



10-423/10-623 Generative AI

Machine Learning Department
School of Computer Science
Carnegie Mellon University

Homework 2 Recitation

Diffusion Models


Variational Inference

Feb. 12, 2024

Agenda

1. Overview of diffusion model
2. Diffusion model math
3. HW2 starter code overview
4. Overview of Fréchet Inception Distance (FID)
5. Helpful functions & practice reading documentation

Diffusion Model



Forward Process:

$$q_\phi(\mathbf{x}_{1:T}) = q(\mathbf{x}_0) \prod_{t=1}^T q_\phi(\mathbf{x}_t | \mathbf{x}_{t-1})$$

(Learned) Reverse Process:


$$p_\theta(\mathbf{x}_{1:T}) = p_\theta(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$$

(Exact) Reverse Process:

$$q_\phi(\mathbf{x}_{1:T}) = q_\phi(\mathbf{x}_T) \prod_{t=1}^T q_\phi(\mathbf{x}_{t-1} | \mathbf{x}_t)$$

The exact reverse process requires inference. And, even though $q_\phi(\mathbf{x}_t | \mathbf{x}_{t-1})$ is simple, computing $q_\phi(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is intractable! Why? Because $q(\mathbf{x}_0)$ might be not-so-simple.

Diffusion Model



adds noise to the image

Forward Process:

$$q_\phi(\mathbf{x}_{1:T}) = q(\mathbf{x}_0) \prod_{t=1}^T q_\phi(\mathbf{x}_t | \mathbf{x}_{t-1})$$

if we could sample from this we'd be done

(Learned) Reverse Process:

removes noise

$$p_\theta(\mathbf{x}_{1:T}) = p_\theta(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$$


goal is to learn this

(Exact) Reverse Process:


$$q_\phi(\mathbf{x}_{1:T}) = q_\phi(\mathbf{x}_T) \prod_{t=1}^T q_\phi(\mathbf{x}_{t-1} | \mathbf{x}_t)$$

The exact reverse process requires inference. And, even though $q_\phi(\mathbf{x}_t | \mathbf{x}_{t-1})$ is simple, computing $q_\phi(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is intractable! Why? Because $q(\mathbf{x}_0)$ might be not-so-simple.


Diffusion Model



How does this actually work?




Generating an image from noise in a single step




Generating an image from noise in 500 steps

Denoising is not image recovery



Denoising is not image recovery



Denoising is not image recovery



Denoising Diffusion Overview

From magic to math



Denoising Diffusion Overview


An alternative perspective

Images as vectors



Color Image
(RGB)

32 by 32 pixel image



32 x 32 x 3 tensor


$$\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_m \end{bmatrix} \cdot$$

3072 dimensional
vector


Denoising Diffusion Overview

An alternative perspective

Images as vectors



Representation of 3D space with 3 axis




In image is a vector in 3072 dimensional space

Denoising Diffusion Overview

An alternative perspective

Manifold hypothesis



If we can find the correct representation of images in high dimensional space, all 32x32 color images of cats will occupy a latent manifold

Latent Variable Models

- For GANs, we assume that there are (unknown) **latent variables** which give rise to our observations
- The **noise vector z** are those latent variables
- After learning a GAN, we can **interpolate** between images in latent z space

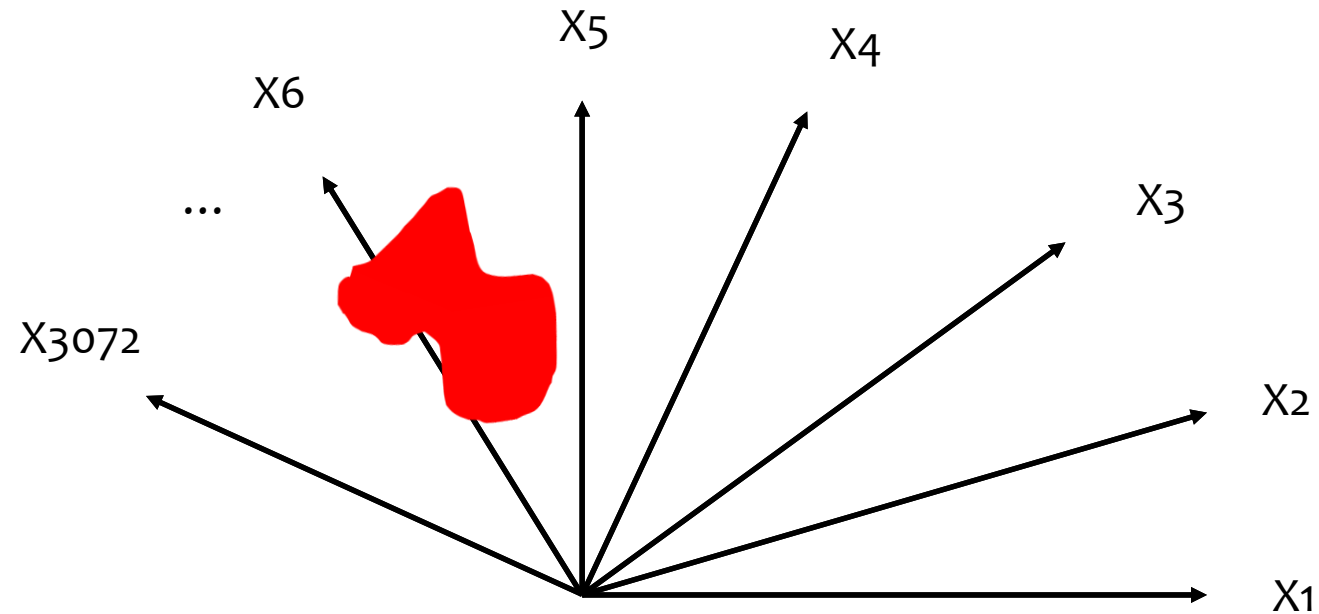


Figure 4: Top rows: Interpolation between a series of 9 random points in Z show that the space learned has smooth transitions, with every image in the space plausibly looking like a bedroom. In the 6th row, you see a room without a window slowly transforming into a room with a giant window. In the 10th row, you see what appears to be a TV slowly being transformed into a window.

