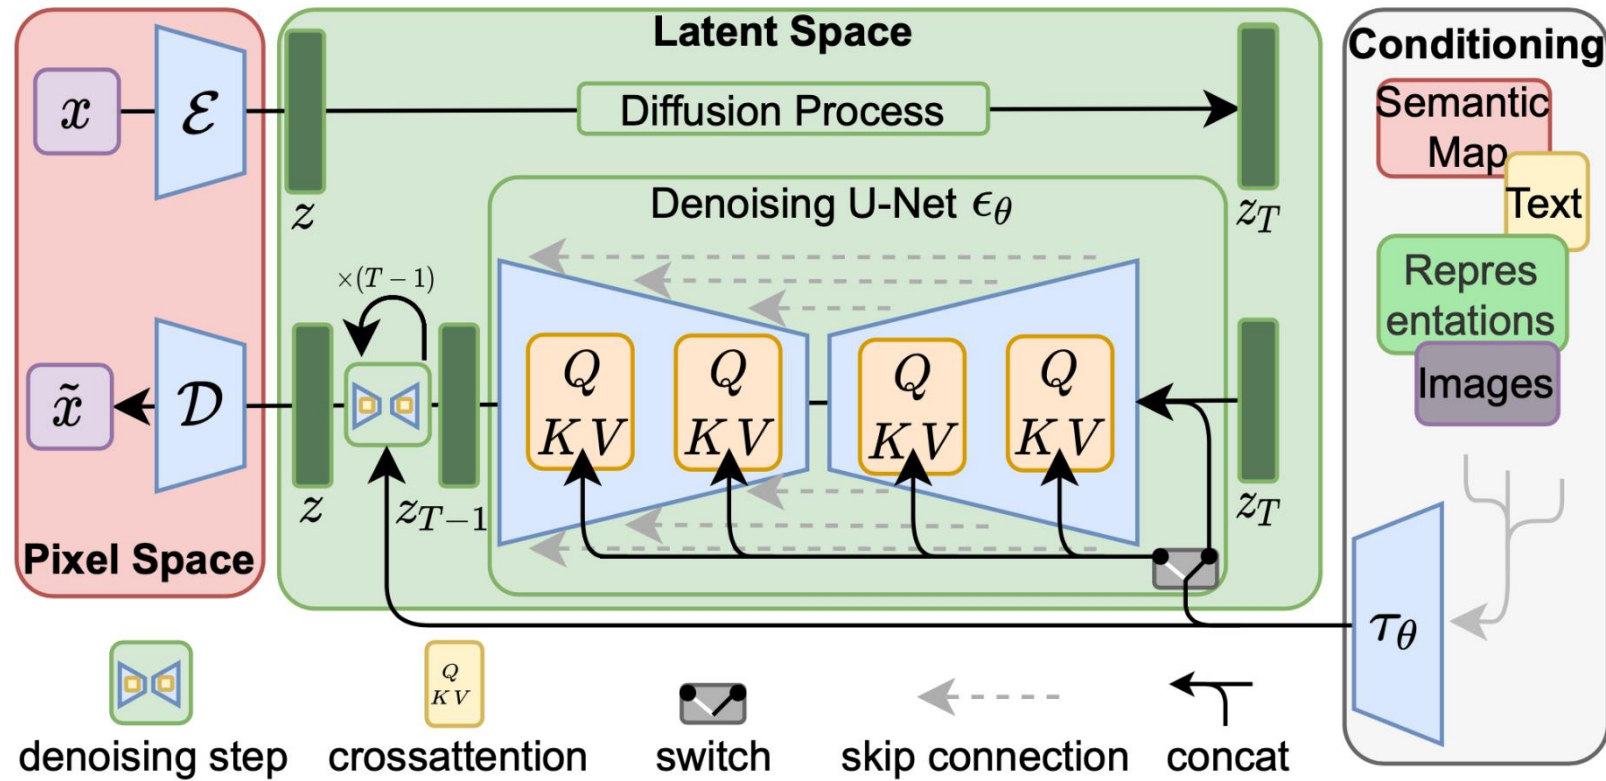


A pre-trained VAE is used to encode and decode high-resolution images into a suitable (reduced) latent format

