Deep latent variable models

TRAIL doctoral seminars

November 5, 2021

Prof. Gilles Louppe
g.louppe@uliege.be

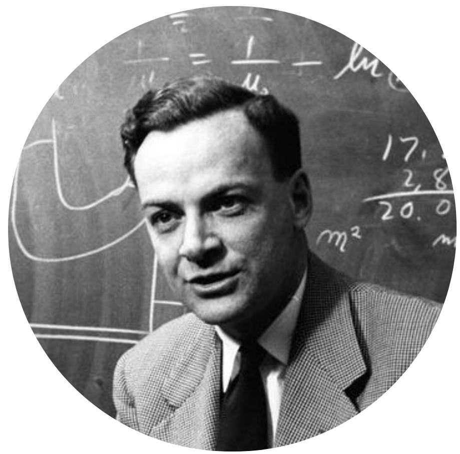

"What I cannot create, I do not understand."

Richard Feynman

A generative model is a probabilistic model p that can be used as a simulator of the data. Its purpose is to generate synthetic but realistic high-dimensional data

$$\mathbf{x} \sim p(\mathbf{x}; \theta),$$

that is as close as possible from the unknown data distribution $p(\mathbf{x})$.

Why generative models

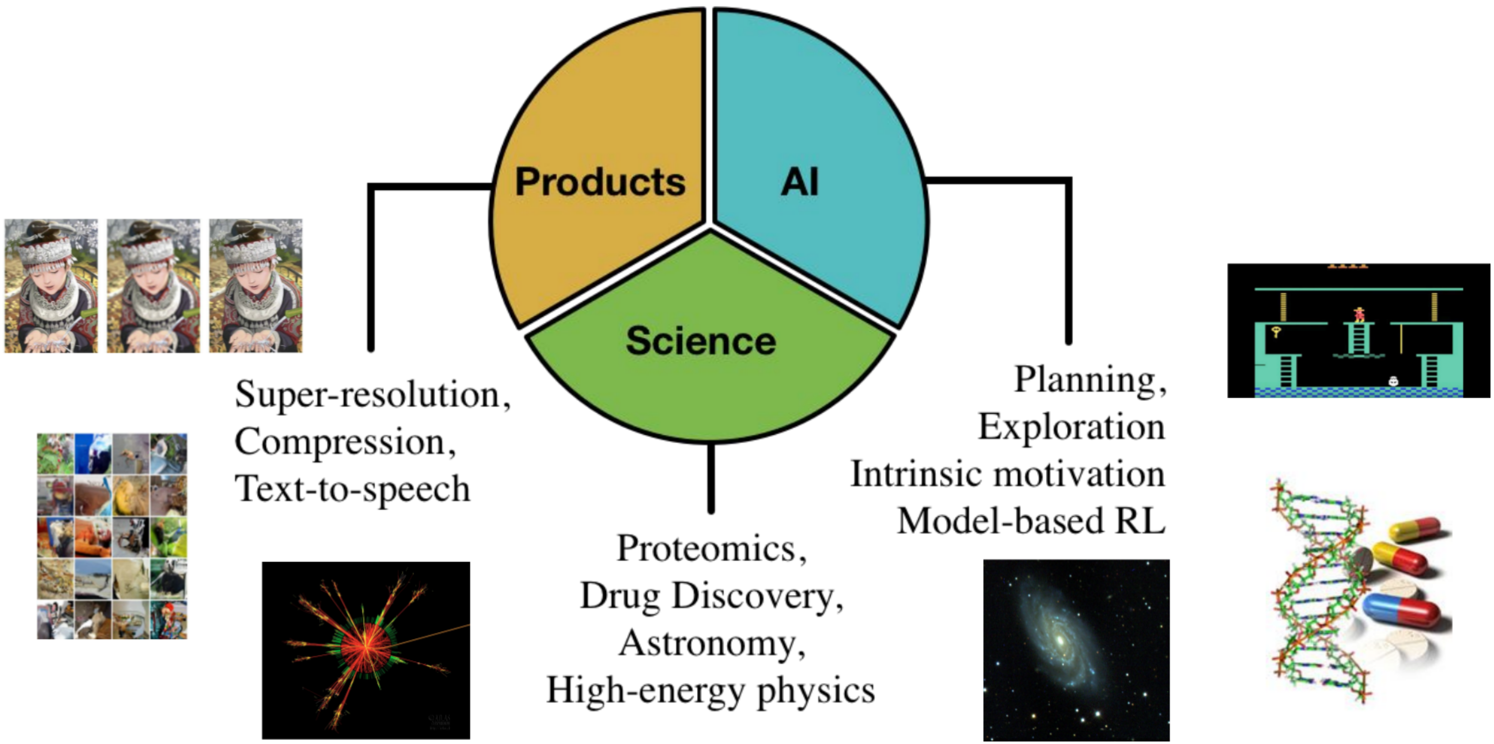

Generative models have a role in many important problems

Outline

I. Variational inference

II. Variational auto-encoders (VAEs)

III. Hierarchical VAEs

IV. Denoising Diffusion Probabilistic Models

V. VAE priors

Part I: Variational inference

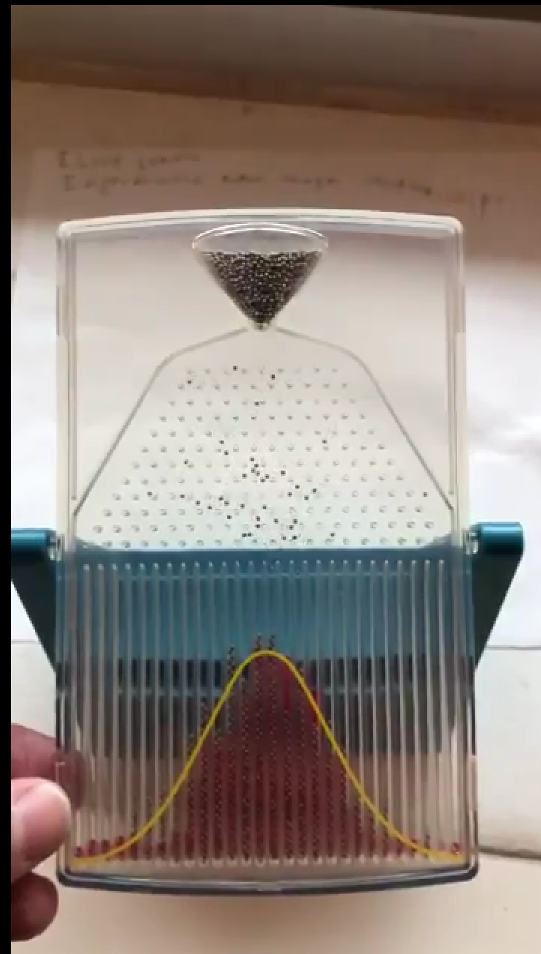

See blackboard.