Denoising Diffusion Overview

An alternative perspective

Manifold hypothesis



If we could just sample this latent manifold, we would be able to generate any cat picture that could ever exist

But how to sample it?

Denoising Diffusion Overview

An alternative perspective

Diffusion as sampling

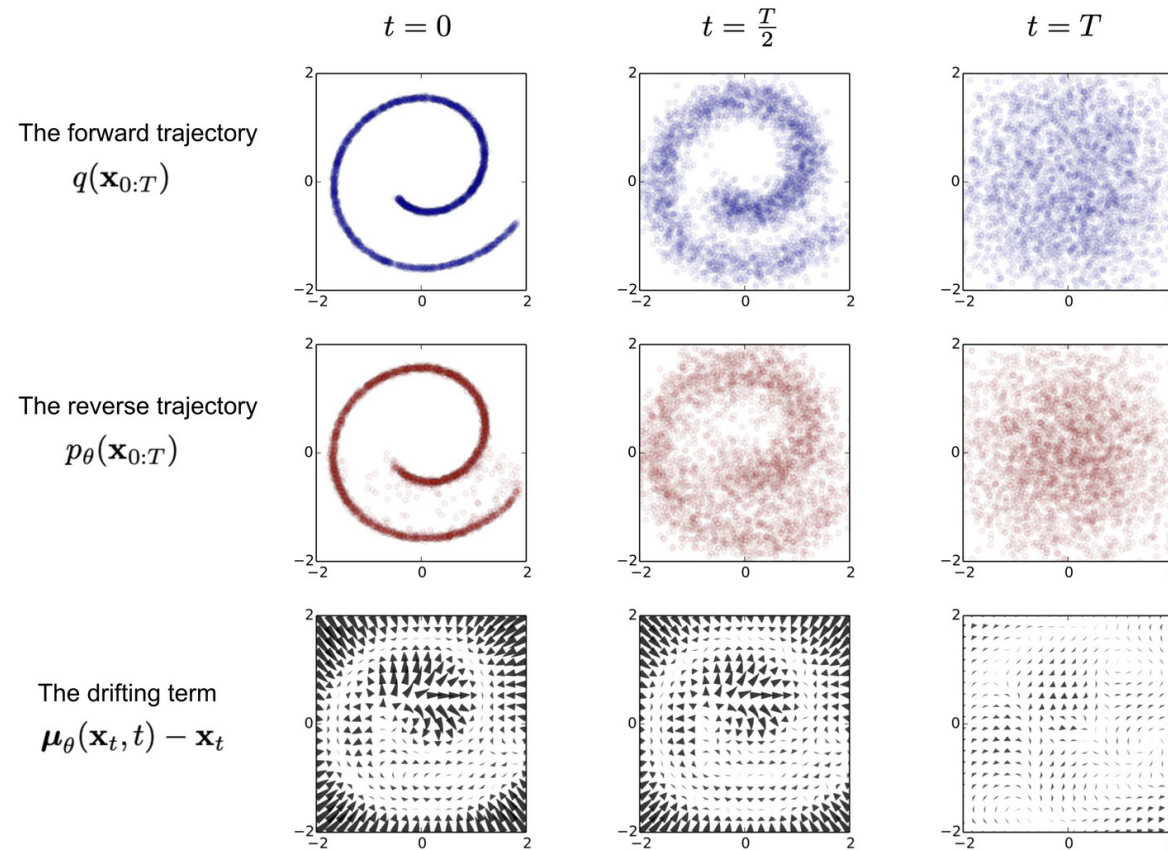
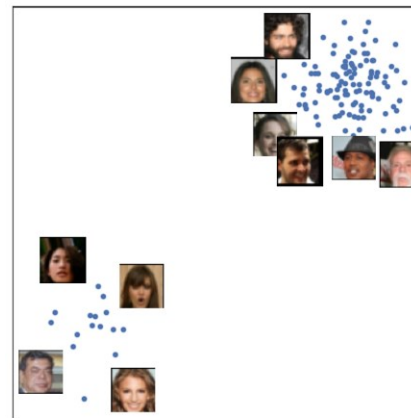



Fig. 3. An example of training a diffusion model for modeling a 2D swiss roll data. (Image source: [Sohl-Dickstein et al., 2015](#))

Denoising Diffusion Overview

An alternative perspective


Sampling the latent manifold using diffusion



Data samples

$$\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \stackrel{\text{i.i.d.}}{\sim} p(\mathbf{x})$$

score matching



Scores

$$\mathbf{s}_\theta(\mathbf{x}) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x})$$

Langevin dynamics



New samples

Score matching with Langevin Dynamics – see Section 3.2 of DDPM paper for more details
arxiv.org/pdf/2006.11239.pdf

Bayes' Rule (Theorem)

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$


$$P(A) = \frac{\text{orange hexagon}}{\text{gray square}}, \quad P(B|A) = \frac{\text{blue diamond}}{\text{orange hexagon}}$$

$$P(B) = \frac{\text{blue circle}}{\text{gray square}}, \quad P(A|B) = \frac{\text{blue diamond}}{\text{blue circle}}$$

$$P(A) \cdot P(B|A) = \frac{\text{orange hexagon with pink slash}}{\text{gray square}} \times \frac{\text{blue diamond}}{\text{orange hexagon with pink slash}} = \frac{\text{blue diamond}}{\text{gray square}}$$

$$P(B) \cdot P(A|B) = \frac{\text{blue circle with pink slash}}{\text{gray square}} \times \frac{\text{blue diamond}}{\text{blue circle with pink slash}} = \frac{\text{blue diamond}}{\text{gray square}}$$

= $P(A) \cdot P(B|A)$, i.e.



$$P(A|B) = \frac{P(A) \cdot P(B|A)}{P(B)}$$

$$P(B|A) = \frac{P(B) \cdot P(A|B)}{P(A)}$$

Evidence Lower Bound (ELBO)