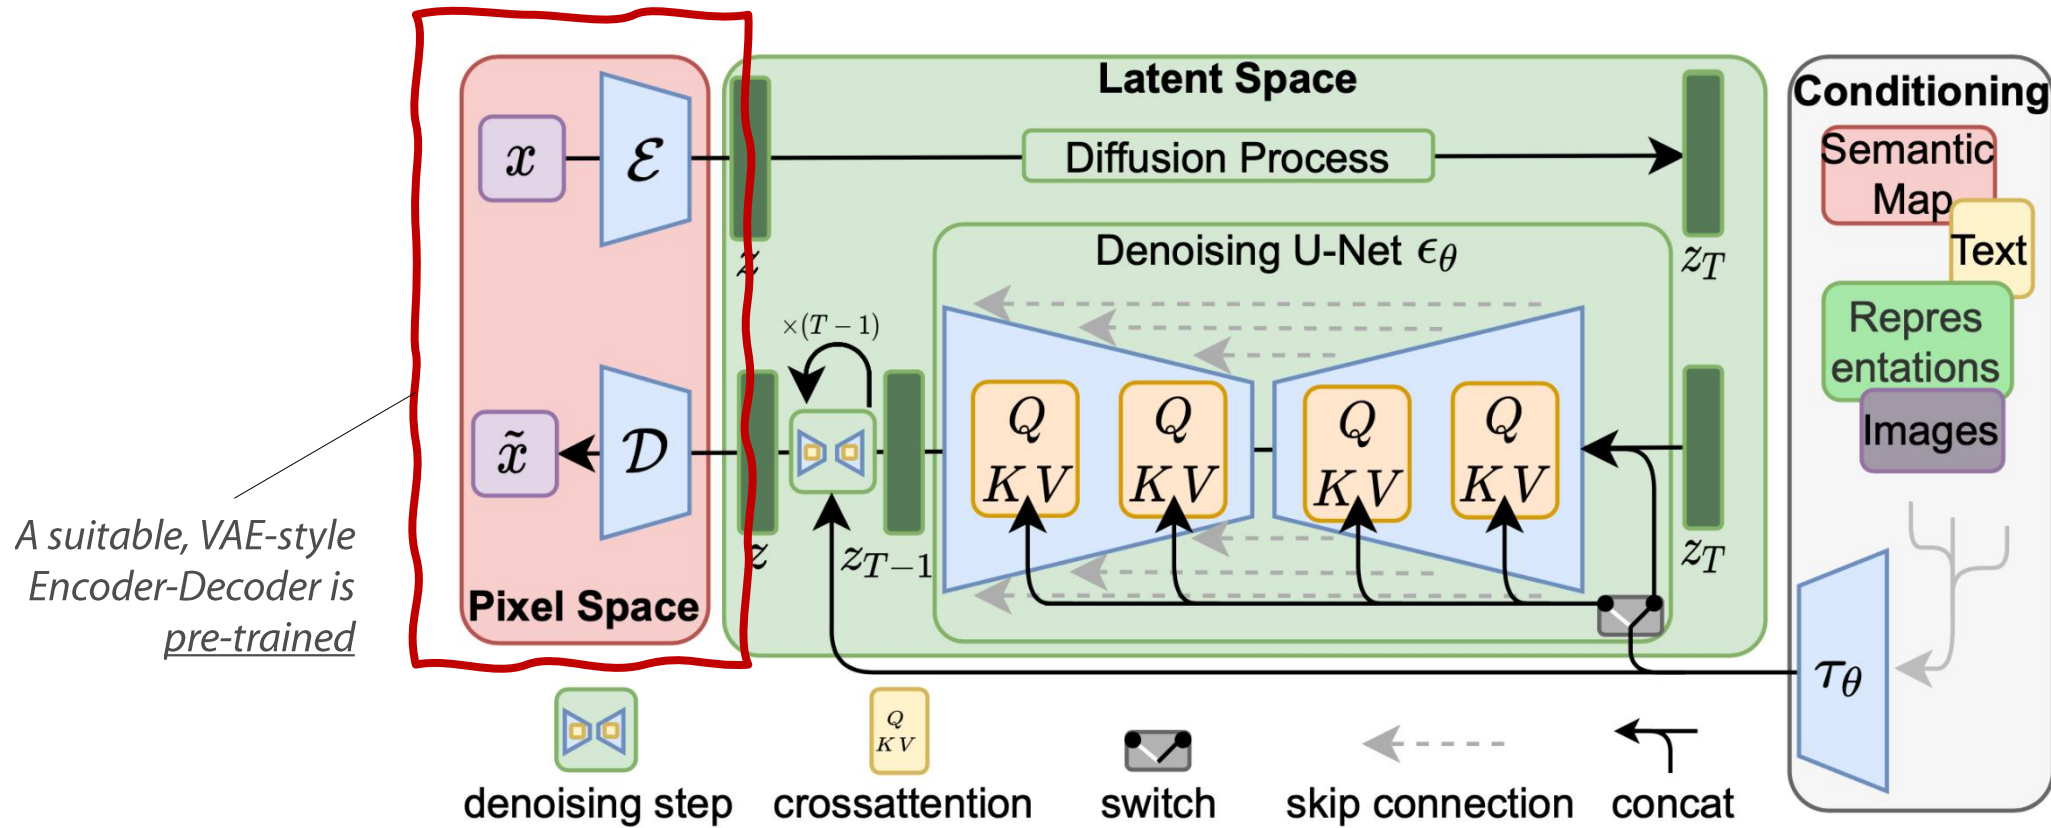
[Rombach et al., 2022 - <https://arxiv.org/pdf/2112.10752>]

