

Lecture 21: Diffusion Models

CS480/680 Intro to Machine Learning

2023-3-28

Pascal Poupart
David R. Cheriton School of Computer Science



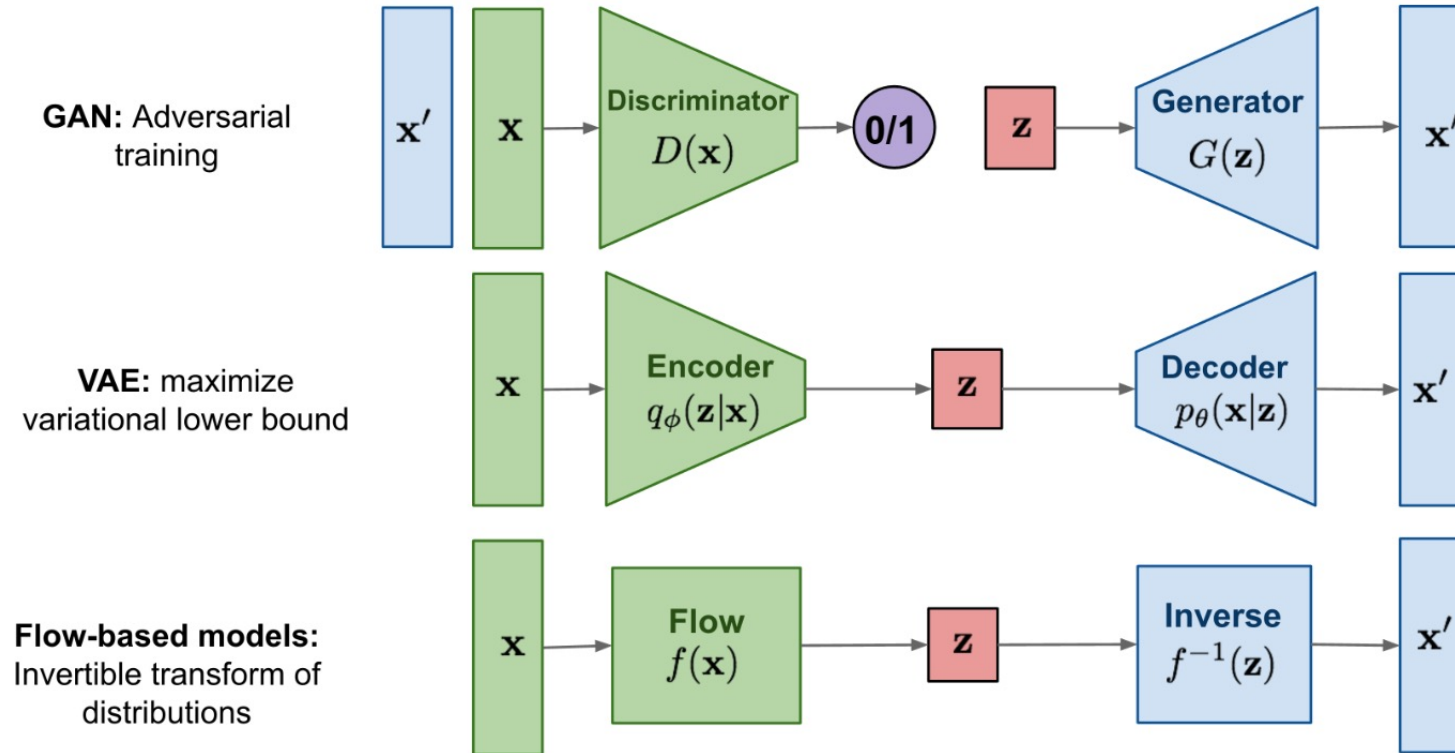
Preview

```
seed_and_prompt_sequence = [  
    3764, 'in the beginning there was nothing, just darkness',  
    1537, 'special effects render of the big bang',  
    6573, 'HD photo of a large amount of spiral galaxies',  
    1791, 'early planet formation in the solar system',  
    9973, 'the Hadean earth was bombarded with asteroids and massive volcanic eruptions',  
    736, 'panoramic view of earth with ocean surrounding newly formed land and volcanos',  
    3639, 'hydrothermal vents at the bottom of the ocean',  
    3559, 'bacteria under a microscope',  
    4724, 'bacteria under a microscope',  
    3359, 'ammonites floating in the ocean',  
    6344, 'the first reptile to leave the ocean and crawl onto the land',  
    6344, 'the first reptile to leave the ocean and crawl onto the land',  
    6813, 'massive brachiosaurus walking amidst a green mountain range',  
    6678, 'the extinction of the dinosaurs be a huge meteorite',  
    7450, 'small mammals thriving in a cave',  
    9766, 'small, prehistoric mammals living in the jungle',  
    5009, 'group of monkeys in the forest',  
    7287, 'HD photograph of neanderthal, the first man',  
    6008, 'cave painting',  
    208, 'cavemen tribe gathered around a fire at night looking at the stars',  
    2222, 'maasai tribe hunting on the savanna with spears',  
    571, 'homo sapiens using stone tools',  
    632, 'a small, tribal village with huts',  
    1332, 'at the dawn of civilization, small villages emerged',  
    2496, 'ancient egypt, the first massive civiliation',  
    1869, 'the height of the roman empire, incredible architecture, by Greg Rutkowski',  
    7559, 'medieval town square',  
    1265, 'medieval city',  
    6628, 'the skyline of New York city',
```

