

# Cornell Bowers C-IS

College of Computing and Information Science

# Deep Learning

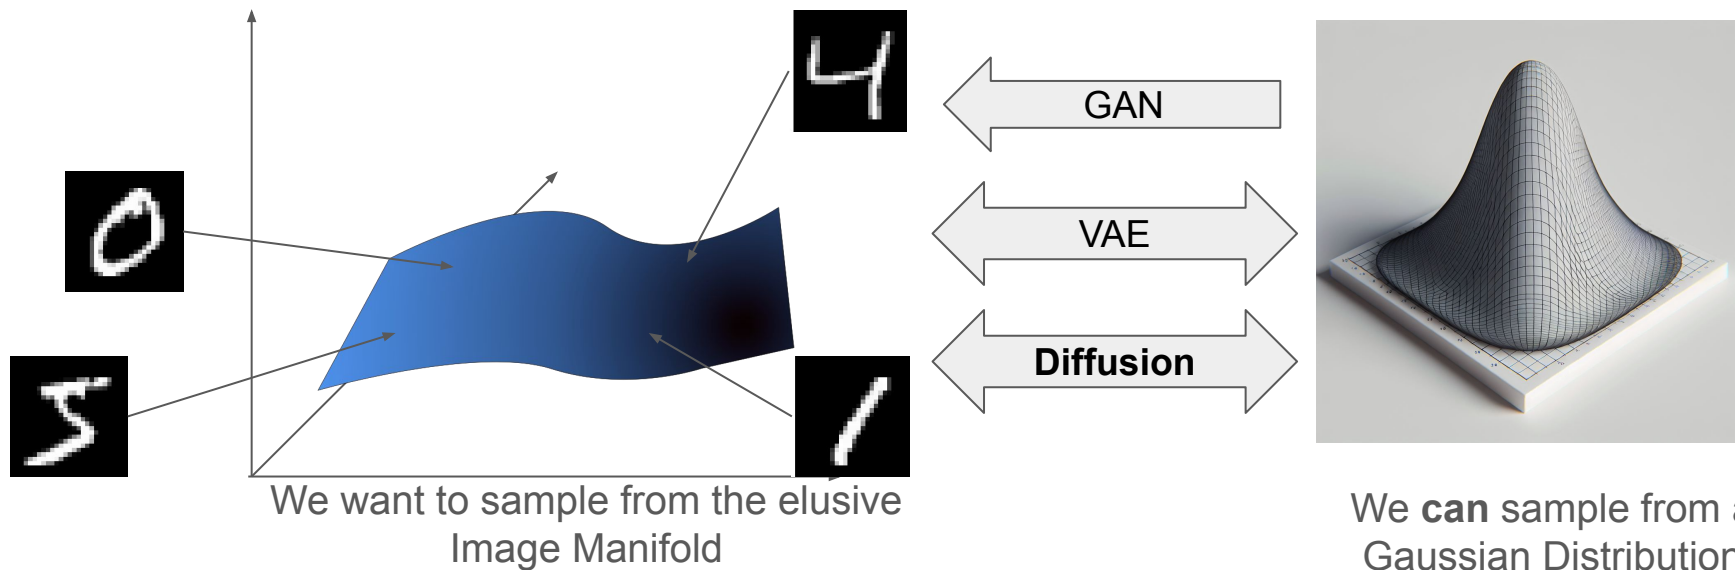
Week 7: Diffusion Models

# Overview

- Recap
- Diffusion model overview
- Forward
- Reverse
- Training Objective


## Recall: Data Manifold

- Data distribution  $\mathbf{P}(\mathbf{X})$  defines a manifold of valid images
- Problem: data manifold takes up **tiny** volume of ambient space
- Naive random samples (e.g. within  $[0,1]^d$ ) are always **off manifold**
- Solution: Sample from a Gaussian, then learn mapping to and from manifold



## Recall: VAE


Back to our AutoEncoder,  
but this time we make  
everything **probabilistic!**



$$\max_{\phi, \theta} \mathbb{E}_{z \sim q_{\phi}(\mathbf{z}|\mathbf{x})} [\log (p_{\theta}(\mathbf{x}|\mathbf{z}))]$$

How likely would it be to encode  $x$ ,  
decode the result, and recover  $x$ ?