Jensen's Inequality

If function $f(\cdot)$ is concave, then:

$$\mathbb{E}[f(X)] \leq f(\mathbb{E}[X])$$

Evidence Lower Bound (ELBO)

ELBO via Jensen's Inequality

$$\log p(\mathbf{x}) = \log \int p(\mathbf{x}, \mathbf{z}) d\mathbf{z}$$

**We use marginalization
To make latent variable
appear**

Evidence Lower Bound (ELBO)

ELBO via Jensen's Inequality

$$\begin{aligned}\log p(\mathbf{x}) &= \log \int p(\mathbf{x}, \mathbf{z}) d\mathbf{z} \\ &= \log \int q(\mathbf{z}) \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z}\end{aligned}$$

We need this term for the expectation of Jensen's

Evidence Lower BOund (ELBO)

ELBO via Jensen's Inequality

$$\begin{aligned}\log p(\mathbf{x}) &= \log \int p(\mathbf{x}, \mathbf{z}) d\mathbf{z} \\ &= \log \int q(\mathbf{z}) \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z}\end{aligned}$$

We cancel out to preserve
The equality

Evidence Lower Bound (ELBO)

ELBO via Jensen's Inequality

$$\begin{aligned}\log p(\mathbf{x}) &= \log \int p(\mathbf{x}, \mathbf{z}) d\mathbf{z} \\ &= \log \int q(\mathbf{z}) \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z}\end{aligned}$$

It is a concave function!

Evidence Lower BOund (ELBO)

ELBO via Jensen's Inequality

$$\begin{aligned}\log p(\mathbf{x}) &= \log \int p(\mathbf{x}, \mathbf{z}) d\mathbf{z} \\ &= \log \int q(\mathbf{z}) \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z} \\ &\geq \int q(\mathbf{z}) \log \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z}\end{aligned}$$

We use Jensen's to swap log and expectation

Evidence Lower BOund (ELBO)

ELBO

$$\log p(\mathbf{x}) \geq \int q(\mathbf{z}) \log \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z}$$

ELBO is your best friend

Evidence Lower BOund (ELBO)

ELBO

$$\begin{aligned}\log p(\mathbf{x}) &\geq \int q(\mathbf{z}) \log \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} d\mathbf{z} \\ &= \mathbb{E}_q \left[\log \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z})} \right]\end{aligned}$$



Maximize this

ELBO is your best friend

The Reparameterization Trick

Reparameterization is a method of generating random numbers by transforming some base distribution $p(\epsilon)$ to a desired distribution $p_{\theta}(z)$

$$\epsilon \sim p(\epsilon) \longrightarrow g(\epsilon; \theta) \longrightarrow p_{\theta}(z)$$

The Reparameterization Trick

Reparameterization is a method of generating random numbers by transforming some base distribution $p(\epsilon)$ to a desired distribution $p_{\theta}(z)$

$$\epsilon \sim p(\epsilon) \longrightarrow g(\epsilon; \theta) \longrightarrow p_{\theta}(z)$$

A simple distribution to sample from

The Reparameterization Trick

Reparameterization is a method of generating random numbers by transforming some base distribution $p(\epsilon)$ to a desired distribution $p_{\theta}(z)$

$$\epsilon \sim p(\epsilon) \longrightarrow g(\epsilon; \theta) \longrightarrow p_{\theta}(z)$$

A simple transformation

The Reparameterization Trick

Gaussian Distribution:

We want samples from $x \sim \mathcal{N}(\mu, \sigma^2 \mathbf{I})$

The Reparameterization Trick

Gaussian Distribution

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

We sample standard Normal

The Reparameterization Trick

Gaussian Distribution

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

$$\mathbf{x} = \boldsymbol{\mu} + \sigma \epsilon$$

We apply linear transformation

The Reparameterization Trick

Gaussian Distribution


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

$$\mathbf{x} = \boldsymbol{\mu} + \sigma \epsilon$$

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I})$$

The transformed sample comes from the desired Gaussian distribution

Data Scaling



Starter Code Walkthrough

How do we implement a diffusion model?

We have 5 ingredients

- U-Net Model
- Trainer code
- Noise scheduler
- Training implementation
- Sampling implementation

```
✓ handout
  > __pycache__
  > data
  📄 diffusion.py
  📄 main.py
  ≡ requirements.txt
  📄 run_in_colab.ipynb
  📄 trainer.py
  📄 unet.py
  > latex_template
  📄 Homework-2.pdf
```

Starter Code Walkthrough

U-Net

- U-Net's role here is to model the denoising function at each step of the reverse diffusion process.
- Multi-scale features