Conditioning on Labels

Latent Diffusion Models with Conditioning

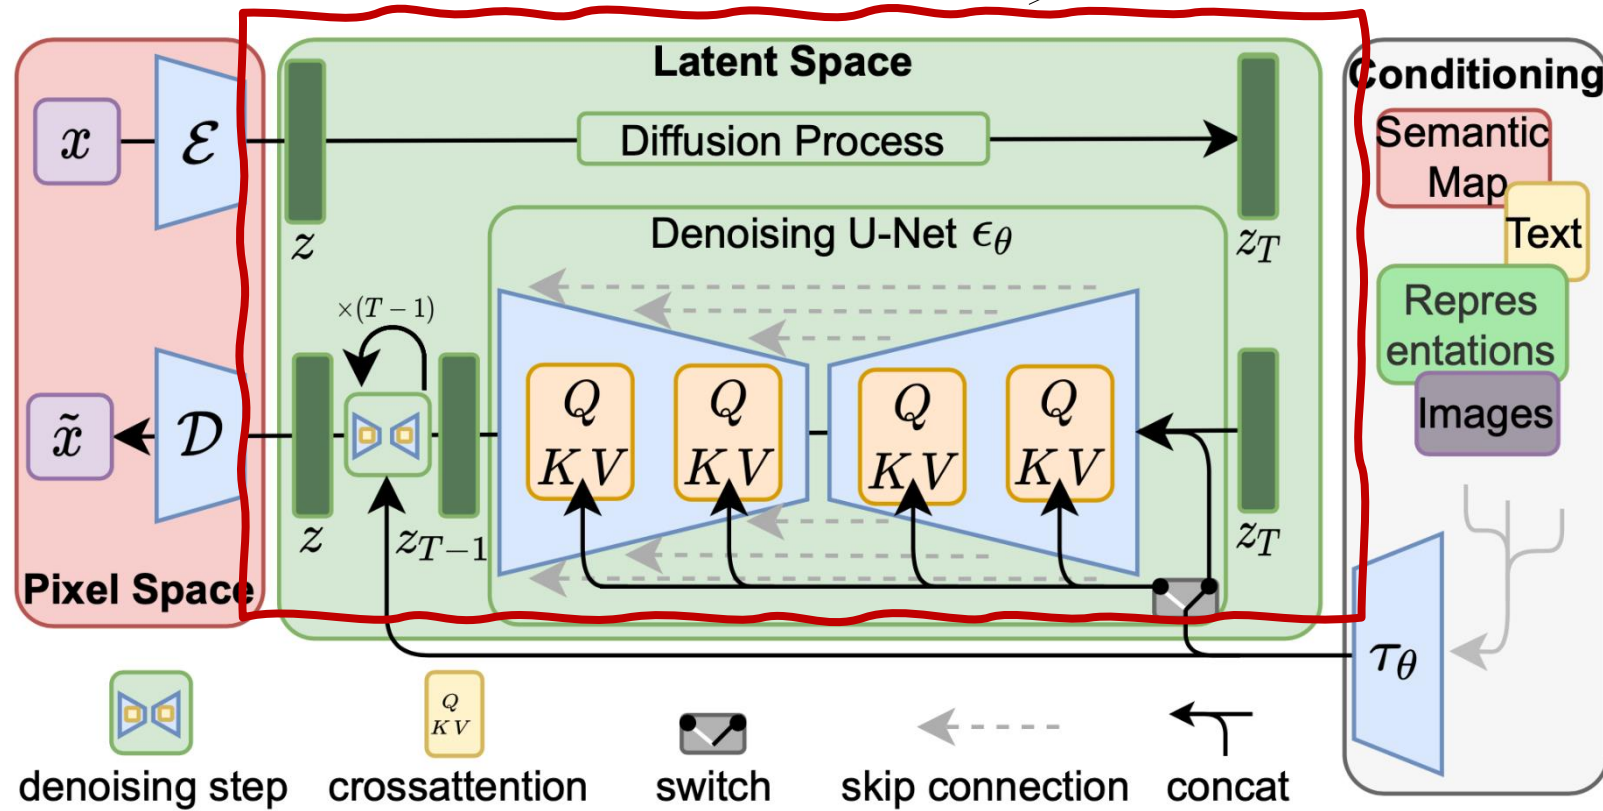


Latent Diffusion Models with Conditioning