# Probabilistic Encoder (Gaussian)





**Problem:** backpropagation through sampling process?

$$\max_{\phi, \theta} \mathbb{E}_{z \sim q_{\phi}(\mathbf{z}|\mathbf{x})} [\log(p_{\theta}(\mathbf{x}|\mathbf{z}))]$$

## Recall: The Reparameterization Trick


$$\mathcal{N}(\mu, \text{diag}(\sigma^2)) = \mu + \sigma \odot \mathcal{N}(0, I)$$



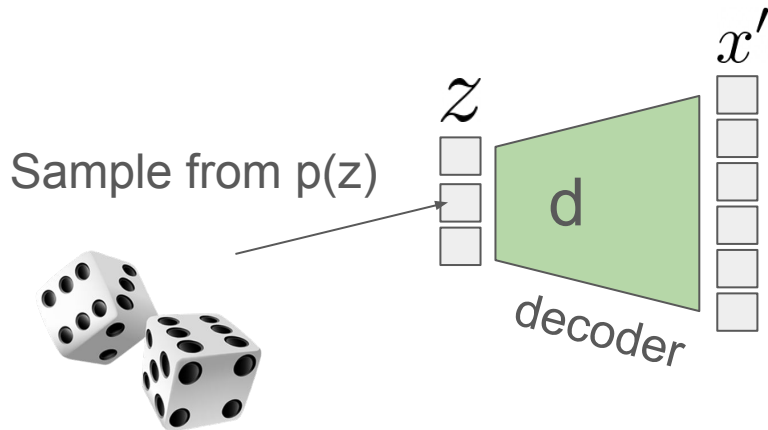
$$=$$


# Recall: How do we sample in latent space?

Solution: Regularize all distributions to be close to the standard normal  $\mathbf{N}(\mathbf{0}; \mathbf{I})$ .



$$\text{maximize} \quad \underbrace{\mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x|z)]}_{\text{reconstruction term}} - \underbrace{D_{\text{KL}}(q_\phi(z|x) \parallel p(z))}_{\text{prior matching term}}$$



# KL Divergence (a.k.a. relative entropy)

$$\begin{aligned}
 \mathbf{D}(p \parallel q) &:= \mathbb{E}_{x \sim p} \left[ \log \frac{p(x)}{q(x)} \right] \\
 &= \mathbb{E}_{x \sim p} \left[ \log \frac{1}{q(x)} \right] - \left[ \log \frac{1}{p(x)} \right] \\
 &\quad \text{Cross Entropy!} \qquad \qquad \qquad \text{( constant wrt } q \text{ )}
 \end{aligned}$$

Distribution 1  $\uparrow$   $p$       Distribution 2  $\uparrow$   $q$

- non-negative  $\mathbf{D}(p \parallel q) \geq 0$
- zero means same  $\mathbf{D}(p \parallel q) = 0 \iff p = q$
- not symmetric
- has many other, uniquely nice properties...




# Recall: Evidence Lower Bound (ELBO)

Data likelihood  $\geq$  Reconstruction – KL Divergence


$$\log p(\mathbf{x}) \geq \underbrace{\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} [\log p_{\theta}(\mathbf{x}|\mathbf{z})]}_{\text{reconstruction term}} - \underbrace{D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}))}_{\text{prior matching term}}$$

(We are **maximizing** this lower bound.)


If we maximize ELBO, we get closer to max to  $P(\mathbf{x})$ .