Xander Steenbrugge created the amazing **Voyage through Time** video below using stable diffusion with the input prompts shown in the figure.



Recap




From <https://lilianweng.github.io/>

Diffusion Model

- Stochastic autoencoder (encoder introduces noise and decoder denoises the data)
 - Generates better data than variational autoencoders
 - Easier to train than generative adversarial networks and does not suffer from mode collapse
 - Special type of stochastic flow that is not restricted to invertible transformations

Forward Diffusion Process



Distribution of the
noised images

Output

Mean μ_t

Variance Σ_t

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t I)$$

Notations:

t : time step (from 0 to T)

x_0 : a data sampled from the real data distribution $q(x)$ (i.e. $x_0 \sim q(x)$)

β_t : variance schedule ($0 \leq \beta_t \leq 1$, and $\beta_0 = \text{small number}$, $\beta_T = \text{large number}$)

I : identity matrix

From Steins (medium.com)

Stochastic Transformation

- Recall the reparameterization trick:
 - When $x \sim P(x) = N(x|\mu, \sigma^2)$
then $x = \sigma\epsilon + \mu$ where $\epsilon \sim N(\epsilon|0,1)$
- Since $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = N(\mathbf{x}_t|\sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$
Then $\mathbf{x}_t = \sqrt{1 - \beta_t}\mathbf{x}_{t-1} + \sqrt{\beta_t}\epsilon$ where $\epsilon \sim N(\epsilon|\mathbf{0}, \mathbf{I})$

Stochastic Transformation

- We can speed up the noise process by computing \mathbf{x}_t in one step:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}$$

where $\boldsymbol{\epsilon} \sim N(\boldsymbol{\epsilon} | \mathbf{0}, \mathbf{1})$