Part II: Variational auto-encoders

See blackboard.

See code example 1.

Examples

Consider as data $\mathbf{d}$ the MNIST digit dataset:

(a) 2-D latent space (b) 5-D latent space (c) 10-D latent space (d) 20-D latent space

Figure 5: Random samples from learned generative models of MNIST for different dimensionalities of latent space.

(Kingma and Welling, 2013)

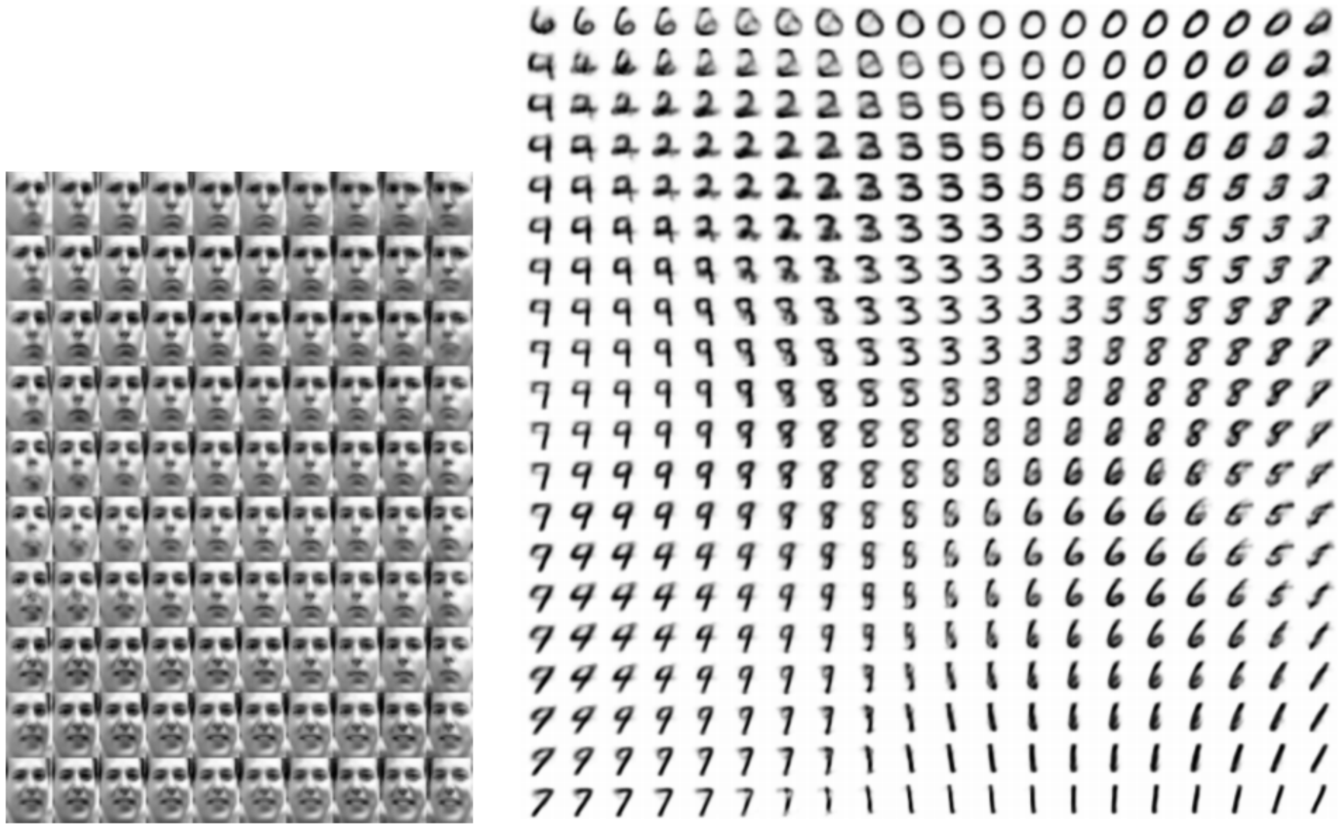

(a) Learned Frey Face manifold (b) Learned MNIST manifold

Figure 4: Visualisations of learned data manifold for generative models with two-dimensional latent space, learned with AEVB. Since the prior of the latent space is Gaussian, linearly spaced coordinates on the unit square were transformed through the inverse CDF of the Gaussian to produce values of the latent variables $\mathbf{z}$. For each of these values $\mathbf{z}$, we plotted the corresponding generative $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})$ with the learned parameters $\boldsymbol{\theta}$.

(Kingma and Welling, 2013)