**GAN:** Adversarial training



**VAE:** maximize variational lower bound



**Diffusion models:**  
Gradually add Gaussian noise and then reverse



# Progress In Generative Modeling

VAEs, 2013



GANs, 2014



PixelCNN, 2016




BigGAN, 2019



Imagen, 2022



# Text-to-Image Diffusion Models



A bread, an apple, and a knife on a table



a robot cooking dinner in the kitchen



A teddy bear and a stuffed raccoon sitting on a wooden chair side by side




The oil painting shows a cow standing near a tree with red leaves



A traditional tea house in a tranquil garden with blooming cherry blossom trees



a painting of trees near a peaceful lake




A heart made of wood



an old man with green eyes and a long grey beard



A painting of an adorable rabbit sitting on a colorful splash



an afrofuturist lady wearing gold jewelry



a black basketball shoe with a lightning bolt on it



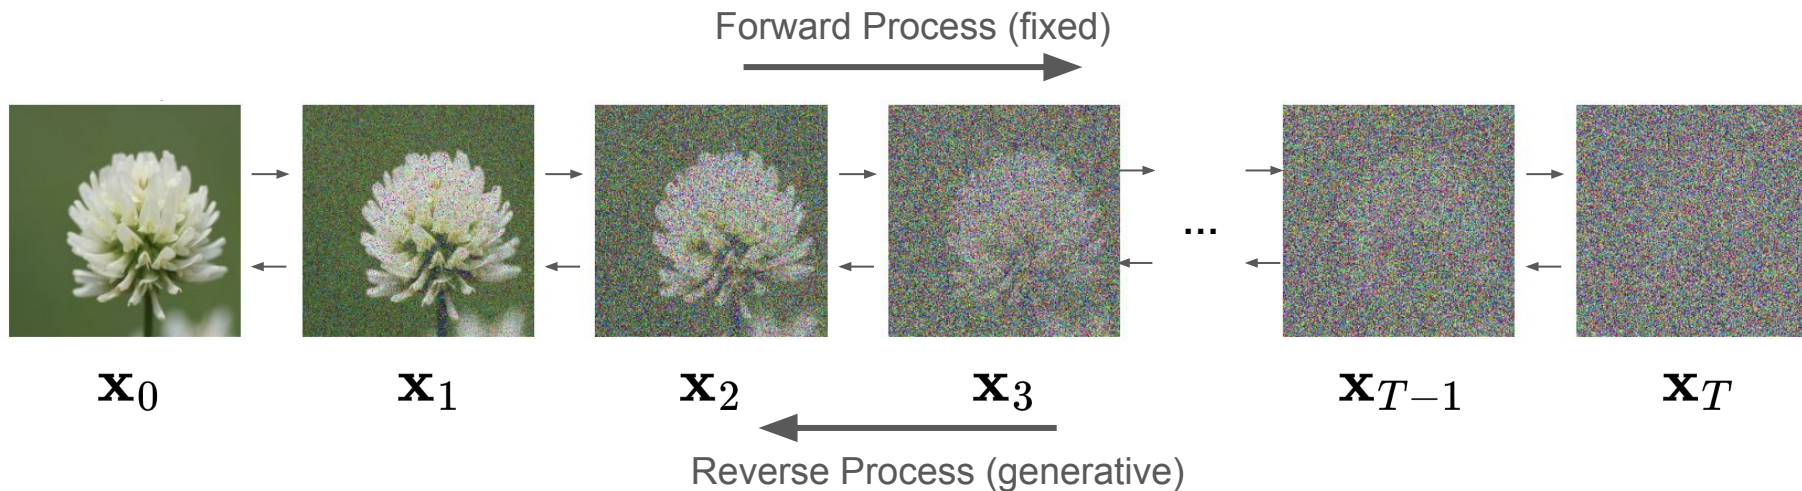
A cool orange cat wearing sunglasses playing a guitar with a group of dancing bananas

# Diffusion Overview

# Denosing Diffusion Models


Denosing diffusion models consist of two processes:

- Forward diffusion process that gradually adds noise to input
- Reverse denosing process that learns to generate data by denoising



# Diffusion Models

We define a mapping to Gaussian noise (forward process)  
Want to **learn the reverse mapping to generate data** (reverse process)







# Markov Chain Implications



$\mathbf{x}_0$



$\mathbf{x}_1$




$\mathbf{x}_2$




$\mathbf{x}_3$

...