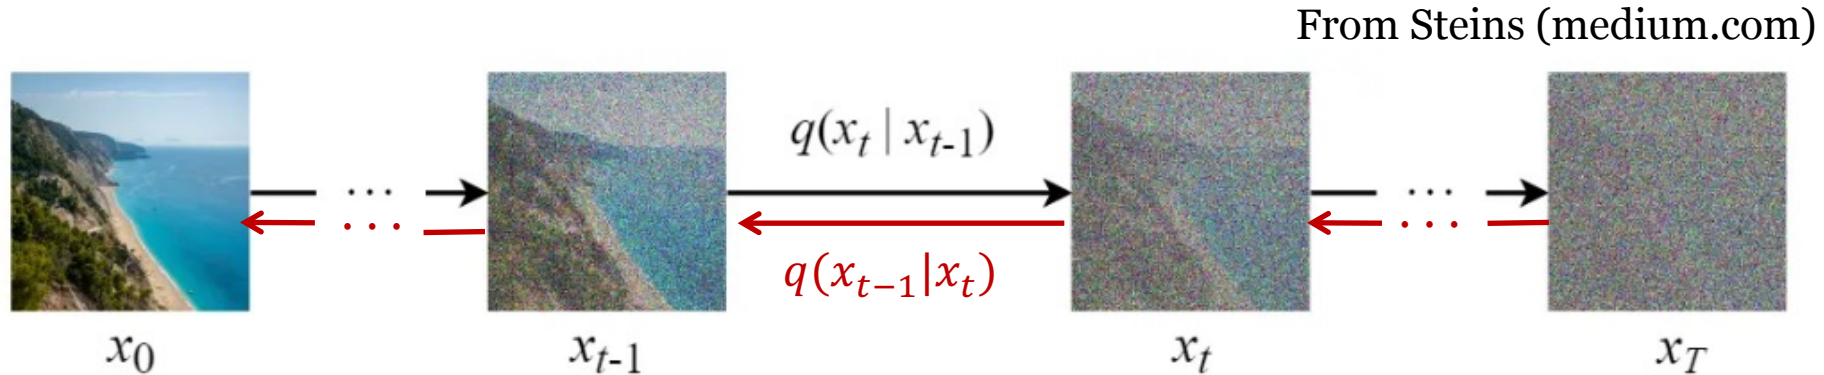
$$\bar{\alpha}_t = \prod_{i=1}^t \alpha_i \text{ and } \alpha_i = 1 - \beta_i$$

See <https://medium.com/@steinsfu/diffusion-model-clearly-explained-cd331bd41166> for derivation

- In the limit, \mathbf{x}_∞ is a random vector from an isotropic Gaussian

$$\lim_{t \rightarrow \infty} \mathbf{x}_t = \sqrt{\bar{\alpha}_\infty} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_\infty} \boldsymbol{\epsilon} = \boldsymbol{\epsilon} \text{ since } \bar{\alpha}_\infty \rightarrow 0$$

Reverse Denoising Process



- Forward factorization: $q(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T) = q(\mathbf{x}_0) \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1})$
- Reverse factorization: $q(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T) = \prod_{t=1}^T q(\mathbf{x}_{t-1} | \mathbf{x}_t) q(\mathbf{x}_T)$
 - Since joint distribution is Gaussian then $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is also Gaussian
 - $q(\mathbf{x}_{t-1} | \mathbf{x}_t) = N(\mathbf{x}_{t-1} | \tilde{\mu}_t(\mathbf{x}_t, t), \sigma_t I)$

Reverse Conditional Gaussian

- The reverse conditional $q(\mathbf{x}_{t-1}|\mathbf{x}_t) = N(\mathbf{x}_{t-1}|\tilde{\mu}_t(\mathbf{x}_t, t), \sigma_t \mathbf{I})$ does not have a closed form, but Ho, Jain and Abbeel (2020) derived the following approximation for $\tilde{\mu}_t$:

$$\tilde{\mu}_t(\mathbf{x}_t, t) \approx \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_t \right)$$

where $\boldsymbol{\epsilon}_t = N(\boldsymbol{\epsilon}|\mathbf{0}, \mathbf{I})$ is the noise introduced at step t

- We do not know $\boldsymbol{\epsilon}_t$, but we can train a neural network $\epsilon_\theta(\mathbf{x}_t, t)$ to approximate it:

$$\text{Minimize } L(\theta) = \|\boldsymbol{\epsilon}_t - \epsilon_\theta(\mathbf{x}_t, t)\|_2^2$$

Training Algorithm

Repeat


- $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
- $t \sim \text{uniform}(\{1, \dots, T\})$
- $\epsilon \sim N(\mathbf{0}, I)$
- $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$
- $\theta \leftarrow \theta + \eta \nabla_{\theta} \|\epsilon - \epsilon_{\theta}(\mathbf{x}_t, t)\|_2^2$

Until convergence

Training Algorithm

For each training step:

1. Randomly select a time step & encode it



2. Add noise to image



$$x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$$

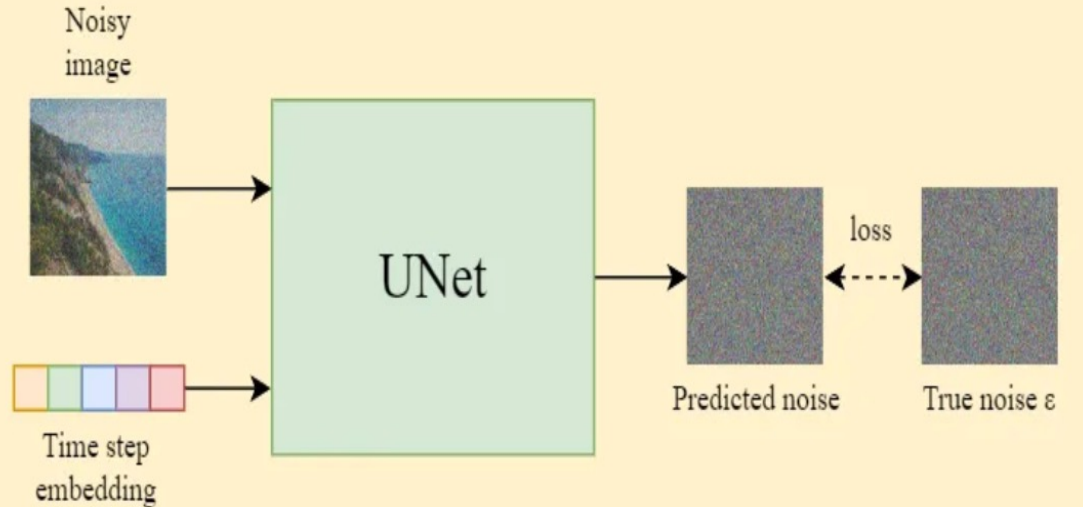
Adjust the amount of noise according to the time step t

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

$$\alpha_t = 1 - \beta_t$$

$$\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$$

3. Train the UNet




From Steins (medium.com)

UNet

Special type of fully convolutional neural network

From Cai, Wu and Chen 2022



Data Generation Algorithm

- $\mathbf{x}_T \sim N(\mathbf{x}|\mathbf{0}, I)$
- For $t = T, \dots, 1$ do
 - $\epsilon \sim N(\epsilon|\mathbf{0}, I)$ if $t > 1$, else $\epsilon = \mathbf{0}$
 - $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_{\theta}(\mathbf{x}_t, t) \right) + \sqrt{\sigma_t} \epsilon$
- Return \mathbf{x}_0