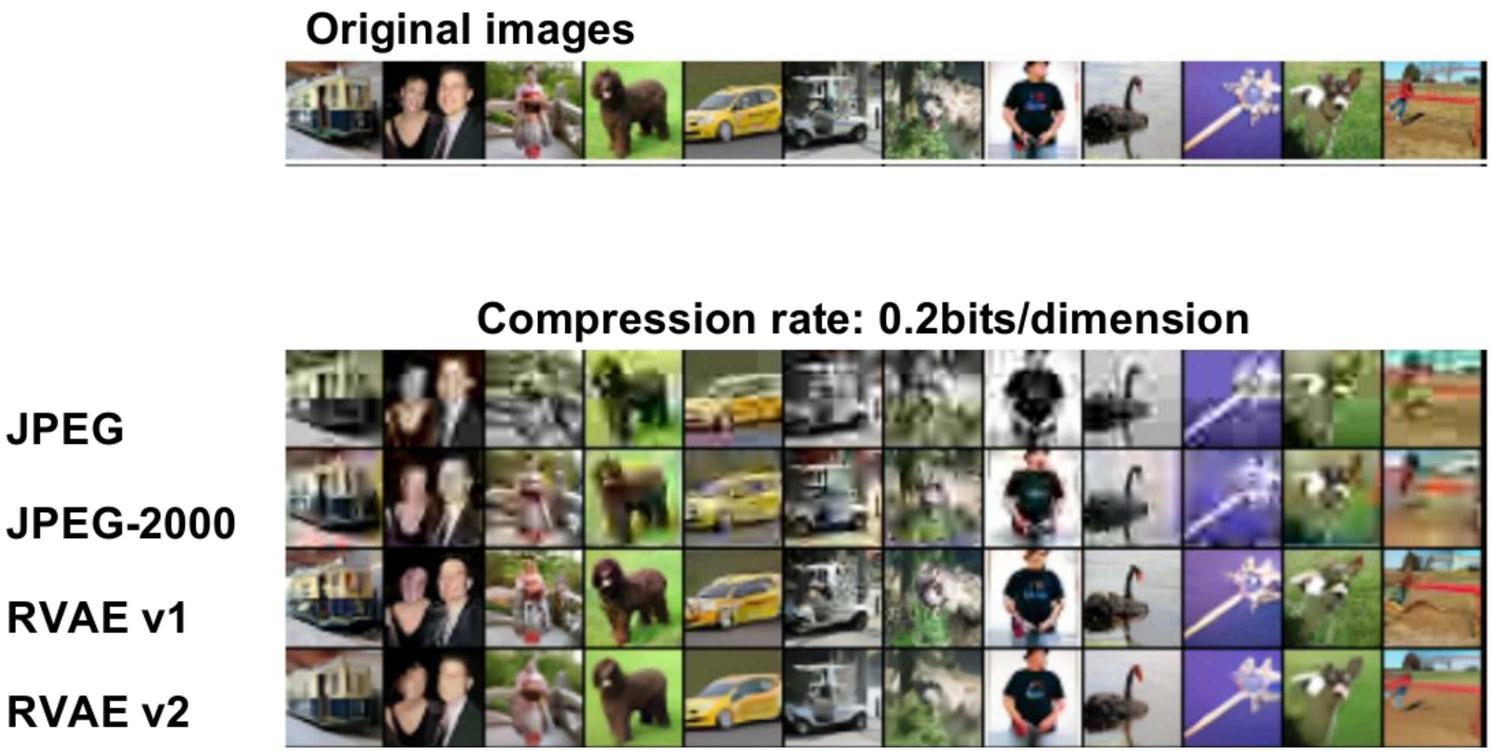

Hierarchical **compression of images and other data**, e.g., in video conferencing systems (Gregor et al, 2016).

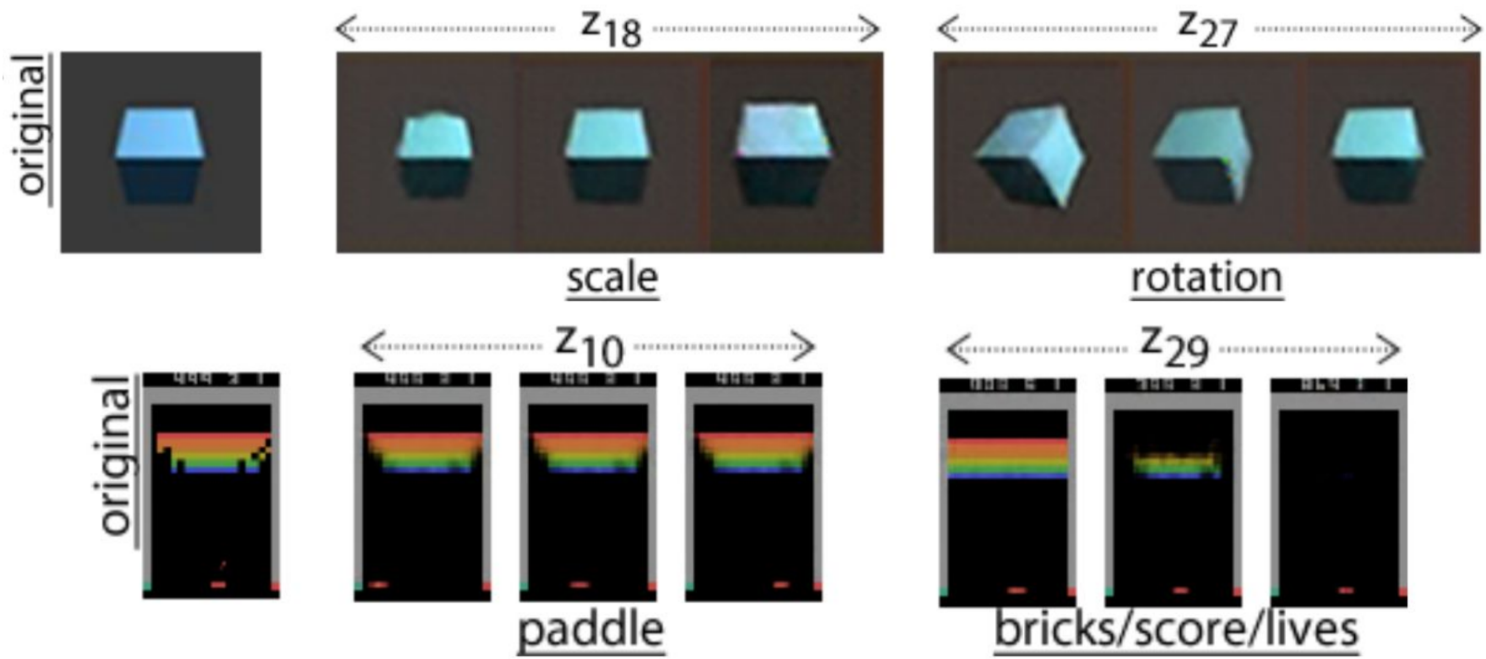

Understanding the factors of variation and invariances (Higgins et al, 2017).

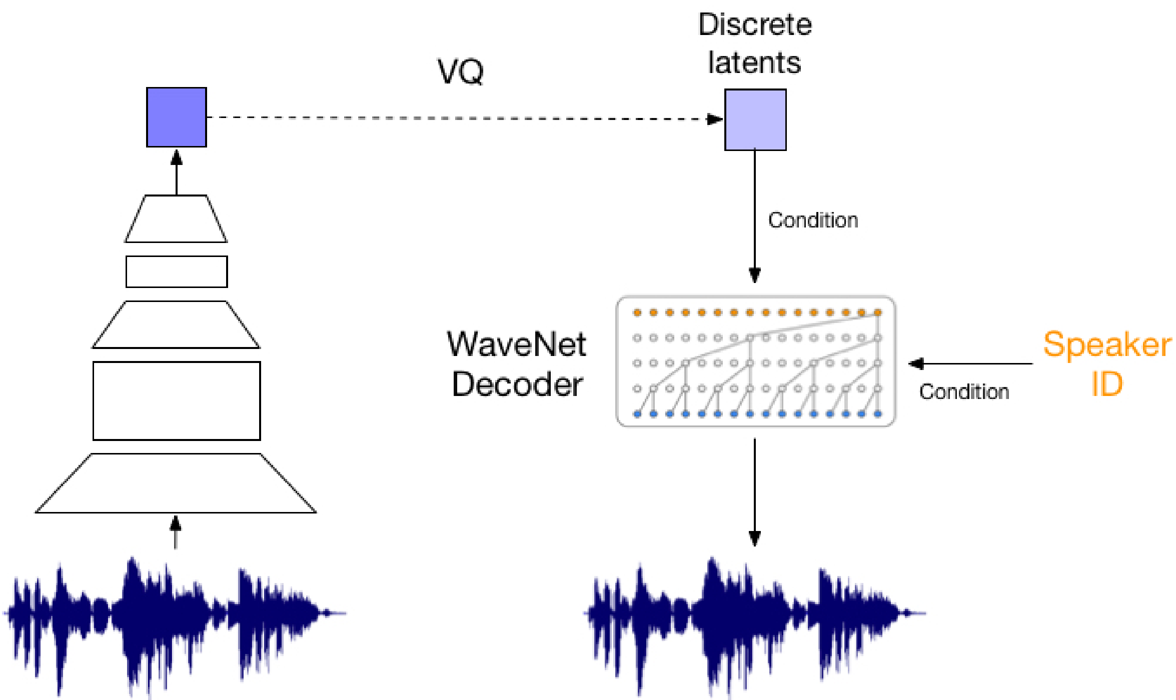

Voice style transfer [demo] (van den Oord et al, 2017).