$\mathbf{x}_{T-1}$



$\mathbf{x}_T$

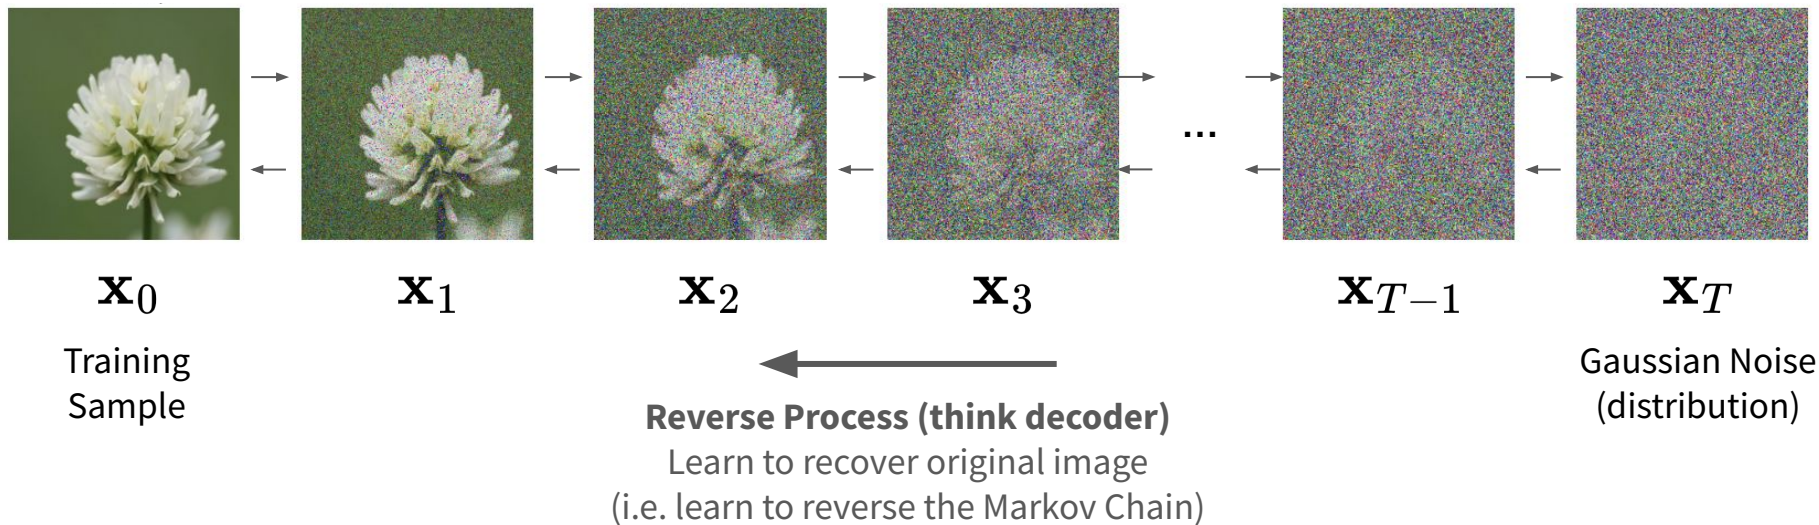


Direction of dependence

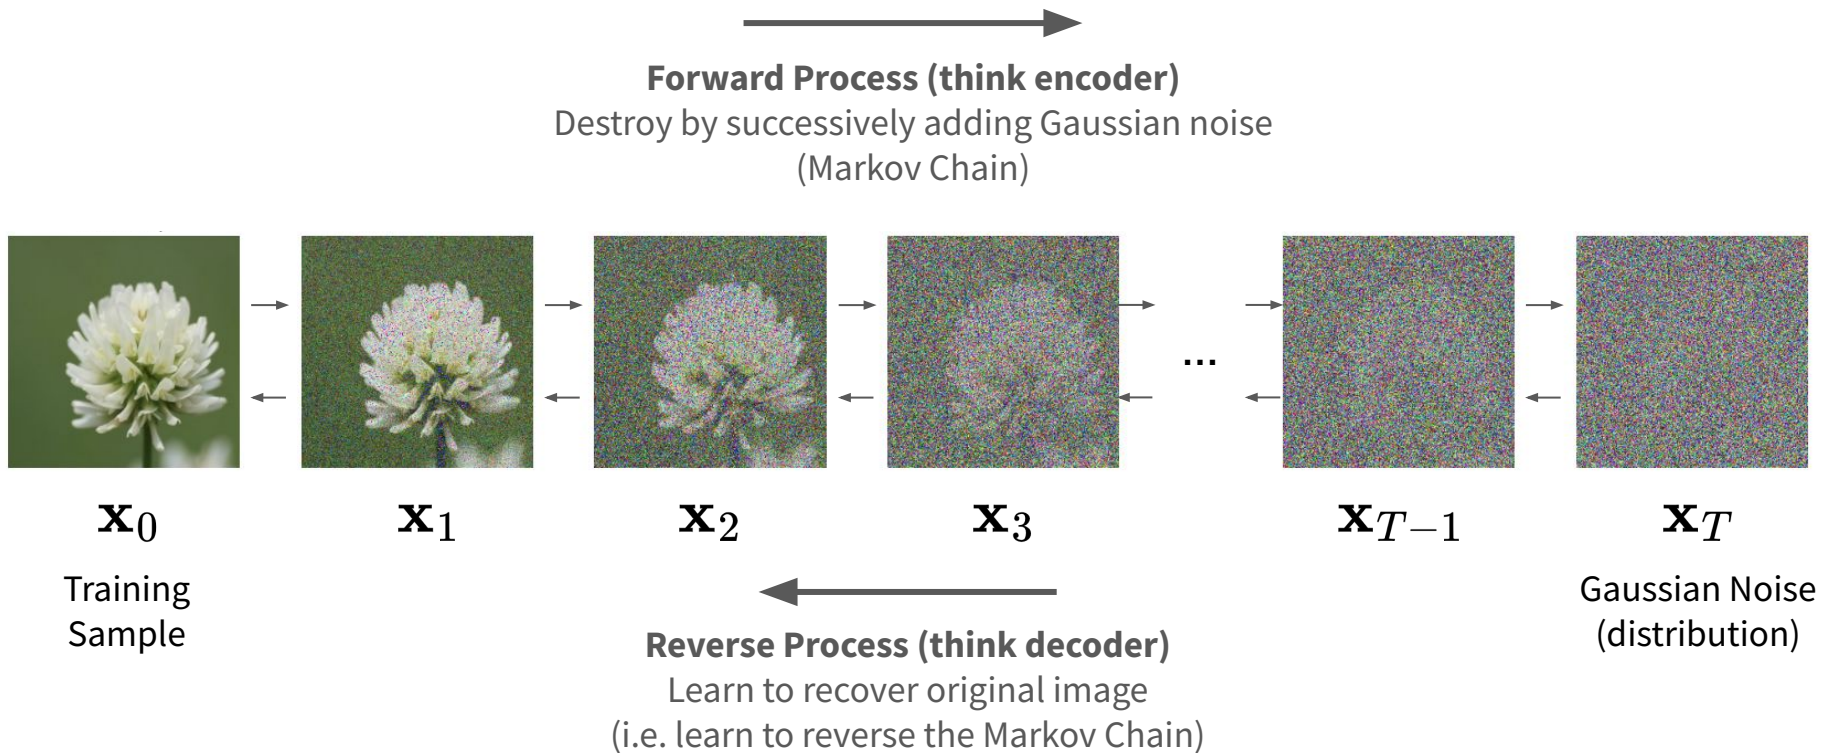
$$q(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0) = q(\mathbf{x}_t | \mathbf{x}_{t-1}) \quad \text{T or F?}$$