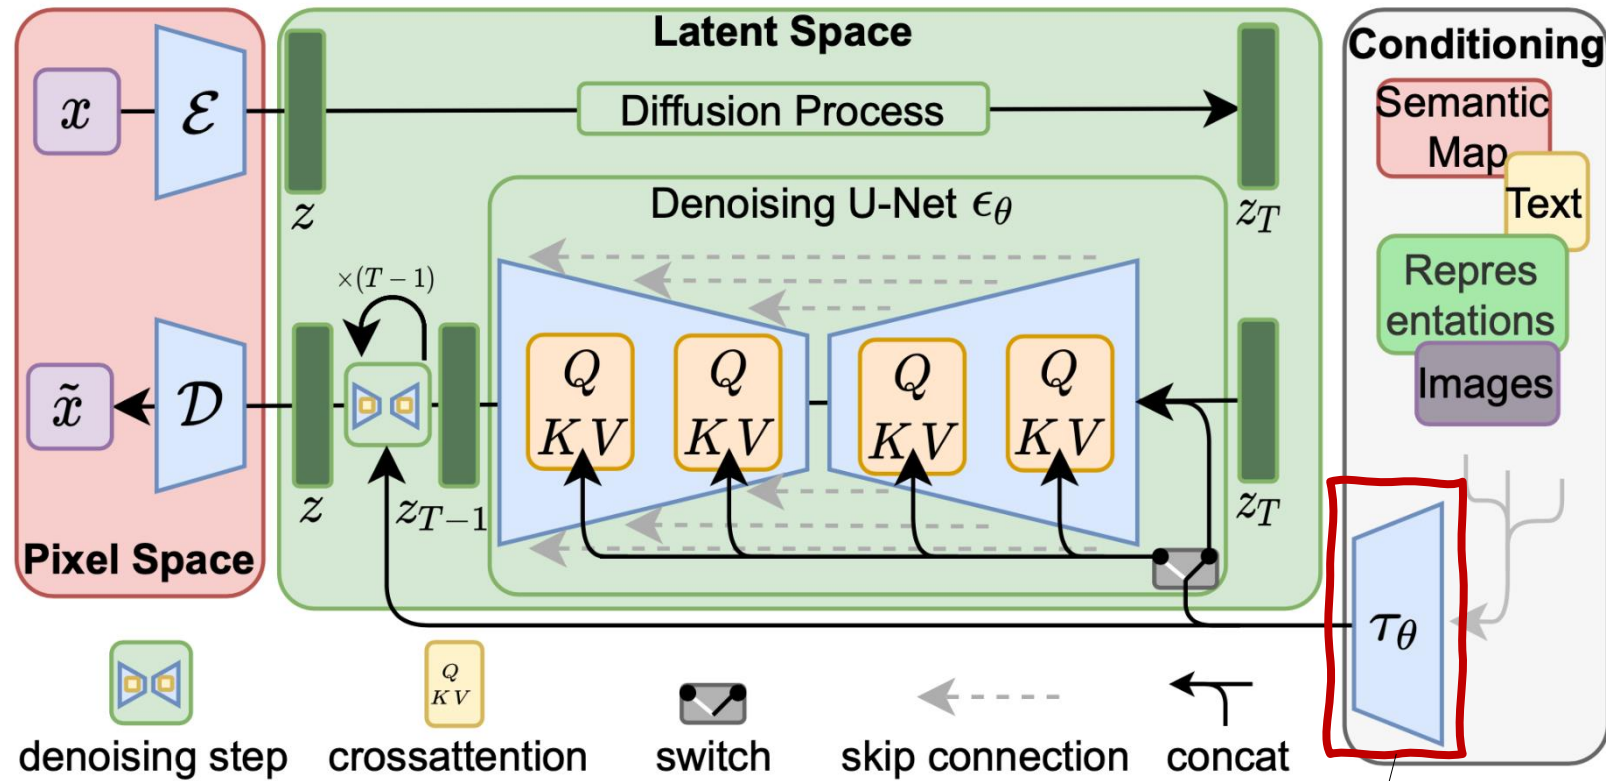


Latent Diffusion Models with Conditioning

The latent diffusion model is then pre-trained (without conditioning)