Data Generation Algorithm

1. Sample a Gaussian noise

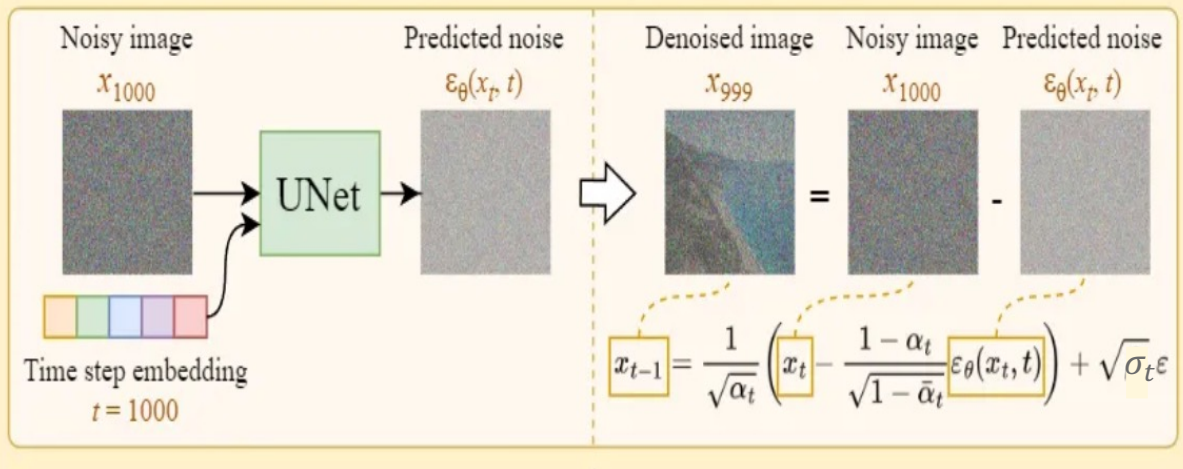
$$x_T \sim N(0, I)$$

E.g. $T = 1000$

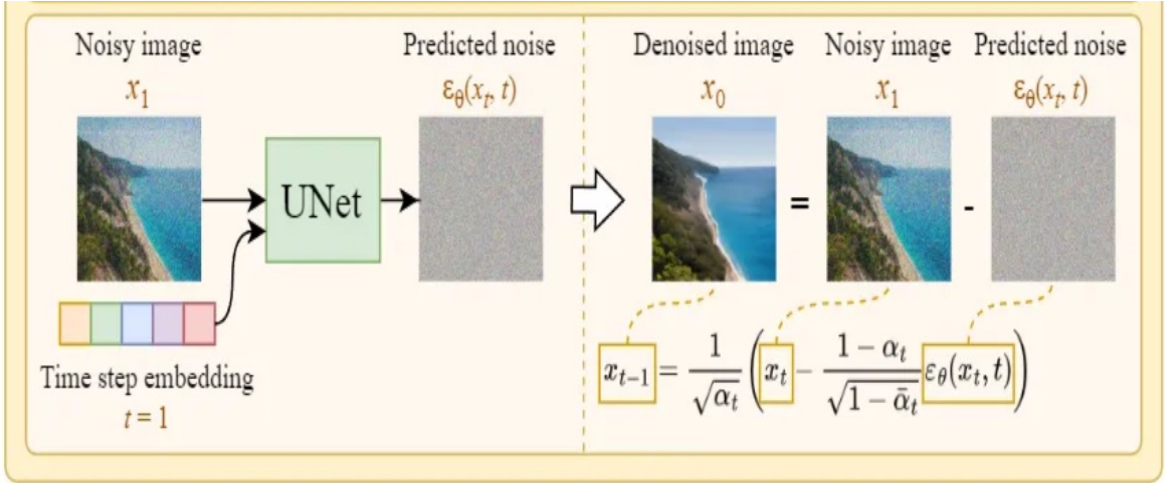
$$x_{1000} \sim N(0, I)$$




2. Iteratively denoise the image



...



3. Output the denoised image



From Steins (medium.com)

Results

Ho, Jain and Abbeel (2020)