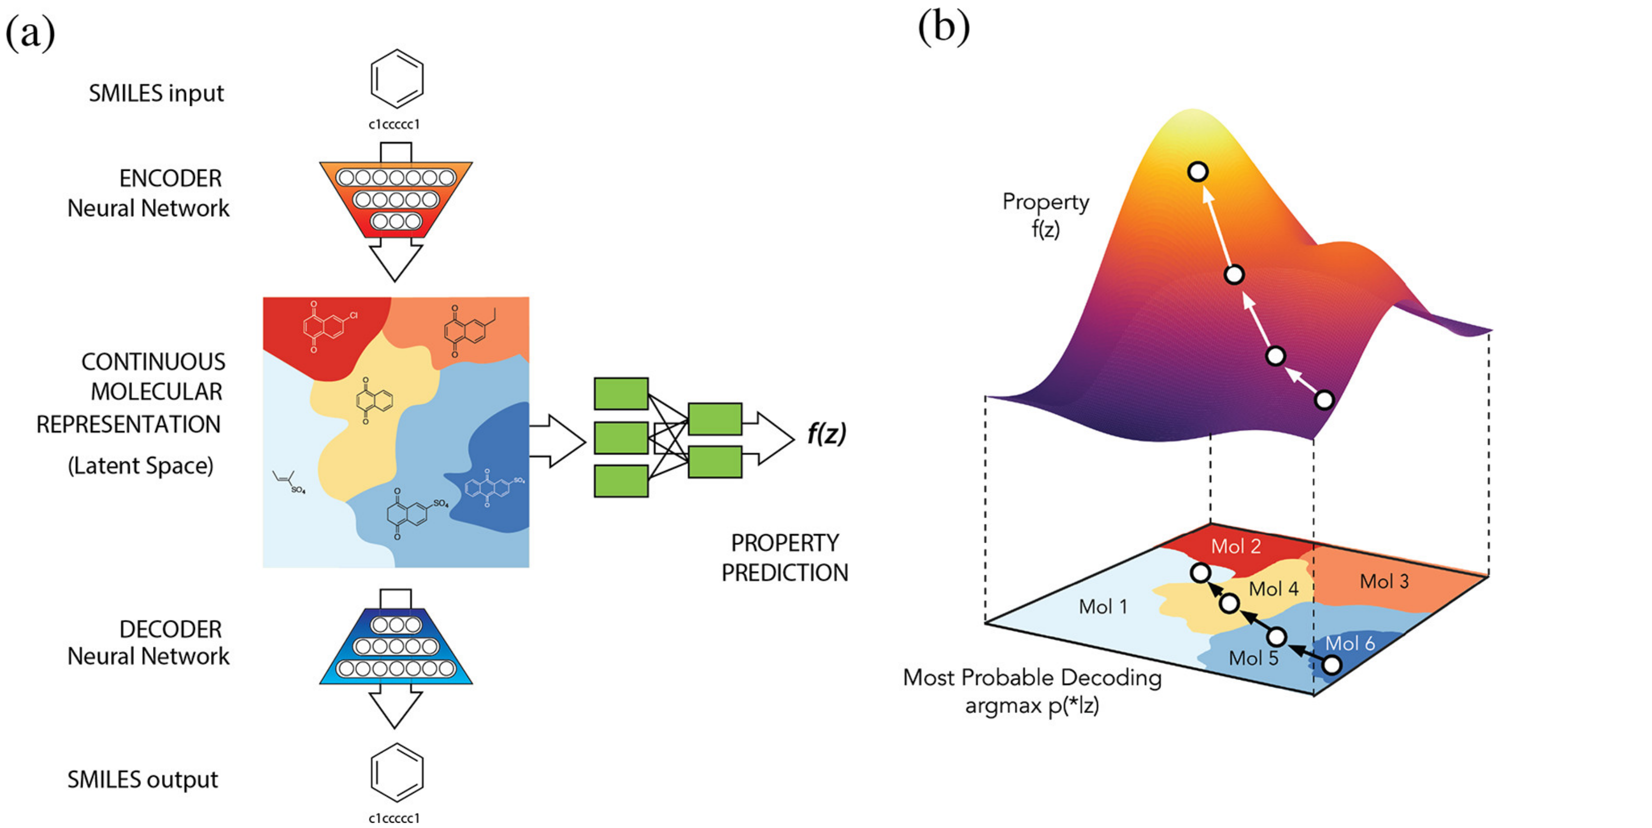

Design of new molecules with desired chemical properties (Gomez-Bombarelli et al, 2016).