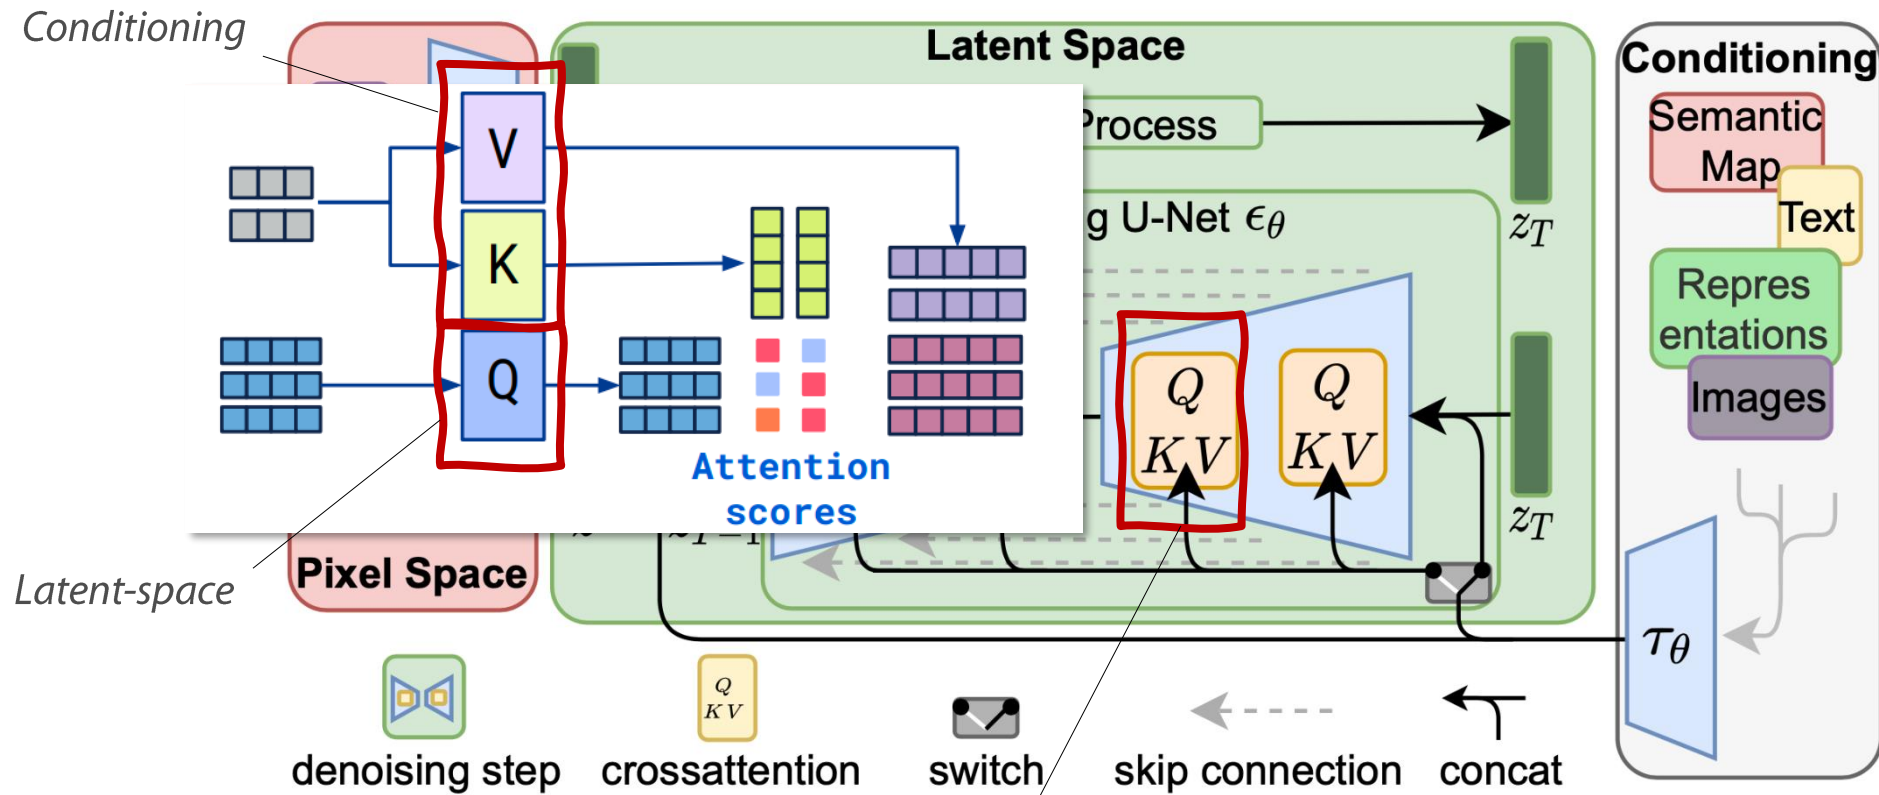


Latent Diffusion Models with Conditioning