Figure 3: LSUN Church samples. FID=7.89




Figure 4: LSUN Bedroom samples. FID=4.90





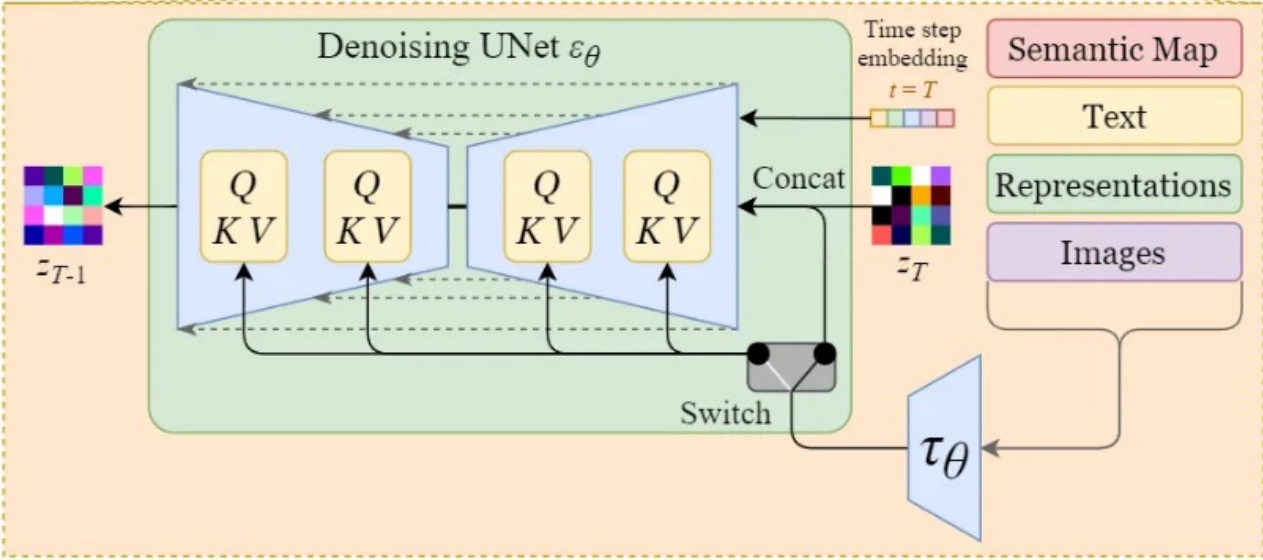
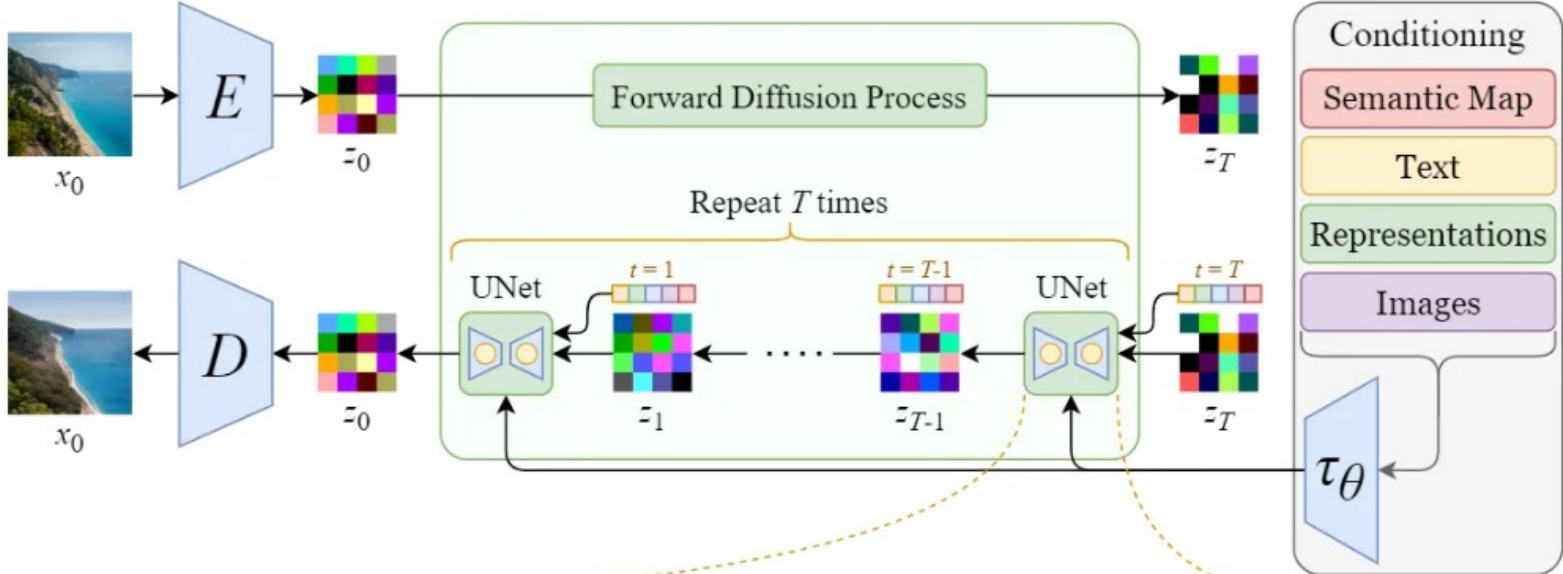
Figure 8: Interpolations of CelebA-HQ 256x256 images with 500 timesteps of diffusion.

Latent Diffusion Model (a.k.a. Stable Diffusion)

- Rombach, Blattman et al., 2022
- **Speed up:** performing the diffusion in a low dimensional latent space
- **Conditional generation:** condition denoising on text, images, etc.



