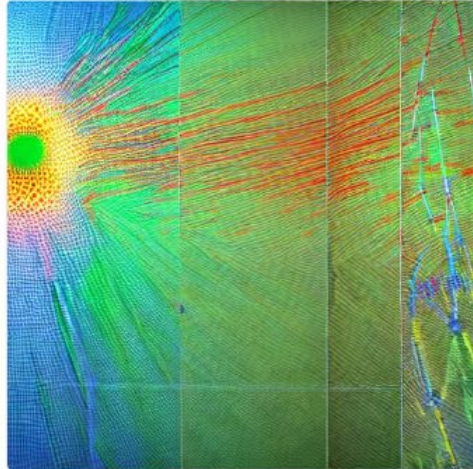




New Generative AI models

(a big thank to Sotiris
Anagnostidis for the presentation
skeleton)

A neural network model used for generating art

Generate image



Overview

- Fundamentals
- GANs
- Autoregressive models
- Diffusion models

What is a generative model?

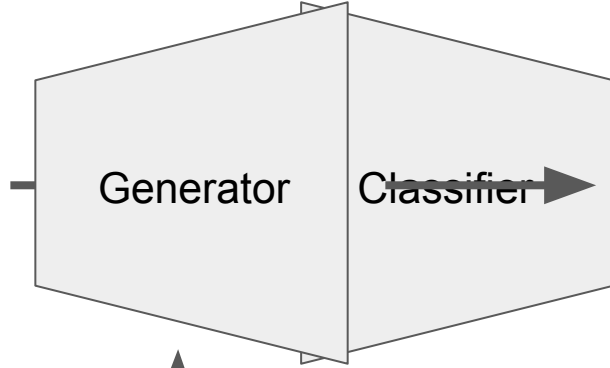

- An algorithm that generates data
- A statistical model of the joint distribution of some data $p(x, y, \dots)$

S

What is a generative model?



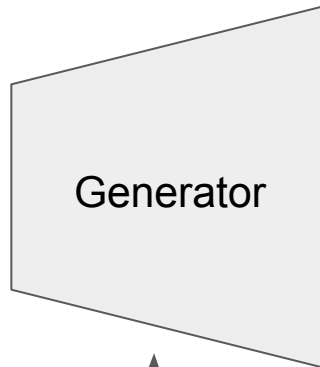
A generative model is a type of machine learning model that is capable of generating new examples that are similar to a training dataset. Generative models are designed to learn the underlying distribution of a dataset and then use this knowledge to generate new examples that belong to the same distribution. These models are typically used in tasks such as image generation, text generation, and audio generation.



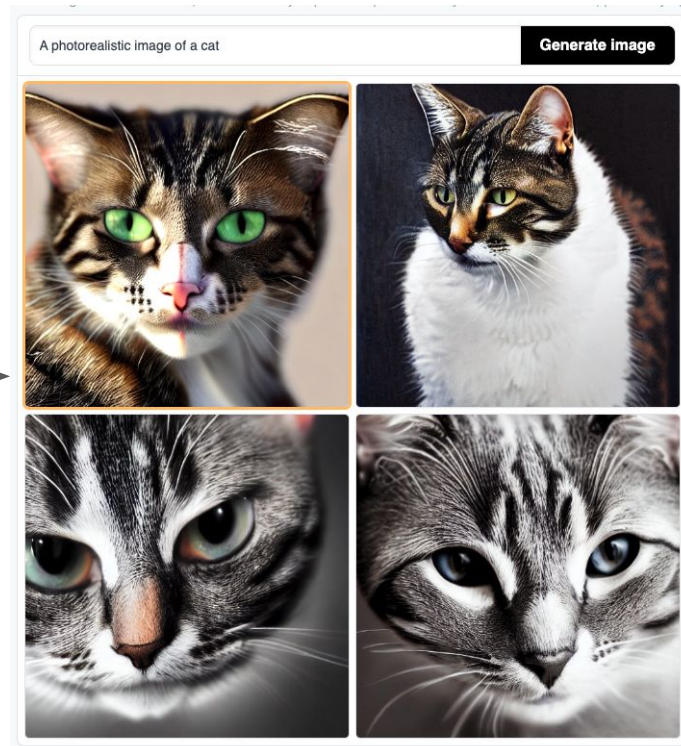
$z \sim N(0, 1)$



Cat




$z \sim N(0, 1)$




What is the goal of generative modeling?