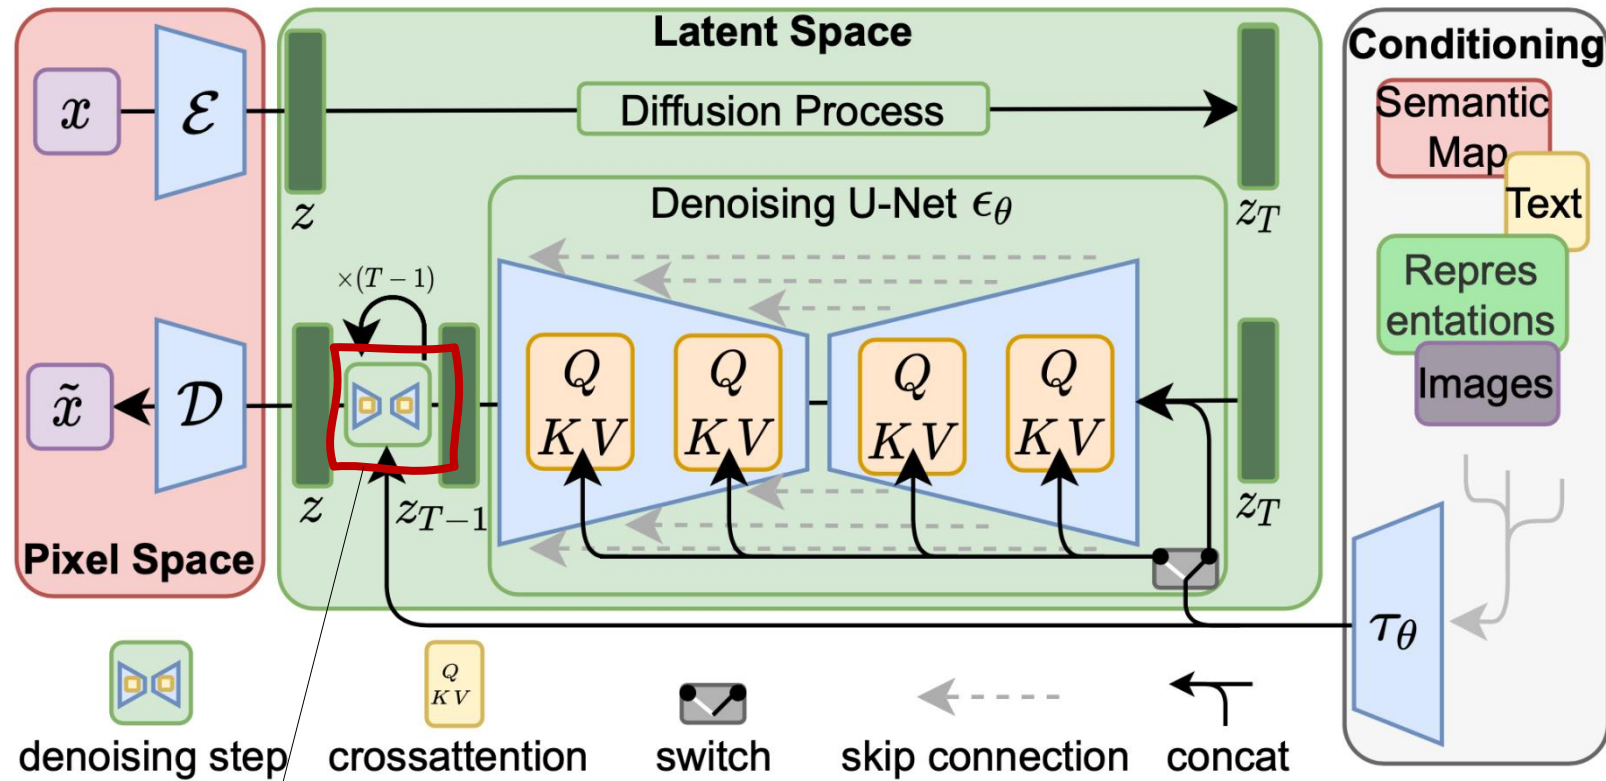
A suitable encoder of the conditioning elements is pre-trained separately