Full Architecture



From Steins (medium.com)

Results

- Rombach, Blattman et al., 2022

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

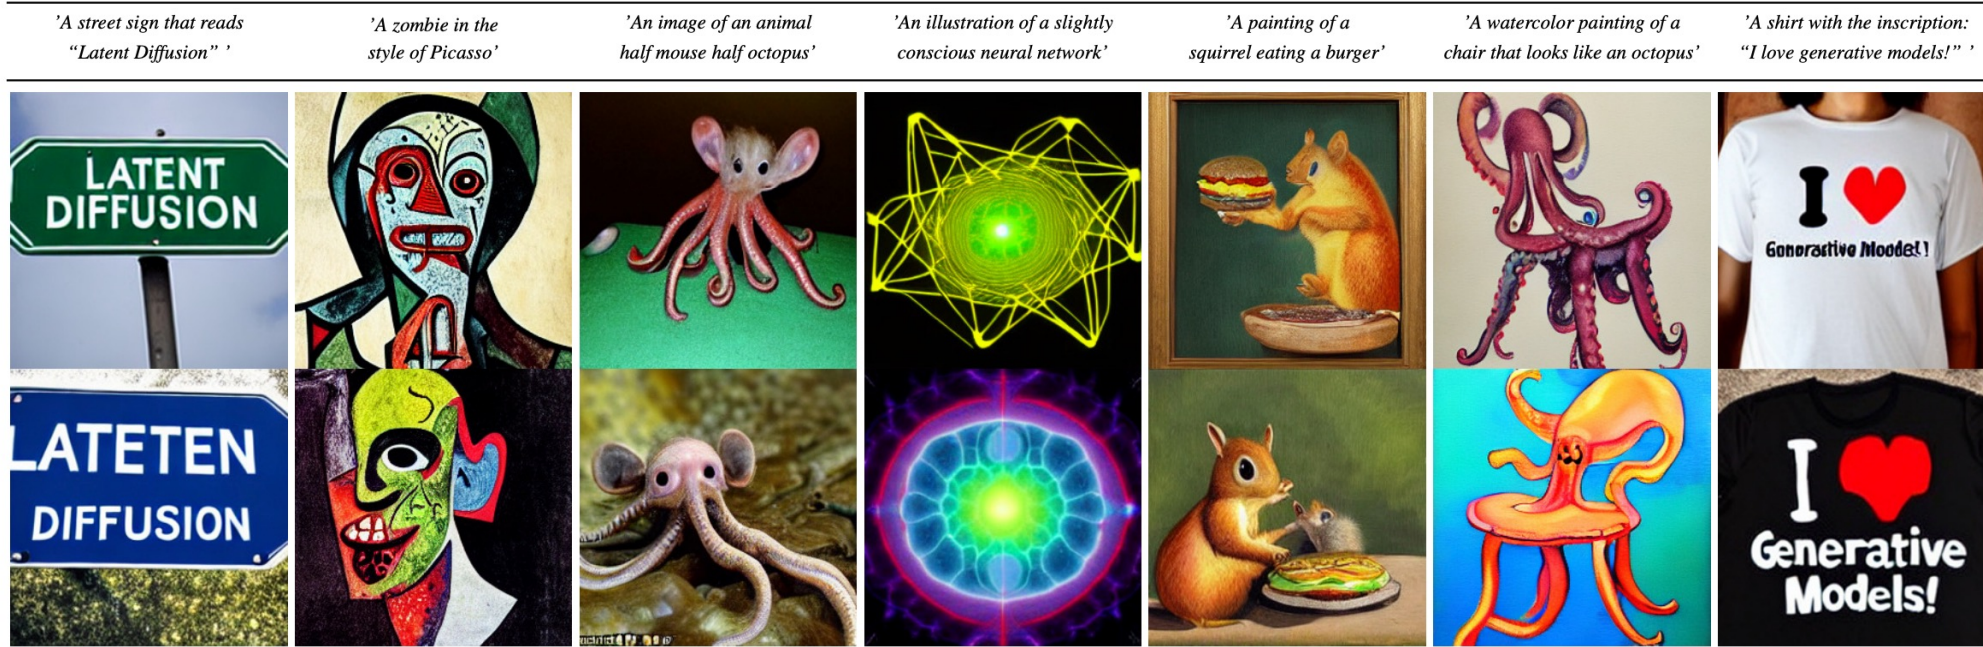


Figure 5. Samples for user-defined text prompts from our model for text-to-image synthesis, *LDM-8 (KL)*, which was trained on the LAION [78] database. Samples generated with 200 DDIM steps and $\eta = 1.0$. We use unconditional guidance [32] with $s = 10.0$.