- Make synthetic data that “looks like” real data
- How to measure “looks like”?
- Has high probability under a density model fit to real data (true density of the data).
- Goal is not to replicate training data (failure mode) but to make new data.

 What is the goal of generative neural network?




The goal of a generative neural network is to learn a model of the distribution of a dataset and use that learned distribution to generate new samples that are similar to the ones in the original dataset. This can be useful for a variety of tasks, such as creating new images that are similar to a dataset of images, or generating new text that is similar to a dataset of text. Generative neural networks are a type of unsupervised machine 


Example



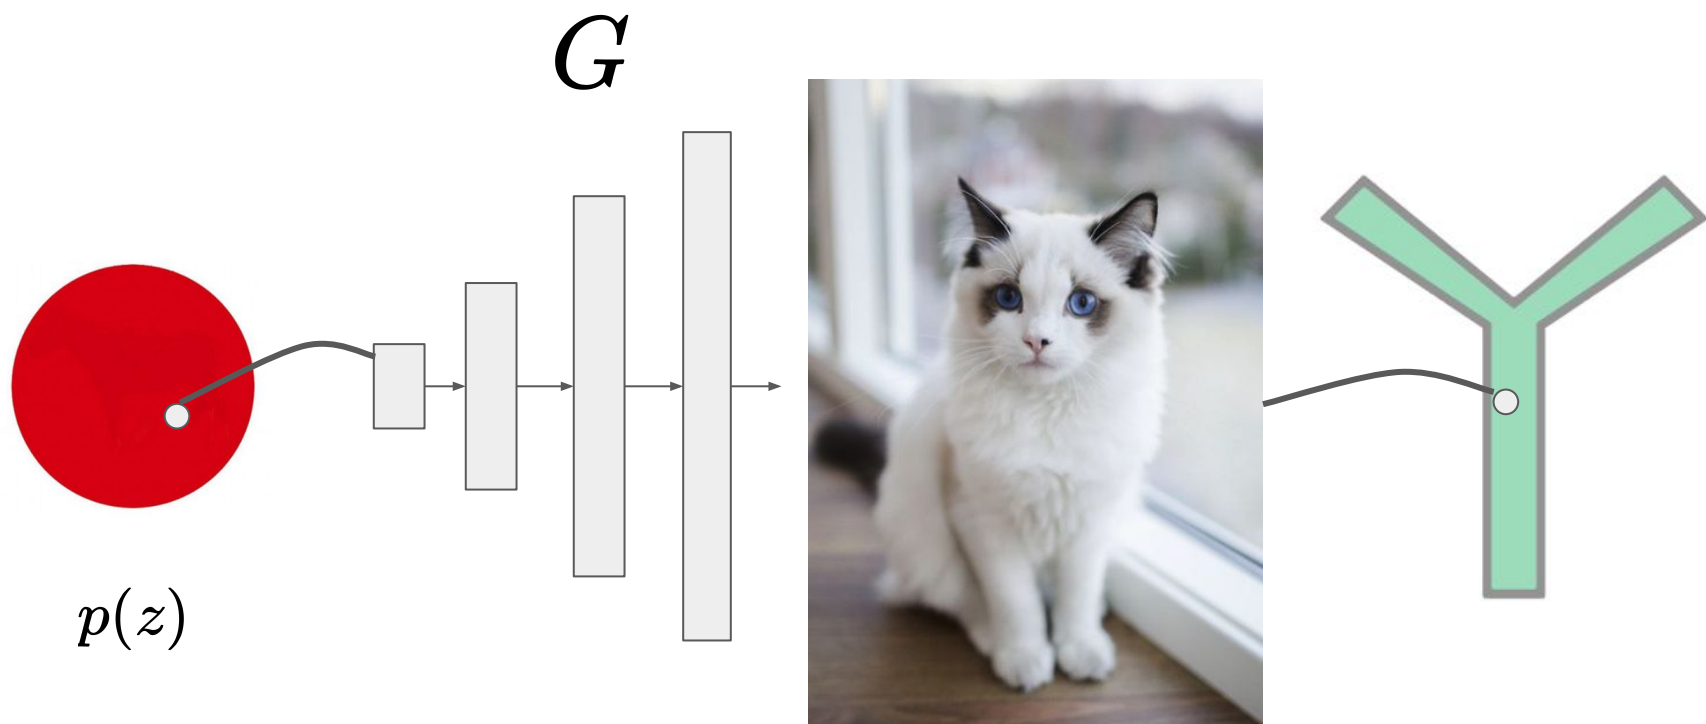
Deep generative models are distribution transformers