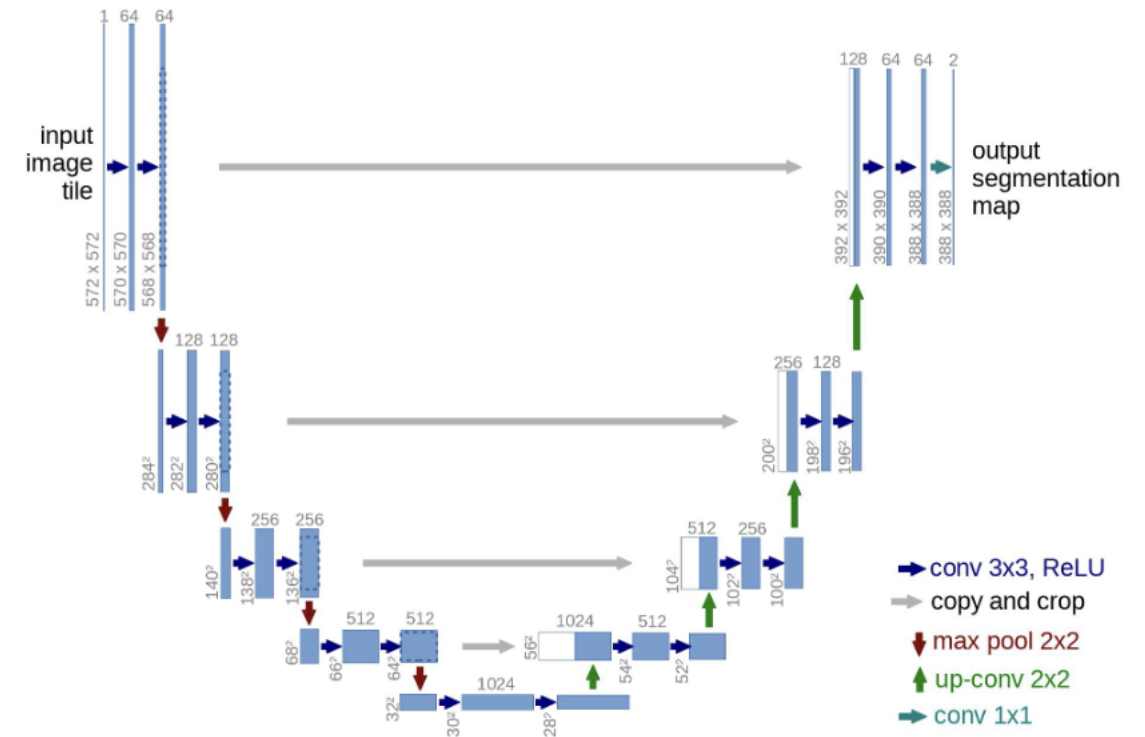



Figure 2: The structure of U-Net.

Starter Code Walkthrough

Noise Scheduler

- Control the amount of noise we add in each step of the diffusion forward process.



- we adopt the improved cosine-based variance schedule, introduced in (Nichol & Dhariwal, 2021).

$$\alpha_t = \text{clip} \left(\frac{\bar{\alpha}_t}{\bar{\alpha}_{t-1}}, 0.001, 1 \right), \bar{\alpha}_t = \frac{f(t)}{f(0)},$$
$$\text{where } f(t) = \cos \left(\frac{t/T + s}{1 + s} \cdot \frac{\pi}{2} \right)^2,$$

Starter Code Walkthrough

How do we train a denoising U-Net model?


Algorithm 1 Training

```
1: repeat  
2:    $\mathbf{x}_0 \sim q(\mathbf{x}_0)$   
3:    $t \sim \text{Uniform}(\{1, \dots, T\})$   
4:    $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$   
5:    $\mathbf{x}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$   $\triangleright$  forward diffusion process  
6:   Take optimizer step on  $L_1$  loss,  $\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(\mathbf{x}_t, t)\|_1$   
7: until converged
```

1. Take a training image
2. Pick a random time step
3. Run forward diffusion to generate a noisy version at that time step
4. Use our model to predict the noise that was added
5. Calculate the loss between the actual noise and the predicted noise

Starter Code Walkthrough


How do we sample an image?




t=4



t=3



t=1



t=0

We want a model that can revert images with any amount of noise $t=n$ to the previous step $t=n-1$

How do we achieve this?

Starter Code Walkthrough


How do we train a denoising U-Net model?




t=4



t=3



t=1



t=0

We want a model that can revert images with any amount of noise $t=n$ to the previous step $t=n-1$

We adopt the Option C Sampling algorithm from the lecture

Algorithm 1 Sampling (Option C)

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 2: **for** $t \in \{1, \dots, T\}$ **do**
 - 3: $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 4: $\hat{\mathbf{x}}_0 \leftarrow (\mathbf{x}_0 + (1 - \bar{\alpha}_t)\epsilon_{\theta}(\mathbf{x}_t, t)) / \sqrt{\bar{\alpha}_t}$
 - 5: $\hat{\boldsymbol{\mu}}_t \leftarrow \alpha_t^{(0)} \hat{\mathbf{x}}_0 + \alpha_t^{(t)} \mathbf{x}_t$
 - 6: $\mathbf{x}_{t-1} \leftarrow \hat{\boldsymbol{\mu}}_t + \sigma_t^2 \epsilon$
 - 7: **return** \mathbf{x}_0
-

Starter Code Walkthrough

How do we train a denoising U-Net model?

Algorithm 2 Sampling

```
1:  $\mathbf{x}_T \sim \mathcal{N}(0, I)$ 
2: for  $t = T, \dots, 1$  do
3:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, I)$  if  $t > 1$ , else  $\mathbf{z} = 0$ 
4:    $\boldsymbol{\epsilon}_t \leftarrow \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$  ▷ predicted noise
5:    $\hat{\mathbf{x}}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_t)$  ▷ estimated  $\hat{\mathbf{x}}_0$ 
6:    $\hat{\mathbf{x}}_0 \leftarrow \text{clamp}(\hat{\mathbf{x}}_0, -1, 1)$  ▷ rectify  $\hat{\mathbf{x}}_0$ 
7:    $\tilde{\boldsymbol{\mu}}_t \leftarrow \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0$  ▷ posterior mean of  $x_{t-1}$ 
8:    $\sigma_t^2 \leftarrow \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} (1 - \alpha_t)$  ▷ posterior variance of  $x_{t-1}$ 
9:    $\mathbf{x}_{t-1} \leftarrow \tilde{\boldsymbol{\mu}}_t + \sigma_t \mathbf{z}$  ▷ reverse diffusion process
return  $\mathbf{x}_0$ 
```

1. Start from a noisy image
2. Denoise in a loop

Starter Code Walkthrough

What functions do we need to write?