Latent Diffusion Models with Conditioning



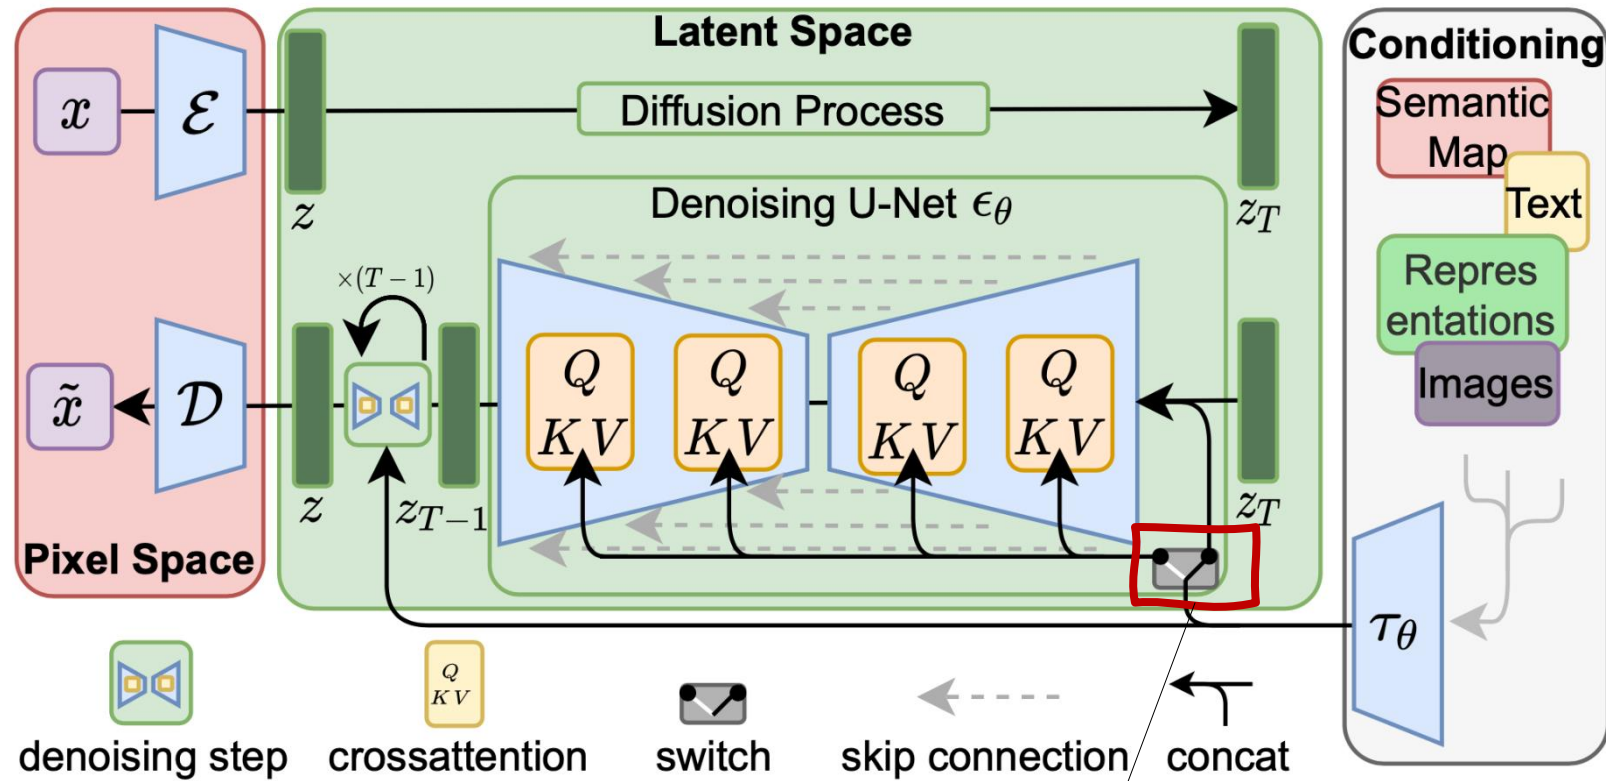
Latent-space representations and embedded condition elements are combined via cross-attention