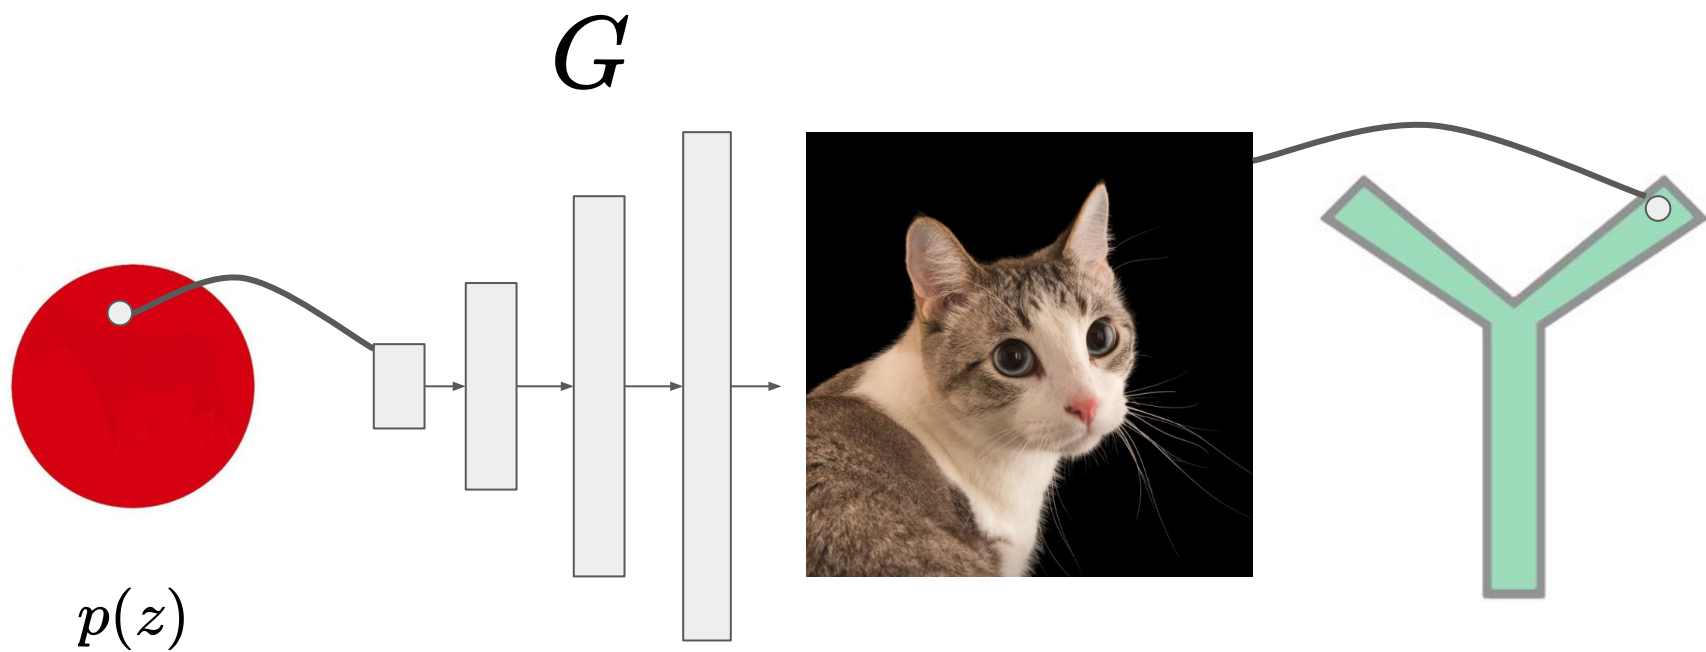
$p(z)$




Deep generative models are distribution transformers



Deep generative models are distribution transformers





Generative Adversarial Networks (GAN)




- G tries to synthesize fake images that fool D
- D tries to identify the fakes

Generative Adversarial Networks (GAN)




$$\operatorname{argmax}_D \mathbb{E}_{z,x} [\log D(G(z)) + \log(1 - D(x))]$$

Generative Adversarial Networks (GAN)