- Training
 - Forward
 - P_loss
 - Q_sample
- Sampling
 - Sample
 - P_sample_loop
 - P_sample

Starter Code Walkthrough

Flags

Configuration Parameter	Example Flag Usage
Model image size	<code>--image_size 32</code>
Model batch size	<code>--batch_size 32</code>
Model data domain of AFHQ dataset	<code>--data_class cat</code>
Directory where the model is stored	<code>--save_folder ./results/</code>
Path of a trained model	<code>--load_path ./results/model.pt</code>
Directory from which to load dataset	<code>--data_path ./data/train/</code>
Number of iterations to train the model	<code>--train_steps 10000</code>
Number of steps of diffusion process, T	<code>--time_steps 300</code>
Number of output channels of the first layer in U-Net	<code>--UNET_dim 16</code>
Learning rate in the training	<code>--learning_rate 1e-3</code>
Frequency of periodic save, sample and (optionally) FID calculation	<code>--save_and_sample_every 1000</code>
Enable FID calculation	<code>--fid</code>
Enable visualization	<code>--visualize</code>

Table 1: Useful flags for `main.py`

FID – Fréchet Inception Distance

How do we measure the quality of a generated image?



Vaguely Cat-like



Decently Fluffy Cats

Fréchet distance

$$F(A, B) = \inf_{\alpha, \beta} \max_{t \in [0, 1]} \left\{ d(A(\alpha(t)), B(\beta(t))) \right\}$$

Fréchet distance for probability distributions


$$d_F(\mu, \nu) := \left(\inf_{\gamma \in \Gamma(\mu, \nu)} \int_{\mathbb{R}^n \times \mathbb{R}^n} \|x - y\|^2 d\gamma(x, y) \right)^{1/2},$$

BUT for two multidimensional Gaussians

$$d_F(\mathcal{N}(\mu, \Sigma), \mathcal{N}(\mu', \Sigma'))^2 = \|\mu - \mu'\|_2^2 + \text{tr} \left(\Sigma + \Sigma' - 2(\Sigma\Sigma')^{1/2} \right)$$

FID – Fréchet Inception Distance

Inception model



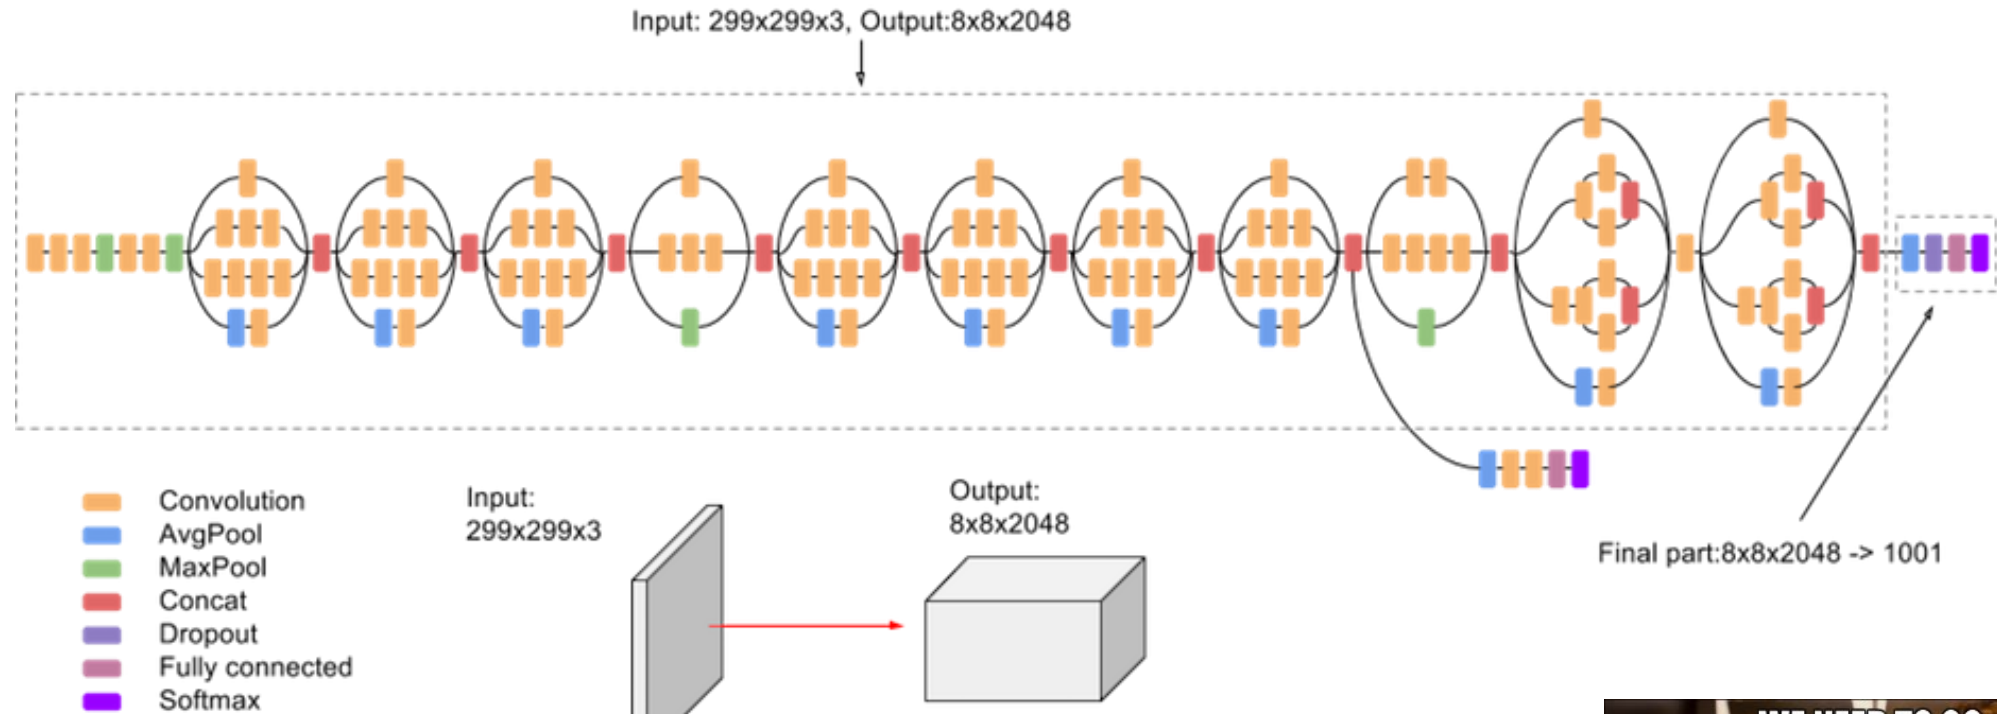
48 layers

SOTA in 2015 on ImageNet top-5 error

Named after an internet meme

FID – Fréchet Inception Distance

Inception model



48 layers

SOTA in 2015 on ImageNet top-5 error

Named after an internet meme



FID – Fréchet Inception Distance

How do we measure the quality of a generated image?