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = q(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad \text{T or F?}$$

# Reverse Process: high level idea




# Putting it together



# Diffusion Sampling

Different draws of initial noise lead to diverse outputs




# Forward Process

## Forward Process Overview

- Destroys original image  $\mathbf{x}_0$  by **successively adding Gaussian noise**
- Desired outcome: At step  $T$ ,  $\mathbf{x}_T$  is a **pure Gaussian noise**
  - i.e. the distribution we map the data manifold to

No training yet!!!



## Details: Forward Process

Start from  $\mathbf{x}_0$  sampled from some real-world distribution of images


For timestamps until T:

$\mathbf{x}_t$  sampled from normal distribution conditioned on  $\mathbf{x}_{t-1}$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}) \quad \{\beta_t \in (0, 1)\}_{t=1}^T$$

$$q(\mathbf{x}_T) \approx \mathcal{N}(0, \mathbf{I})$$

noise schedule: how fast we move towards Gaussian noise



## Details: Forward Process

Can we extend this to sampling  $\mathbf{x}_t$  in a closed form?


Let  $\alpha_t := 1 - \beta_t$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{\alpha_t} \mathbf{x}_{t-1}, (1 - \alpha_t) \mathbf{I})$$

**Re-parametrization trick!**


$$\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_{t-1} + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}_{t-1}$$

$$\boldsymbol{\epsilon}_{t-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$





## Details: Forward Process



Inductively, we can say

$$\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_{t-1} + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}_{t-1}$$

$$= \sqrt{\alpha_t \alpha_{t-1}} \mathbf{x}_{t-2} + \sqrt{1 - \alpha_t \alpha_{t-1}} \bar{\boldsymbol{\epsilon}}_{t-2}$$