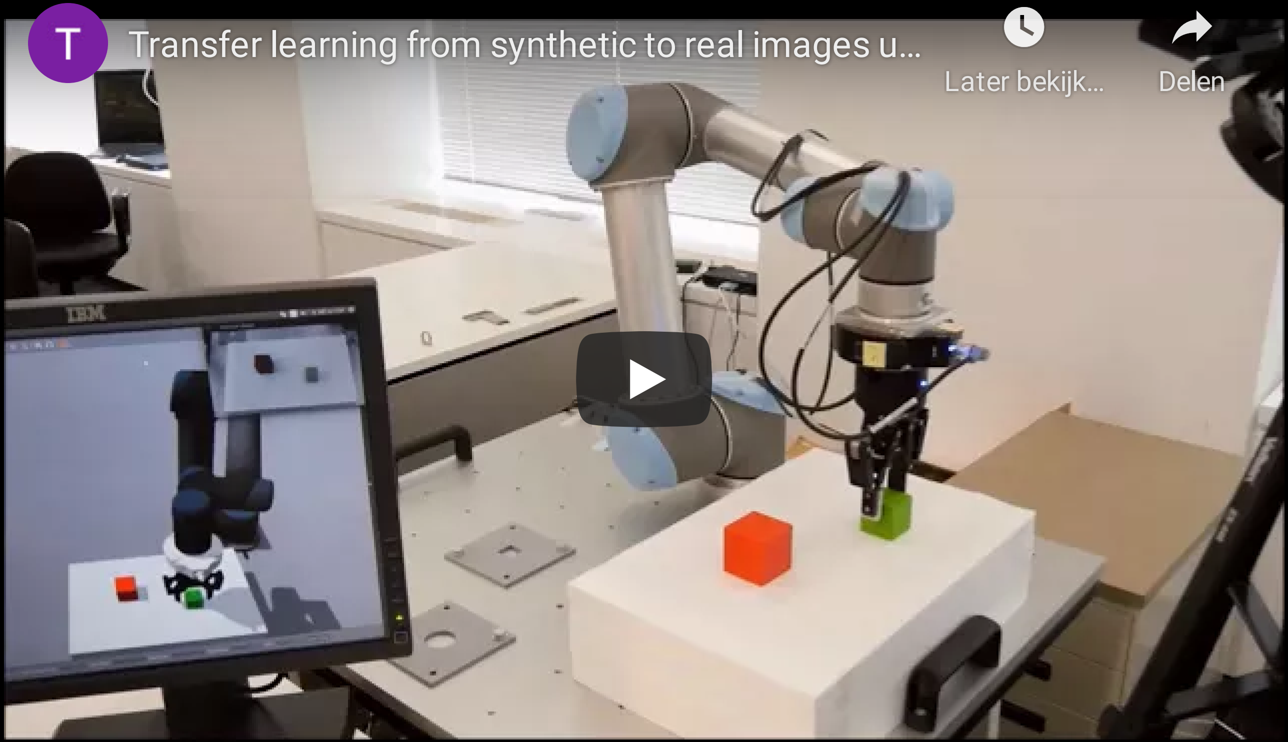

Bridging the **simulation-to-reality** gap (Inoue et al, 2017).

Part III: Hierarchical VAEs

See blackboard.

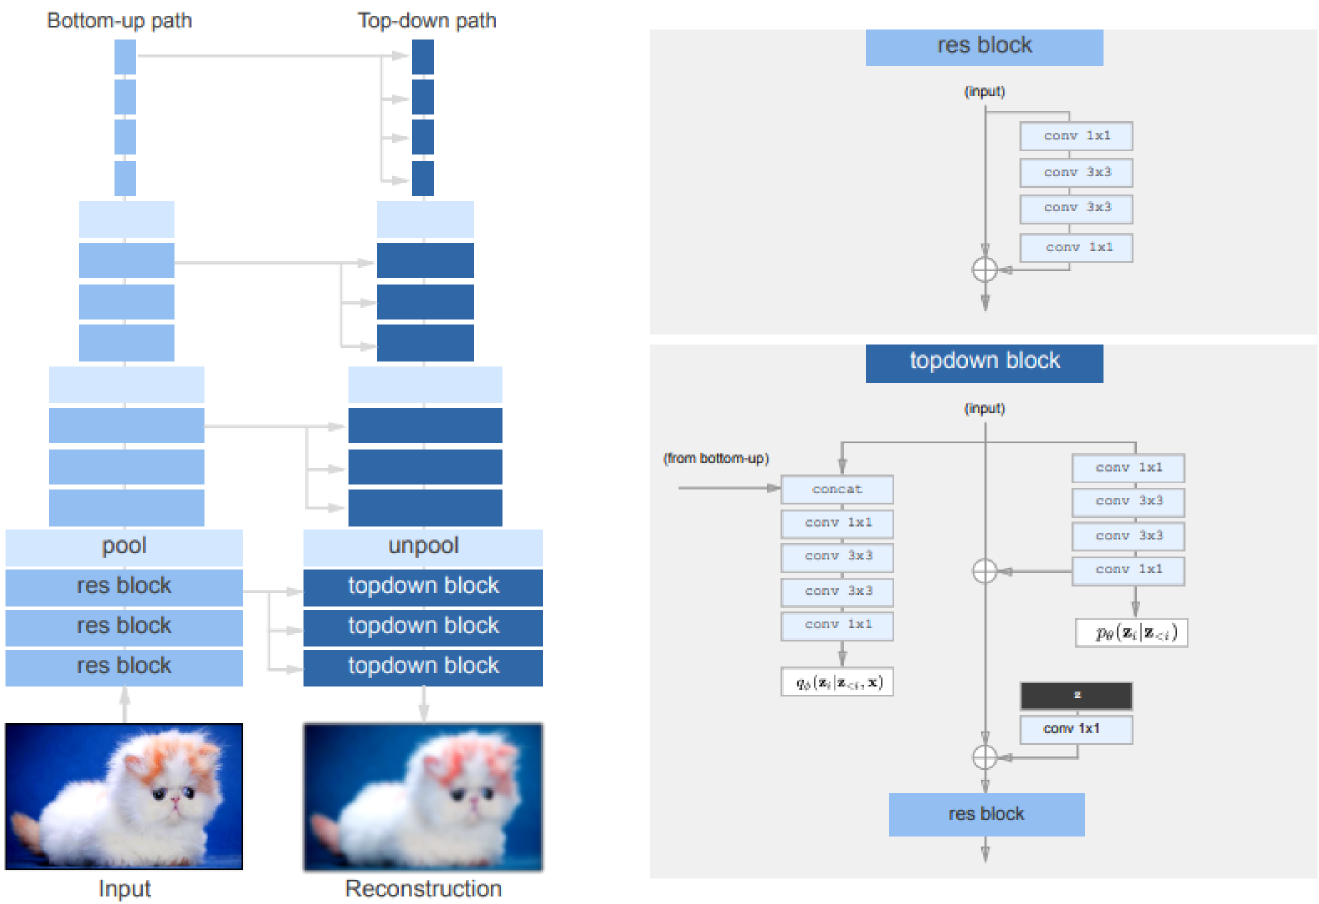

VDVAE: Very Deep VAEs (Child, 2020-2021).

VDVAE samples (Child, 2020-2021).