High score



Low score



Why do I need a huge NN and so much math to tell me if my cat photos are fluffy or ugly?

FID – Fréchet Inception Distance

Looking into the details

“Earthmover Distance”

- Wasserstein Distance
- Kantorovich-Rubinstein Metric
- Cramér distance
- Mallows distance
- Fréchet Distance
 - Wasserstien-2 Distance

These are all closely related ideas!




If transported optimally, how much probability mass would need to be moved to change one distribution into the other?

An alternative to the Kullback-Leibler Divergence

FID – Fréchet Inception Distance

Looking into the details


Inception module -- how does that work?



FID – Fréchet Inception Distance

Looking into the details

Inception module -- how does that work?



FID – Fréchet Inception Distance

Looking into the details

Inception module -- how does that work?

Massive feature extraction system

Find a features representation such that the images can be linearly separated

Linear
function of
inputs

FID – Fréchet Inception Distance

Putting it all together

Summary

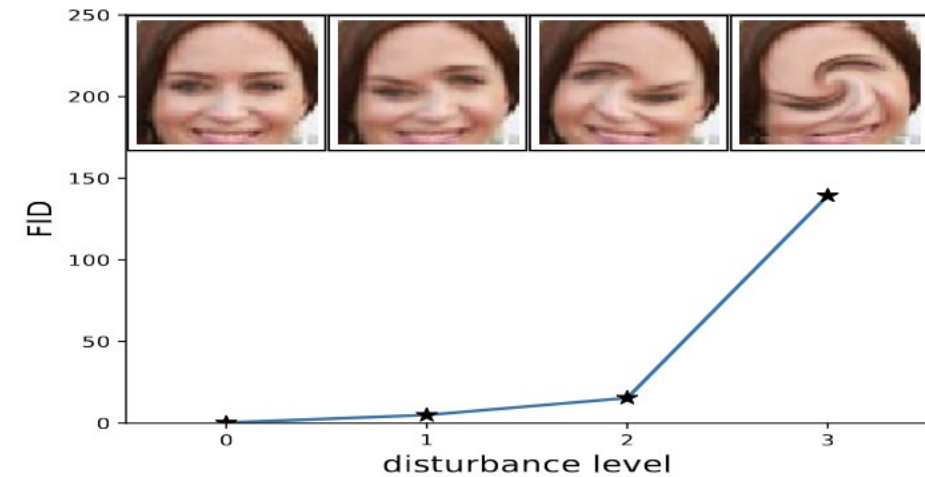

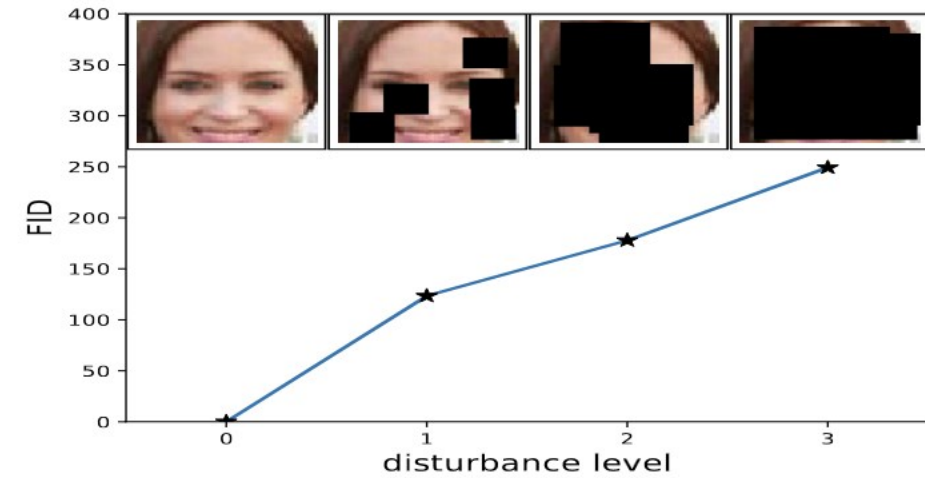

1. Extract the features from real images and generated images using an Inceptionv3 model.
2. Approximate the distribution of features as multivariate Gaussians (max entropy).
3. Find the distance between the two distribution of features.

Code snippet

```
Python 3.7.10 (default, Feb 26 2021, 18:47:35)
[GCC 7.3.0] :: Anaconda, Inc. on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> █
```

FID – Fréchet Inception Distance

Results



Helpful methods and functions for your homework

torch.clamp

TORCH.CLAMP

```
torch.clamp(input, min=None, max=None, *, out=None) → Tensor
```

Clamps all elements in `input` into the range `[min , max]`. Letting `min_value` and `max_value` be `min` and `max`, respectively, this returns:

$$y_i = \min(\max(x_i, \text{min_value}_i), \text{max_value}_i)$$

If `min` is `None`, there is no lower bound. Or, if `max` is `None` there is no upper bound.

• NOTE

If `min` is greater than `max` `torch.clamp(..., min, max)` sets all elements in `input` to the value of `max`.