$$= \dots$$

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

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

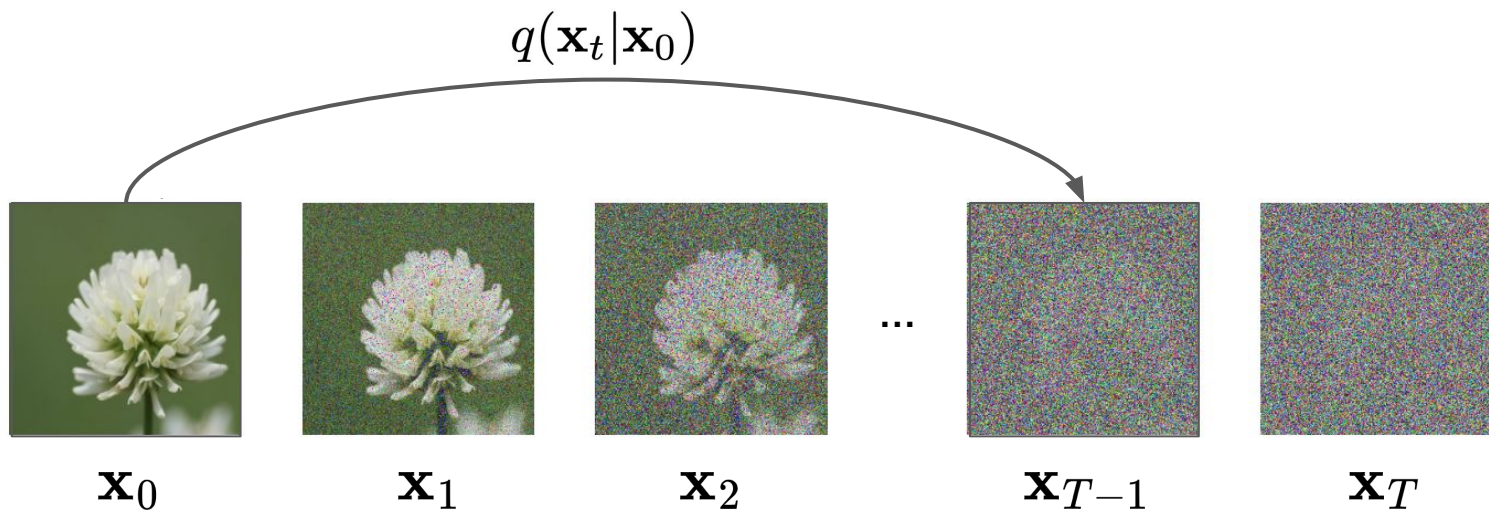
Merged noise.  
epsilon is still  $\sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\boxed{\phantom{\mathbf{x}_0}}, \boxed{\phantom{\mathbf{I}}})$$

## Details: Forward Process

Can sample  $\mathbf{x}_t$  in closed-form as  $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I})$

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \bar{\alpha}_t \in (0, 1)$$



## Aside: Noise Schedules

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \bar{\alpha}_t \in (0, 1)$$

- Define the noise schedule in terms of  $\bar{\alpha}_t \in (0, 1)$ 
  - Some monotonically decreasing function from 1 to 0
- Cosine Noise schedule:

$$\bar{\alpha}_t = \cos(.5\pi t/T)^2$$

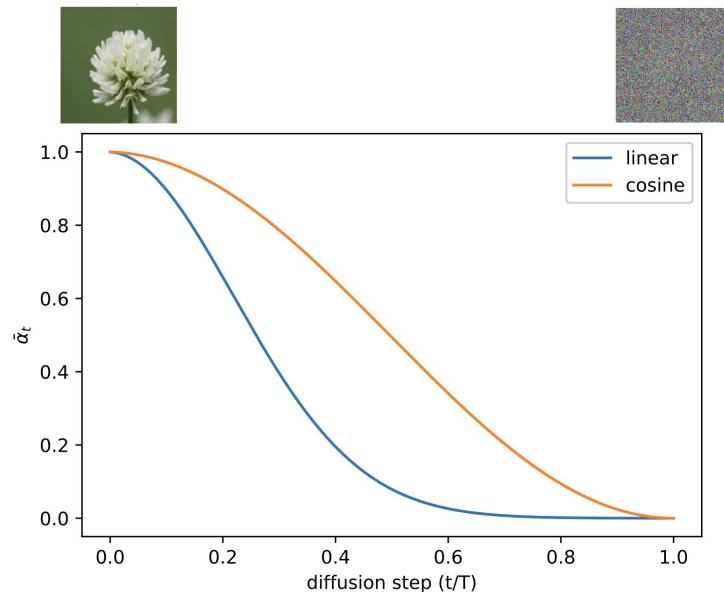



Figure 5.  $\bar{\alpha}_t$  throughout diffusion in the linear schedule and our proposed cosine schedule.

# Reverse Process



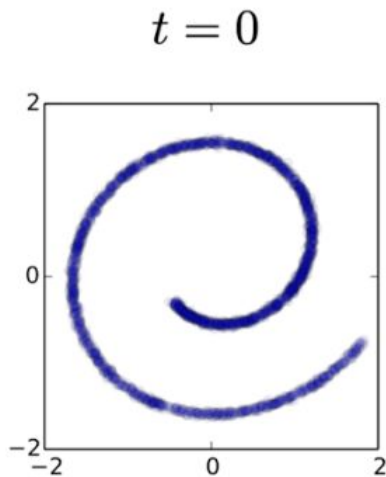
## Reverse Process Overview

- "Learn to reverse what we just destroyed"
  - Learn time reversal of Markov Chain; we **train a model for this**
- Desired outcome: some  $\mathbf{x}_0$  close to the original data distribution