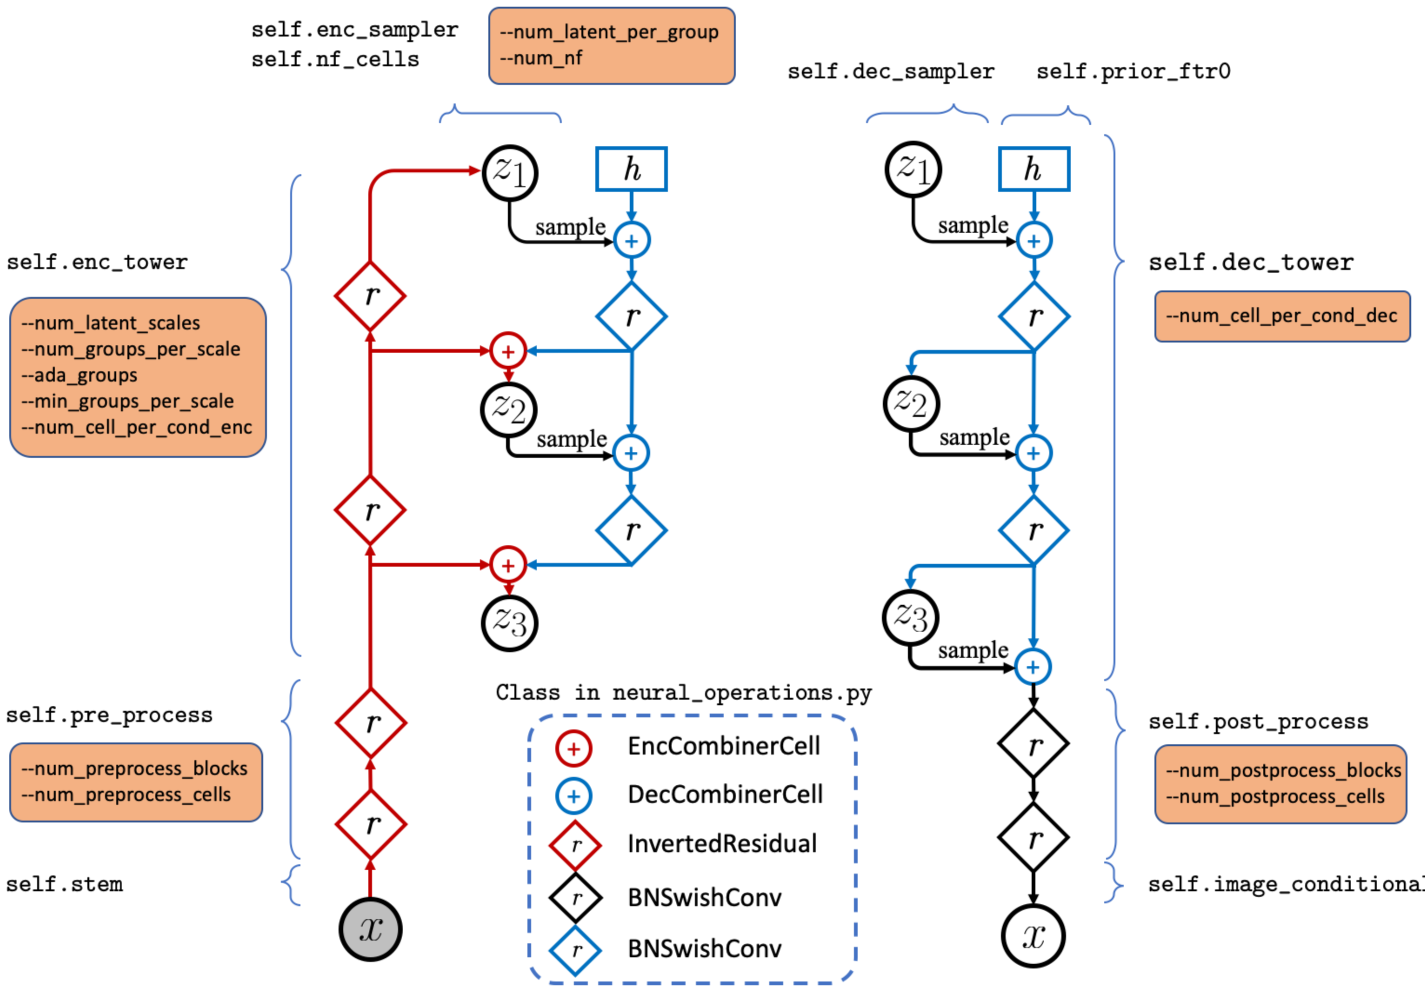

NVAE: A Deep Hierarchical Variational Autoencoder (Vahdat and Kautz, 2020).

NVAE samples (Vahdat and Kautz, 2020).

NVAE: Random walks in latent space. (Vahdat and Kautz, 2020)

Part IV: Denoising diffusion probabilistic models

See blackboard.

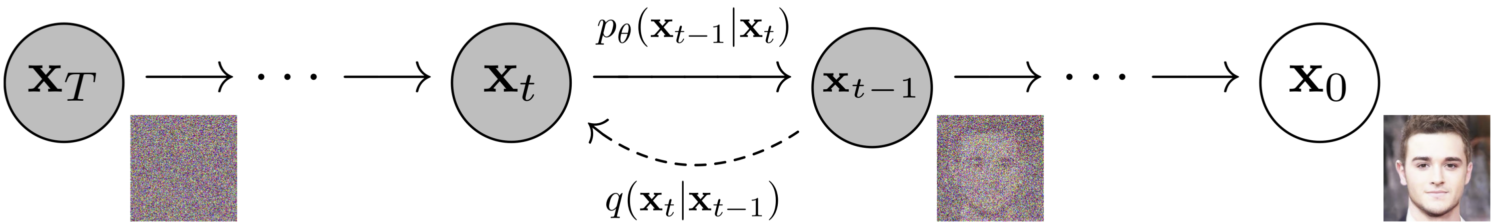

Denoising Diffusion Probabilistic Models (Ho et al, 2020).