Parameters

- **input** (*Tensor*) – the input tensor.
- **min** (*Number or Tensor, optional*) – lower-bound of the range to be clamped to
- **max** (*Number or Tensor, optional*) – upper-bound of the range to be clamped to

Keyword Arguments

out (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch
import numpy as np

x = np.array([[1, 2, 3], [4, 5, 6], [7, 8, 9]])
y = torch.clamp(x, min=4, max=6)
print(y)
```

Helpful methods and functions for your homework

torch.clamp

TORCH.CLAMP

```
torch.clamp(input, min=None, max=None, *, out=None) → Tensor
```

Clamps all elements in `input` into the range `[min , max]`. Letting `min_value` and `max_value` be `min` and `max`, respectively, this returns:

$$y_i = \min(\max(x_i, \text{min_value}_i), \text{max_value}_i)$$

If `min` is `None`, there is no lower bound. Or, if `max` is `None` there is no upper bound.

• NOTE

If `min` is greater than `max` `torch.clamp(..., min, max)` sets all elements in `input` to the value of `max`.

Parameters

- **input** (*Tensor*) – the input tensor.
- **min** (*Number or Tensor, optional*) – lower-bound of the range to be clamped to
- **max** (*Number or Tensor, optional*) – upper-bound of the range to be clamped to

Keyword Arguments

out (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch
import numpy as np

x = np.array([[1, 2, 3], [4, 5, 6], [7, 8, 9]])
y = torch.clamp(x, min=4, max=6)
print(y)
```

TypeError: clamp() received an invalid combination of arguments - got (numpy.ndarray, max=int, min=int), but expected one of: * (Tensor input, Tensor min, Tensor max, *, Tensor out) ***** (Tensor input, Number min, Number max, *, Tensor out)

Helpful methods and functions for your homework

torch.clamp

TORCH.CLAMP

```
torch.clamp(input, min=None, max=None, *, out=None) → Tensor
```

Clamps all elements in `input` into the range `[min , max]`. Letting `min_value` and `max_value` be `min` and `max`, respectively, this returns:

$$y_i = \min(\max(x_i, \text{min_value}_i), \text{max_value}_i)$$

If `min` is `None`, there is no lower bound. Or, if `max` is `None` there is no upper bound.

• NOTE

If `min` is greater than `max` `torch.clamp(..., min, max)` sets all elements in `input` to the value of `max`.

Parameters

- **input** (*Tensor*) – the input tensor.
- **min** (*Number or Tensor, optional*) – lower-bound of the range to be clamped to
- **max** (*Number or Tensor, optional*) – upper-bound of the range to be clamped to

Keyword Arguments

out (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch

x = torch.tensor([[1, 2, 3], [4, 5, 6], [7, 8, 9]])
y = torch.clamp(x, min=4, max=6)
print(y)
```


Helpful methods and functions for your homework

torch.clamp

TORCH.CLAMP

```
torch.clamp(input, min=None, max=None, *, out=None) → Tensor
```

Clamps all elements in `input` into the range `[min , max]`. Letting `min_value` and `max_value` be `min` and `max`, respectively, this returns:

$$y_i = \min(\max(x_i, \text{min_value}_i), \text{max_value}_i)$$

If `min` is `None`, there is no lower bound. Or, if `max` is `None` there is no upper bound.

• NOTE

If `min` is greater than `max` `torch.clamp(..., min, max)` sets all elements in `input` to the value of `max`.

Parameters

- **input** (*Tensor*) – the input tensor.
- **min** (*Number or Tensor, optional*) – lower-bound of the range to be clamped to
- **max** (*Number or Tensor, optional*) – upper-bound of the range to be clamped to

Keyword Arguments

out (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch

x = torch.tensor([[1, 2, 3], [4, 5, 6], [7, 8, 9]])
y = torch.clamp(x, min=4, max=6)
print(y)
```

Output:

```
tensor([[4, 4, 4], [4, 5, 6], [6, 6, 6]])
```

Helpful methods and functions for your homework

torch.cumprod

TORCH.CUMPROD

```
torch.cumprod(input, dim, *, dtype=None, out=None) → Tensor
```

Returns the cumulative product of elements of `input` in the dimension `dim`.

For example, if `input` is a vector of size `N`, the result will also be a vector of size `N`, with elements.

$$y_i = x_1 \times x_2 \times x_3 \times \dots \times x_i$$

Parameters

- **input** (*Tensor*) – the input tensor.
- **dim** (*int*) – the dimension to do the operation over

Keyword Arguments

- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. If specified, the input tensor is casted to `dtype` before the operation is performed. This is useful for preventing data type overflows. Default: None.
- **out** (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch

x = torch.tensor([[1, 2, 3, 4, 5]])
y = torch.cumprod(x, 0)
print(y)
```

Helpful methods and functions for your homework

torch.cumprod

TORCH.CUMPROD

```
torch.cumprod(input, dim, *, dtype=None, out=None) → Tensor
```

Returns the cumulative product of elements of `input` in the dimension `dim`.

For example, if `input` is a vector of size `N`, the result will also be a vector of size `N`, with elements.

$$y_i = x_1 \times x_2 \times x_3 \times \dots \times x_i$$

Parameters

- **input** (*Tensor*) – the input tensor.
- **dim** (*int*) – the dimension to do the operation over

Keyword Arguments

- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. If specified, the input tensor is casted to `dtype` before the operation is performed. This is useful for preventing data type overflows. Default: None.
- **out** (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch

x = torch.tensor([[1, 2, 3, 4, 5]])
y = torch.cumprod(x, 0)
print(y)
```

Output:
tensor([[1, 2, 3, 4, 5]])

Helpful methods and functions for your homework

torch.cumprod

TORCH.CUMPROD

```
torch.cumprod(input, dim, *, dtype=None, out=None) → Tensor
```

Returns the cumulative product of elements of `input` in the dimension `dim`.