Latent Diffusion Models with Conditioning



The same step is iterated
 $T - 1$ more times

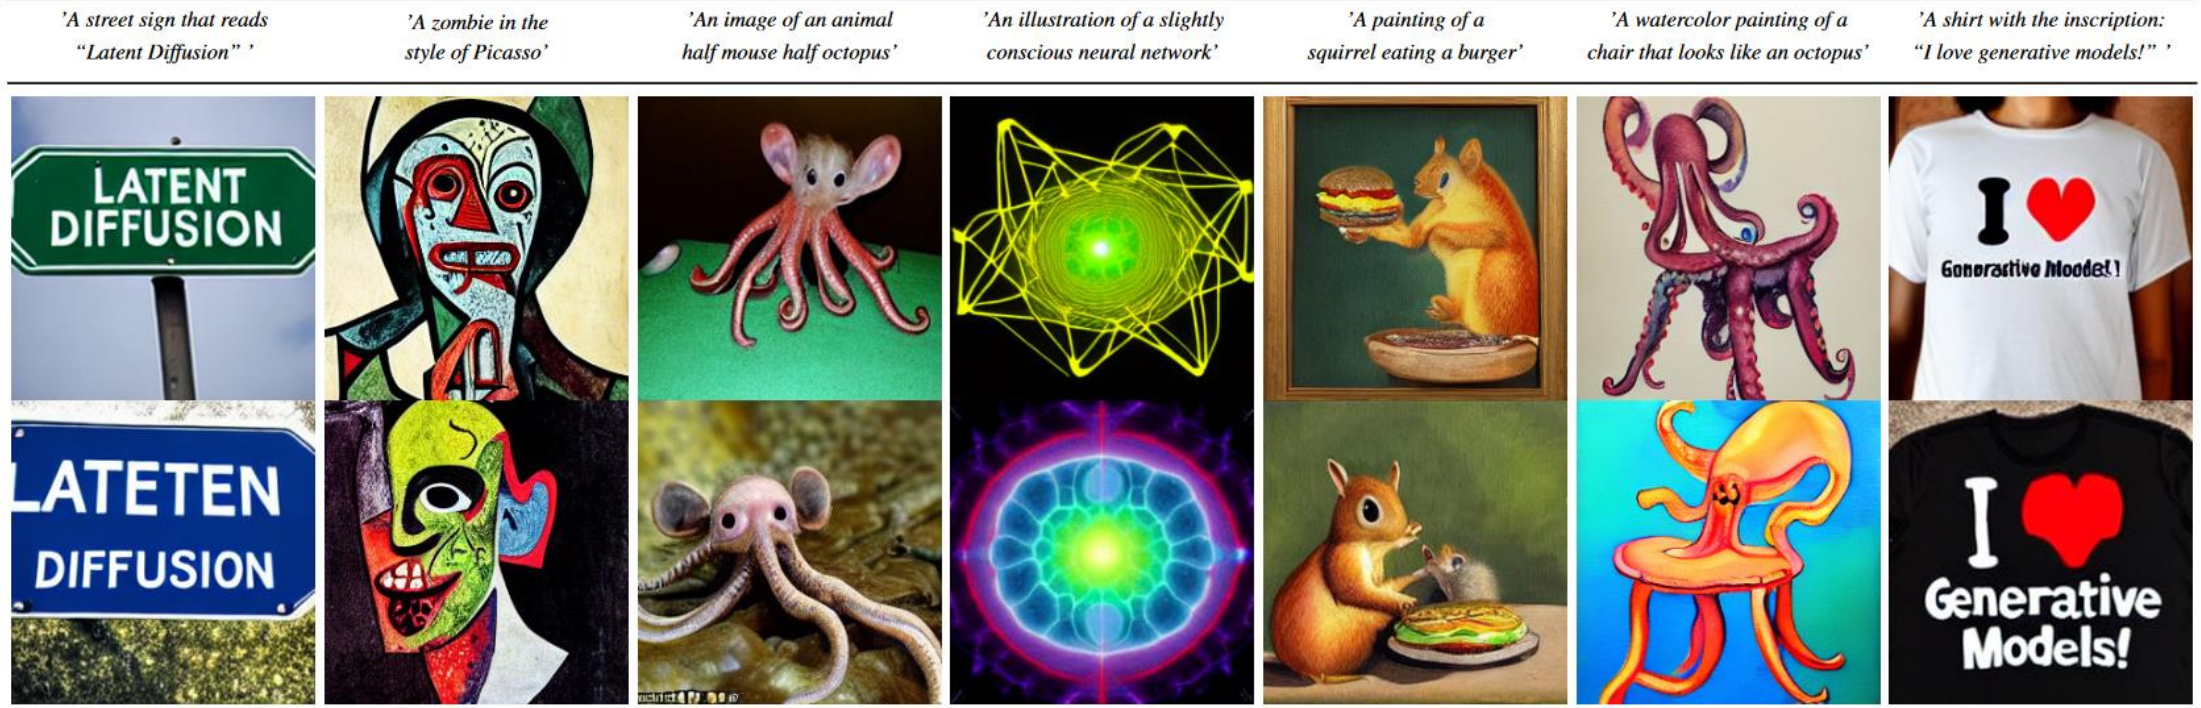
Latent Diffusion Models with Conditioning



The switch is for multi-modality:
if the conditioning element is a class or text, use cross-attention,
if the input is an image, use concatenation

Latent Diffusion Models with Conditioning

Text-to-Image Synthesis on LAION. 1.45B Model.



Links

<https://poloclub.github.io/diffusion-explainer/>

<https://blog.marvik.ai/2023/11/28/an-introduction-to-diffusion-models-and-stable-diffusion/>

<https://theaisummer.com/diffusion-models/>

<https://learnopencv.com/denoising-diffusion-probabilistic-models/>

<https://www.assemblyai.com/blog/diffusion-models-for-machine-learning-introduction/>

<https://lilianweng.github.io/posts/2021-07-11-diffusion-models/>

<https://www.superannotate.com/blog/diffusion-models>

<https://encord.com/blog/diffusion-models/>