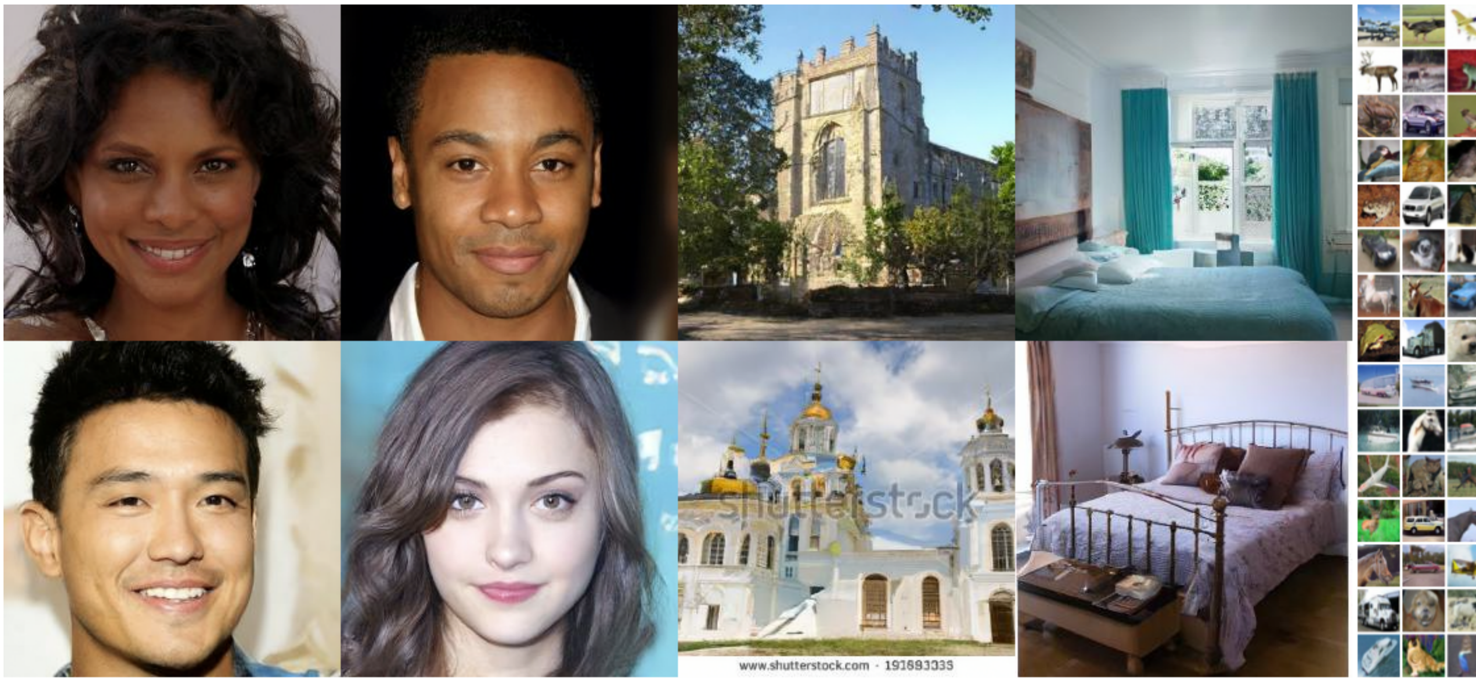

DDPM samples (Ho et al, 2020).

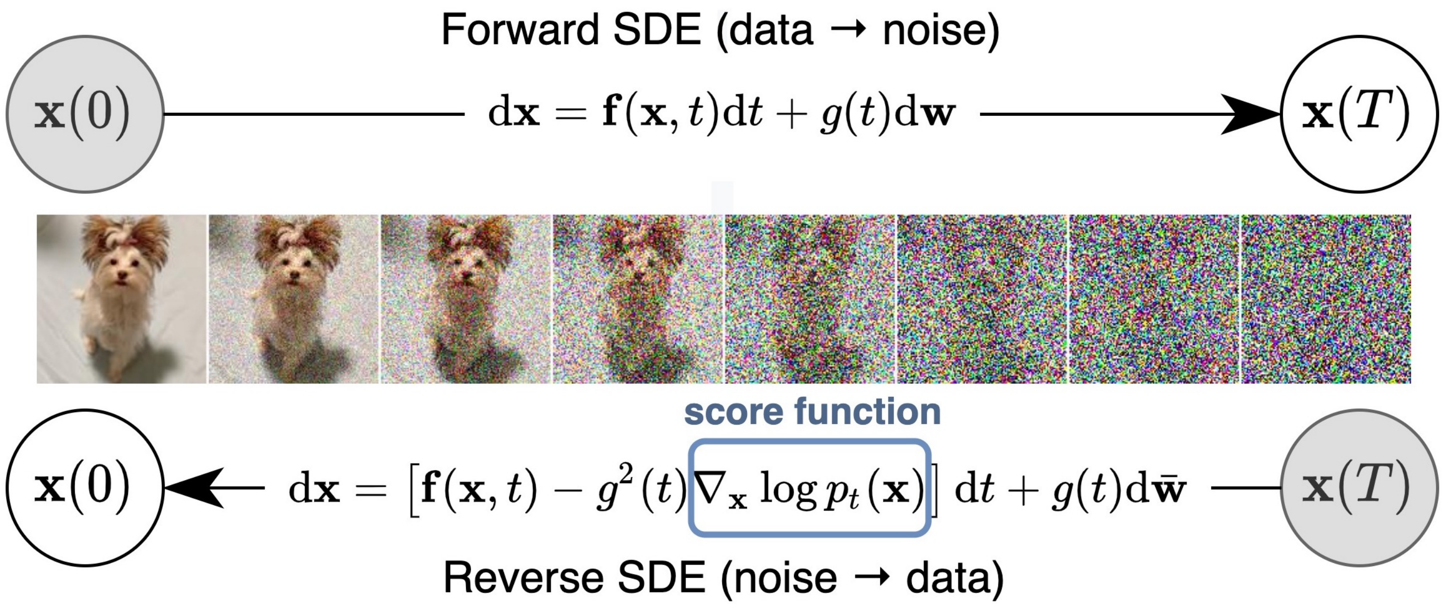

Score-based generative modeling through SDEs (Song et al, 2021).

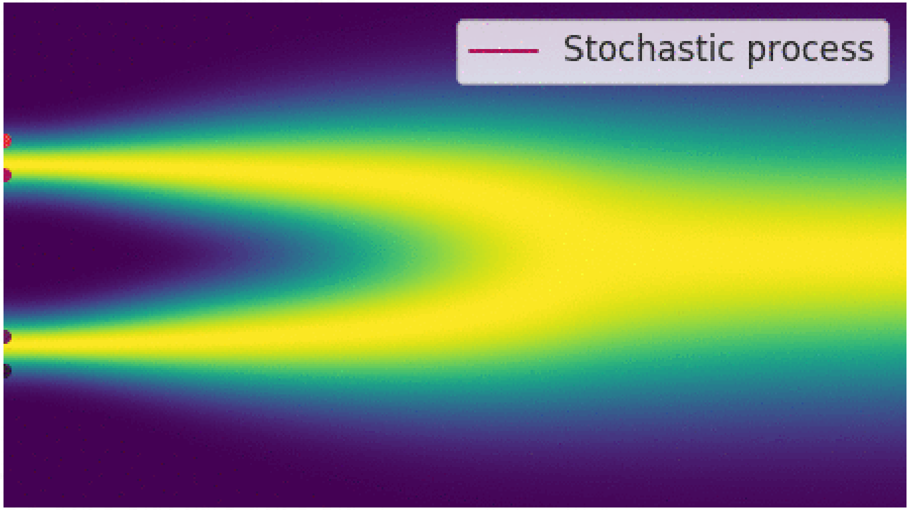

Perturbing data with an SDE (Song et al, 2021)