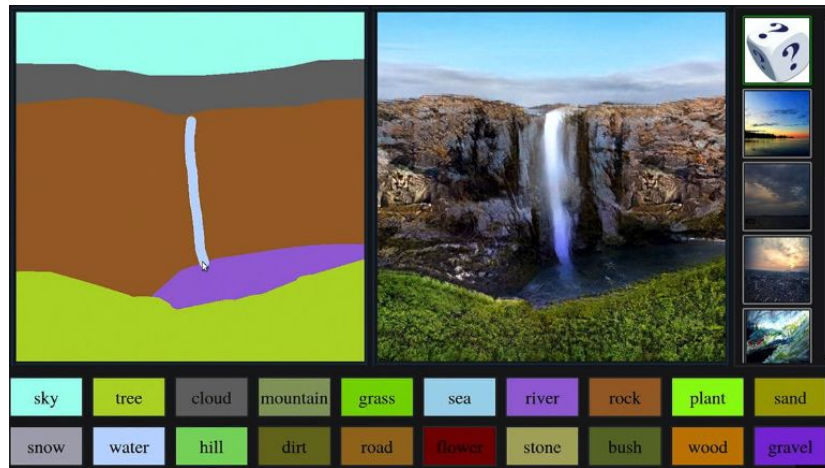
$$\operatorname{argmin}_G \mathbb{E}_{z,x} [\log D(G(z)) + \log(1 - D(x))]$$

Generative Adversarial Networks (GAN)



$$\operatorname{argmin}_G \max_D \mathbb{E}_{z,x} [\log D(G(z)) + \log(1 - D(x))]$$

GANs applications



GANs..again!

Scaling up GANs for Text-to-Image Synthesis


Minguk Kang^{1,3} Jun-Yan Zhu² Richard Zhang³

Jaesik Park¹ Eli Shechtman³ Sylvain Paris³ Taesung Park³

¹POSTECH

²Carnegie Mellon University

³Adobe Research



Autoregressive Models

The cat sits on the _____ Predictor \longrightarrow mat

$$p(X) = p(x_n | x_1, \dots, x_{n-1}) p(x_{n-1} | x_1, \dots, x_{n-2}) \dots p(x_w | x_1) p(x_1)$$

$p(\text{The cat sits on the mat})$

$p(\text{The})$

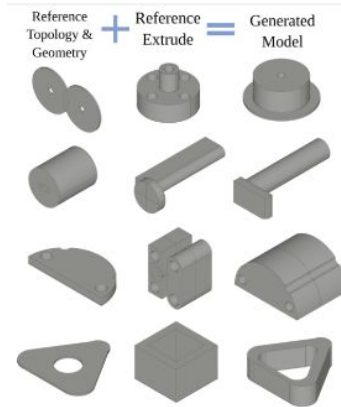
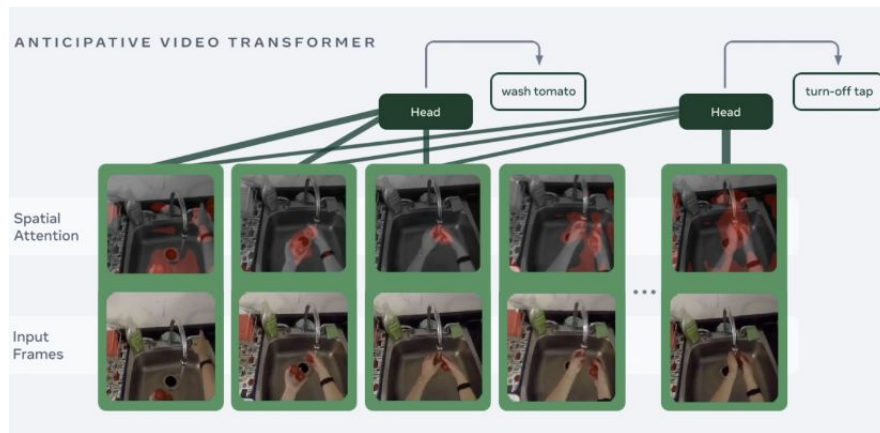
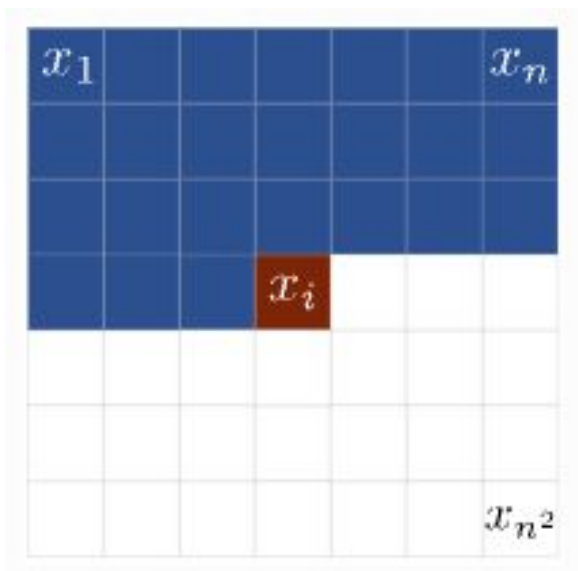