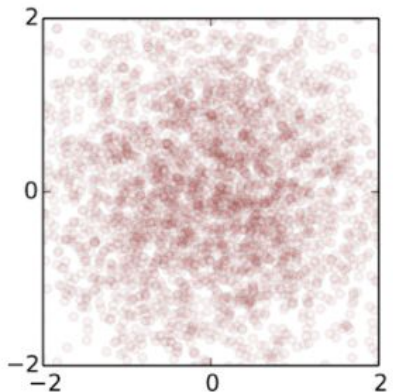


The forward trajectory

$$q(\mathbf{x}_{0:T})$$



The reverse trajectory

$$p_{\theta}(\mathbf{x}_{0:T})$$




## Details: Reverse Process

Start from  $q(\mathbf{x}_T) = \mathcal{N}(0, \mathbf{I})$

Ideally, sample from reversed conditional distribution  $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$


How to compute  $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$  ?



## Details: Reverse Process

How to compute  $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$  ?

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





$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$


$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I})$$

## Details: Reverse Process

$q(\mathbf{x}_{t-1} | \mathbf{x}_t)$  Is **not tractable**. Is  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$  tractable?

I.e. Can we reverse the forward process given the original data?

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \frac{q(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0) q(\mathbf{x}_{t-1} | \mathbf{x}_0)}{q(\mathbf{x}_t | \mathbf{x}_0)}$$



$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I})$$


## Details: Reverse Process

$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$  is **tractable**

Can reverse the forward process given the original data!

Problem: Don't have any "original data  $\mathbf{x}_0$ " during **inference**

We have  $\mathbf{x}_0$  during **training**; train a **generative model**



## Key Idea


We introduce a generative model to approximate the reverse process  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$

$$p(\mathbf{x}_T) = \mathcal{N}(\mathbf{0}, \mathbf{I})$$


$$p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mu_{\theta}(\mathbf{x}_t, t), \sigma_t^2 \mathbf{I})$$

$$\mathbb{E}_{q(\mathbf{x}_t | \mathbf{x}_0)} [D_{\text{KL}}(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) \| p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t))]$$


Learning Objective!




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



$\mathbf{x}_T$



$\mathbf{x}_{T-1}$



$\mathbf{x}_2$



$\mathbf{x}_1$



$\mathbf{x}_0$

# Diffusion Training Objective

# Training

Find the model that **maximizes the likelihood** of the training data




i.e. same as VAEs, variational inference; approximate the true posterior

$$\mathbf{max} \log p(\mathbf{x})$$


# Training Objective

- Bound the likelihood with the ELBO
  - Exactly like VAEs

$$\log p(\mathbf{x}) \geq \underbrace{\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})]}_{\text{reconstruction term}} - \underbrace{D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}))}_{\text{prior matching term}}$$



$$\log p(\mathbf{x}) \geq \underbrace{\mathbb{E}_{q(\mathbf{x}_1|\mathbf{x}_0)} [\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)]}_{\text{reconstruction term}} - \underbrace{D_{\text{KL}}(q(\mathbf{x}_T|\mathbf{x}_0) \parallel p(\mathbf{x}_T))}_{\text{prior matching term}} - \sum_{t=2}^T \underbrace{\mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)} [D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \parallel p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))]}_{\text{denoising matching term}}$$



# Parameterizing the Denoising Model

$$-\sum_{t=2}^T \underbrace{\mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)} [D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \parallel p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))]}_{\text{denoising matching term}}$$

$$: \mathbb{E}_q \left[ \frac{1}{2\sigma_t^2} \|\tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) - \mu_\theta(\mathbf{x}_t, t)\|^2 \right] + C$$

$$\tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) = \frac{1}{\sqrt{1-\beta_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon \right)$$

$$\mu_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{1-\beta_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(\mathbf{x}_t, t) \right)$$

$$L(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2]$$

Simplifying KL Divergence to MSE of means, as distributions are Gaussians with same variance!

Simplifying Bayes Rule...

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0)q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)}$$


Re-parametrize  $\mu_\theta(\mathbf{x}_t, t)$

Loss is MSE of actual to predicted loss!