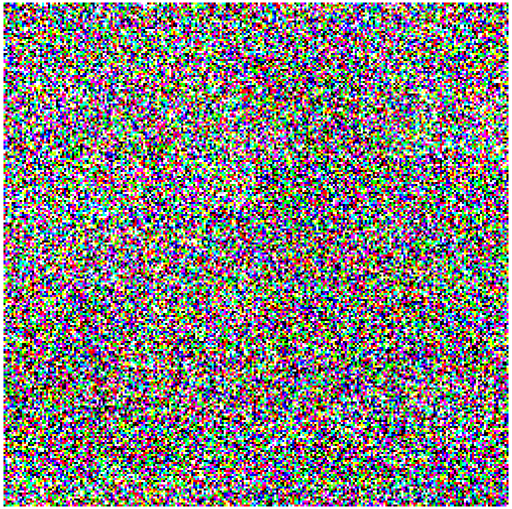

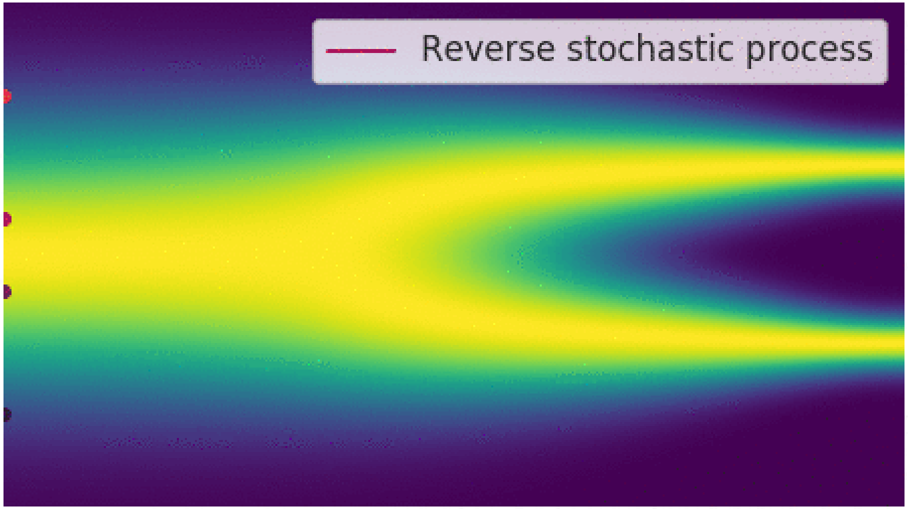

Reversing the SDE for sample generation (Song et al, 2021)

(Song et al, 2021)

Part V: Priors

See blackboard.

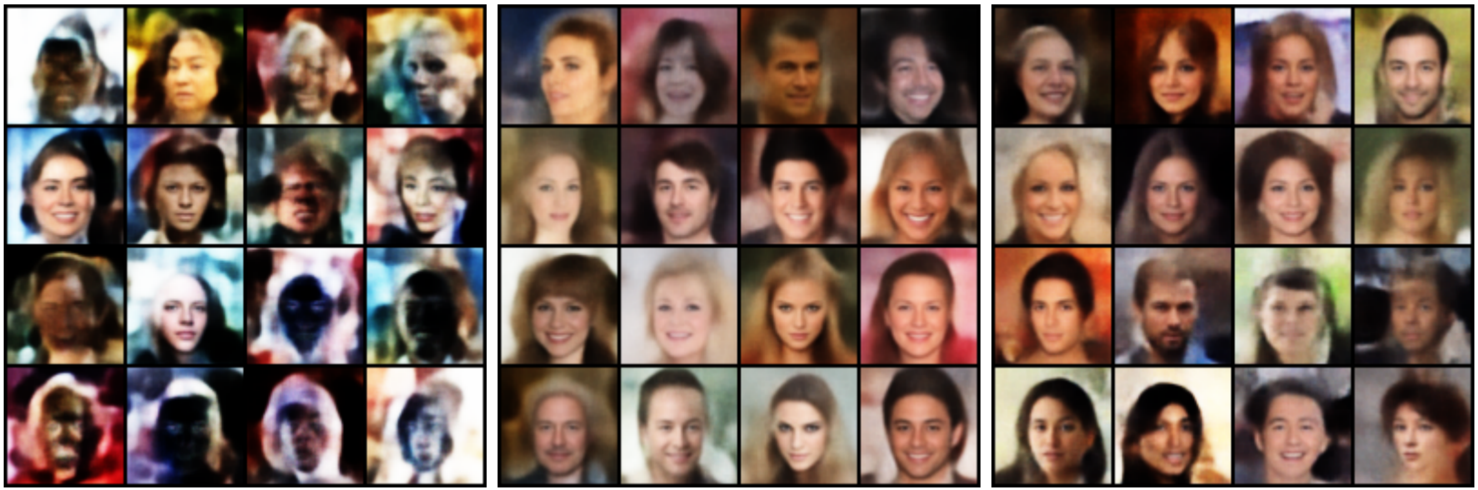

(a) Gaussian prior (b) NF prior (c) Diffused prior

Figure 2. Samples generated by a VAE trained on CelebA for three different prior models. *The diffusion prior leads to better sampling quality than the Gaussian distribution and similar to the NF prior.*

Diffusion priors in VAEs (Wehenkel and Louppe, 2021).

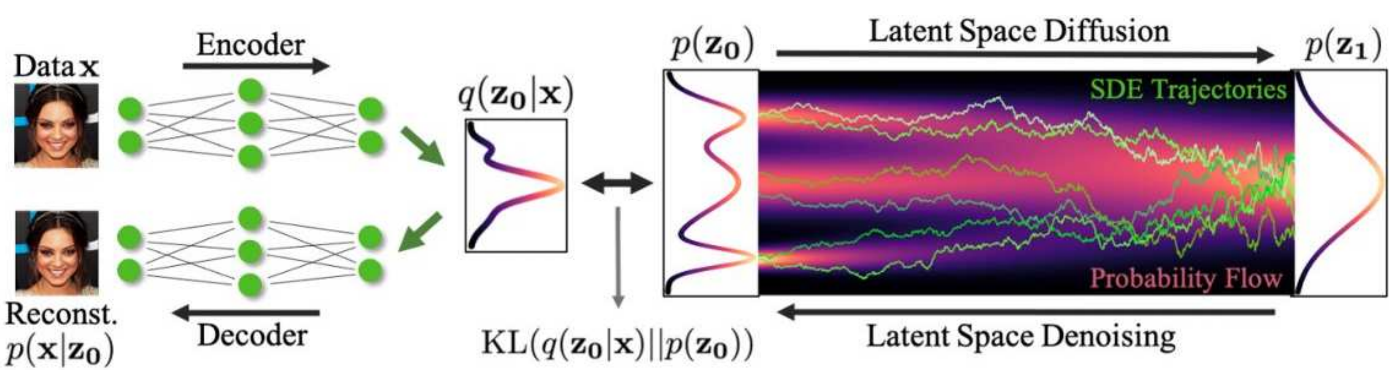

Figure 1: In our latent score-based generative model, data is mapped to latent space via encoder $q(\mathbf{z}_0|\mathbf{x})$ and a diffusion process is applied in the latent space ($\mathbf{z}_0 \rightarrow \mathbf{z}_1$). Synthesis starts from the base distribution $p(\mathbf{z}_1)$ and generates samples in latent space via denoising ($\mathbf{z}_0 \leftarrow \mathbf{z}_1$). Then, the samples are mapped from latent to data space using decoder $p(\mathbf{x}|\mathbf{z}_0)$. The model is trained end-to-end.

LSGM: Latent Score-based Generative Model (Vahdat et al, 2021).

LSGM samples (Vahdat et al, 2021).

The end.