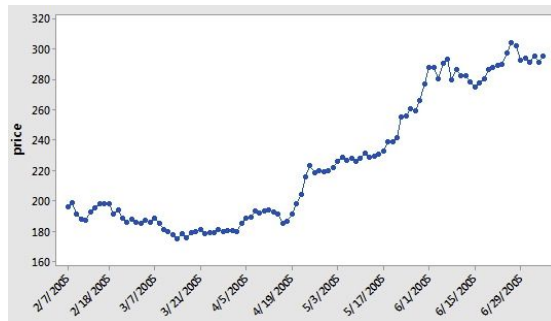
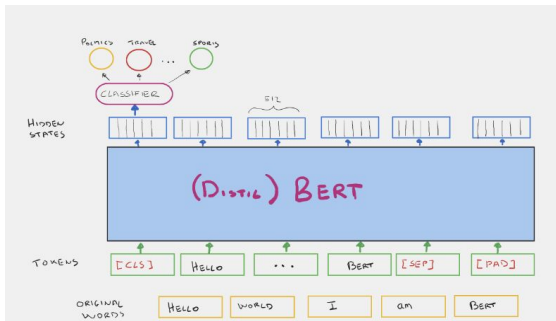
$p(\text{cat} | \text{The})$



$p(\text{sit} | \text{The cat})$

$p(\text{on} | \text{The cat sits})$

$p(\text{the} | \text{The cat sits on })$

$p(\text{mat} | \text{The cat sits on the })$





```

1 function findHighestNumber(array) {
2   var highestNumber = 0;
   for (var i = 0; i < array.length; i++) {
     if (array[i] > highestNumber) {
       highestNumber = array[i];
     }
   }
   return highestNumber;
3 }

```

theorem amc12b_2020_p6


```

(n : ℕ)
(h₀ : 9 ≤ n) :
∃ x : ℕ, (x:ℝ)^2 = (nat.factorial (n + 2) -
  nat.factorial (n + 1)) / nat.factorial n :=
begin
  revert_all,
  intros n hn,
  use n + 1,
  field_simp [nat.factorial_ne_zero, pow_succ'],
  ring_exp
end

```

HatefulMemes (multimodal hateful content classification)

INPUT
Image + OCR



OUTPUT
Hateful Classification

Hateful

VizWiz (visual question answering for the blind)

INPUT
Image + question



Please can you tell me what this item is?

OUTPUT
Answer

Butternut squash red pepper soup


Diffusion Models

Awesome Diffusion Models

Number of papers using diffusion



Yann LeCun liked

 **Zeta Alpha** @ZetaVector · Aug 11

While you wait for Stable Diffusion access... here's text-to-image research *only* from the past 5 months 🤖

Make a Scene: arxiv.org/abs/2203.13131
DALL-E 2: arxiv.org/abs/2204.06125
Imagen: arxiv.org/abs/2205.11487
Cogview 2: arxiv.org/abs/2204.14217
Parti: arxiv.org/abs/2206.10789

1 27 115

How do diffusion models work?

S

How do denoising diffusion probabilistic models work?