# What Network Architecture to Use For $\epsilon_{\theta}$ ?

People often use U-Nets with residual blocks and self-attention layers at low resolutions

Has same input and output image dimensions



Time representation: sinusoidal positional embeddings

Inject time embedding throughout the network (e.g. additive positional embedding)



$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

## Training Algorithm

Repeat until convergence

1.  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$  ← Sample original image from image distribution
2.  $t \sim U\{1, 2, \dots, T\}$  ← Sample random time step uniformly
3.  $\epsilon \sim \mathcal{N}(0, 1)$  ← Sample Gaussian noise
4. Optimizer step on  $L(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2]$   
← Model predicts noise applied at time step  $t$  and calculate loss

## Sampling Algorithm

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$       ← Sample pure Gaussian noise

For  $t = T, T - 1 \dots, 1$

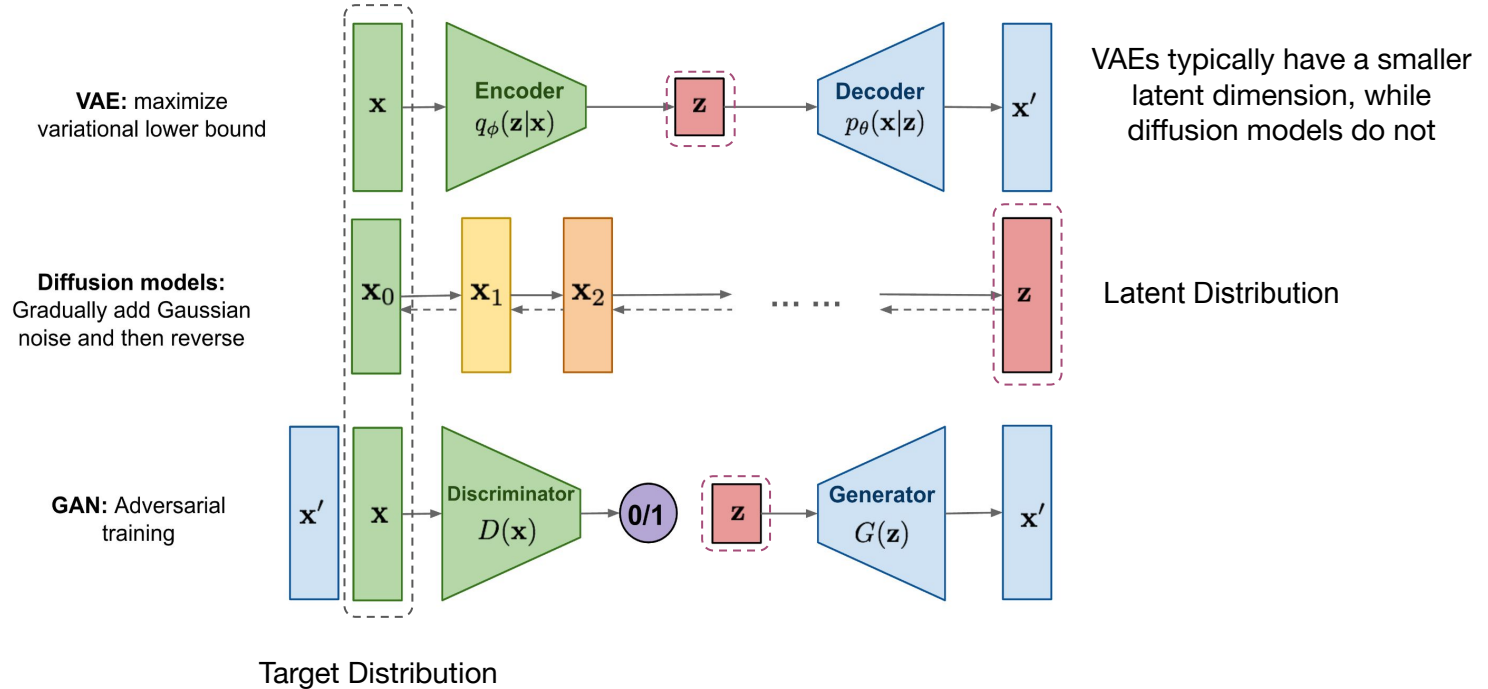
$\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  if  $t > 1$ , else  $\mathbf{z} = \mathbf{0}$       ← Sample Gaussian noise to apply to image

$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$       ← Predict noise applied to image and remove that noise

Return  $\mathbf{x}_0$

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = q(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t))$$

# Generative Modeling



# Recap

- Can bound the likelihood of observed data (i.e. the evidence) with the Evidence Lower Bound (i.e. the ELBO)
- Can learn generative models by maximizing the ELBO
  - VAEs, hierarchical VAEs, Diffusion models
- Learning objective decomposed to each timestep
  - Can be made extremely deep!
  - Can focus on higher noise levels to improve perceptual quality!
- Limitation:
  - Can require many sampling steps for good quality