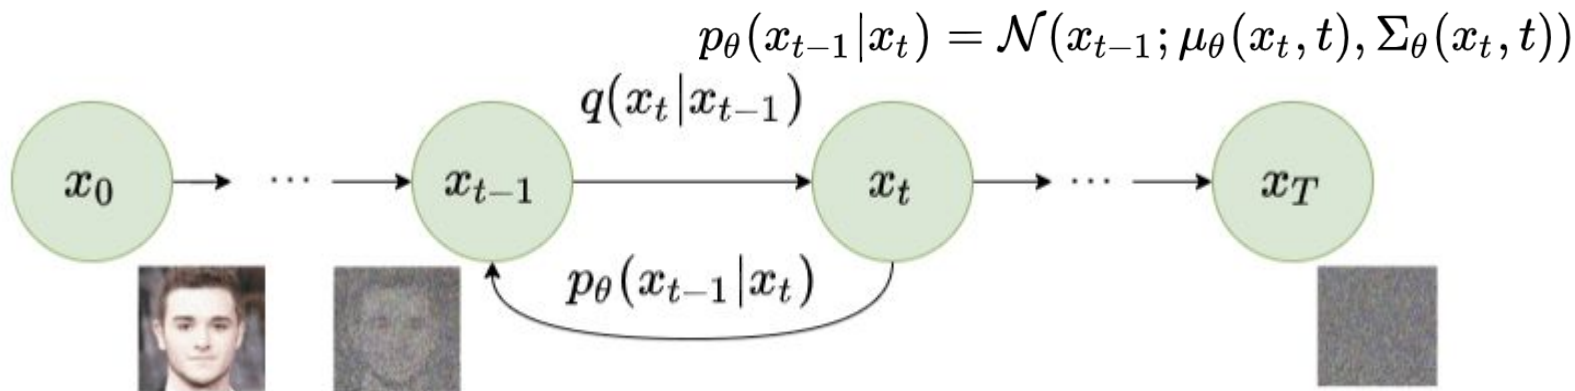

Denoising diffusion probabilistic models are a type of generative model that can be used to generate new data samples that are similar to a training dataset. These models work by using a process called diffusion, which is a way of iteratively refining a set of random noise samples until they resemble the training data.

First, the model is trained on a dataset, and a set of random noise samples are generated. These noise samples are then fed into the model, which uses a series of deterministic and stochastic operations to transform them into samples that are similar to the training data. This process is repeated multiple times, and at each iteration the samples become more and more similar to the training data.

One key aspect of denoising diffusion probabilistic models is that they are designed to be resistant to noise and other forms of corruption in the input data. This allows them to generate high-quality samples even when the input data is noisy or incomplete.

The (vague) details

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



Training DDPM

$$\tilde{a}_t = \prod_{s=0}^t a_s = \prod_{s=0}^t (1 - \beta_s)$$

Forward diffusion: $x_t = \sqrt{1 - \beta_t} x_{t-1} + \sqrt{\beta_t} \epsilon_{t-1} = \sqrt{\tilde{a}_t} x_0 + \sqrt{1 - \tilde{a}_t} \epsilon_0$

Reverse diffusion:

1. Sample from an isotropic Gaussian

x_T



2. $q(x_{t-1} | x_t, x_0) = \mathcal{N}(x_{t-1}; \tilde{\mu}_t(x_t), \tilde{\beta}_t I)$

$$\tilde{\beta}_t = \frac{1 - \tilde{a}_{t-1}}{1 - \tilde{a}_t} \beta_t$$

$$\tilde{\mu}_t(x_t) = \frac{1}{\sqrt{\tilde{a}_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \tilde{a}_t}} \epsilon \right) \rightarrow \epsilon_\theta(x_t, t)$$

3. Train minimizing the Evidence lower bound objective $-\mathbb{E}_{q(x_0)} \log p_\theta(x_0) = \dots$

4. Or simpler...

$$L = \mathbb{E}_{t \sim [1, T], x_0, \epsilon_t} [D(\epsilon_t - \epsilon_\theta(x_t, t))] = \mathbb{E}_{t \sim [1, T], x_0, \epsilon_t} [D(\epsilon_t - \epsilon_\theta(\sqrt{\tilde{a}_t} x_0 + \sqrt{1 - \tilde{a}_t} \epsilon_t, t))]$$

$$D \in L_1, L_2, \dots$$

Training DDPM

Algorithm 1 Training

- 1: **repeat**
 - 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
 - 3: $t \sim \text{Uniform}(\{1, \dots, T\})$
 - 4: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 5: Take gradient descent step on
$$\nabla_{\theta} \left\| \boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\theta}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, t) \right\|^2$$
 - 6: **until** converged
-

Algorithm 2 Sampling

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 - 4: $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$
 - 5: **end for**
 - 6: **return** \mathbf{x}_0
-

Thanks for your attention!

Some References

[CVPR Tutorial on Diffusion Models](#)

[NERF](#)

[Lost of blog posts](#)

[Scaling laws of neural networks](#)