For example, if `input` is a vector of size `N`, the result will also be a vector of size `N`, with elements.

$$y_i = x_1 \times x_2 \times x_3 \times \dots \times x_i$$

Parameters

- **input** (*Tensor*) – the input tensor.
- **dim** (*int*) – the dimension to do the operation over

Keyword Arguments

- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. If specified, the input tensor is casted to `dtype` before the operation is performed. This is useful for preventing data type overflows. Default: None.
- **out** (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch

x = torch.tensor([[1, 2, 3, 4, 5]])
y = torch.cumprod(x, 1)
print(y)
```

Helpful methods and functions for your homework

torch.cumprod

TORCH.CUMPROD

```
torch.cumprod(input, dim, *, dtype=None, out=None) → Tensor
```

Returns the cumulative product of elements of `input` in the dimension `dim`.

For example, if `input` is a vector of size `N`, the result will also be a vector of size `N`, with elements.

$$y_i = x_1 \times x_2 \times x_3 \times \dots \times x_i$$

Parameters

- **input** (*Tensor*) – the input tensor.
- **dim** (*int*) – the dimension to do the operation over

Keyword Arguments

- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. If specified, the input tensor is casted to `dtype` before the operation is performed. This is useful for preventing data type overflows. Default: None.
- **out** (*Tensor, optional*) – the output tensor.

What does the following code snippet return?

```
import torch

x = torch.tensor([[1, 2, 3, 4, 5]])
y = torch.cumprod(x, 1)
print(y)
```

Output:
tensor([[1, 2, 6, 24, 120]])

Helpful methods and functions for your homework

torch.full

TORCH.FULL

```
torch.full(size, fill_value, *, out=None, dtype=None, layout=torch.strided, device=None, requires_grad=False) → Tensor
```

Creates a tensor of size `size` filled with `fill_value`. The tensor's dtype is inferred from `fill_value`.

Parameters

- **size** (*int...*) – a list, tuple, or `torch.Size` of integers defining the shape of the output tensor.
- **fill_value** (*Scalar*) – the value to fill the output tensor with.

Keyword Arguments

- **out** (*Tensor, optional*) – the output tensor.
- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. Default: if `None`, uses a global default (see `torch.set_default_tensor_type()`).
- **layout** (`torch.layout`, optional) – the desired layout of returned Tensor. Default: `torch.strided`.
- **device** (`torch.device`, optional) – the desired device of returned tensor. Default: if `None`, uses the current device for the default tensor type (see `torch.set_default_tensor_type()`). `device` will be the CPU for CPU tensor types and the current CUDA device for CUDA tensor types.
- **requires_grad** (*bool, optional*) – If autograd should record operations on the returned tensor. Default: `False`.

What does the following code snippet return?

```
import torch

x1 = torch.full(2, 3, 3)
x2 = torch.ones(2,3) * 3

print(x1 == x2)
```

Helpful methods and functions for your homework

torch.full

TORCH.FULL

```
torch.full(size, fill_value, *, out=None, dtype=None, layout=torch.strided, device=None, requires_grad=False) → Tensor
```

Creates a tensor of size `size` filled with `fill_value`. The tensor's dtype is inferred from `fill_value`.

Parameters

- **size** (*int...*) – a list, tuple, or `torch.Size` of integers defining the shape of the output tensor.
- **fill_value** (*Scalar*) – the value to fill the output tensor with.

Keyword Arguments

- **out** (*Tensor, optional*) – the output tensor.
- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. Default: if `None`, uses a global default (see `torch.set_default_tensor_type()`).
- **layout** (`torch.layout`, optional) – the desired layout of returned Tensor. Default: `torch.strided`.
- **device** (`torch.device`, optional) – the desired device of returned tensor. Default: if `None`, uses the current device for the default tensor type (see `torch.set_default_tensor_type()`). `device` will be the CPU for CPU tensor types and the current CUDA device for CUDA tensor types.
- **requires_grad** (*bool, optional*) – If autograd should record operations on the returned tensor. Default: `False`.

What does the following code snippet return?

```
import torch

x1 = torch.full(2, 3, 3)
x2 = torch.ones(2,3) * 3

print(x1 == x2)
```

Output:

TypeError: full() received an invalid combination of arguments - got (int, int, int), but expected one of: * (tuple of ints size, Number fill_value, *, tuple of names names, torch.dtype dtype, torch.layout layout, torch.device device, bool pin_memory, bool requires_grad) * (tuple of ints size, Number fill_value, *, Tensor out, torch.dtype dtype, torch.layout layout, torch.device device, bool pin_memory, bool requires_grad)

Helpful methods and functions for your homework

torch.full

TORCH.FULL

```
torch.full(size, fill_value, *, out=None, dtype=None, layout=torch.strided, device=None, requires_grad=False) → Tensor
```

Creates a tensor of size `size` filled with `fill_value`. The tensor's dtype is inferred from `fill_value`.

Parameters

- **size** (*int...*) – a list, tuple, or `torch.Size` of integers defining the shape of the output tensor.
- **fill_value** (*Scalar*) – the value to fill the output tensor with.

Keyword Arguments

- **out** (*Tensor, optional*) – the output tensor.
- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. Default: if `None`, uses a global default (see `torch.set_default_tensor_type()`).
- **layout** (`torch.layout`, optional) – the desired layout of returned Tensor. Default: `torch.strided`.
- **device** (`torch.device`, optional) – the desired device of returned tensor. Default: if `None`, uses the current device for the default tensor type (see `torch.set_default_tensor_type()`). `device` will be the CPU for CPU tensor types and the current CUDA device for CUDA tensor types.
- **requires_grad** (*bool, optional*) – If autograd should record operations on the returned tensor. Default: `False`.

What does the following code snippet return?

```
import torch

x1 = torch.full((2, 3), 3)
x2 = torch.ones(2,3) * 3

print(x1 == x2)
```


Helpful methods and functions for your homework

torch.full

TORCH.FULL

```
torch.full(size, fill_value, *, out=None, dtype=None, layout=torch.strided, device=None, requires_grad=False) → Tensor
```

Creates a tensor of size `size` filled with `fill_value`. The tensor's dtype is inferred from `fill_value`.

Parameters

- **size** (*int...*) – a list, tuple, or `torch.Size` of integers defining the shape of the output tensor.
- **fill_value** (*Scalar*) – the value to fill the output tensor with.

Keyword Arguments

- **out** (*Tensor, optional*) – the output tensor.
- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. Default: if `None`, uses a global default (see `torch.set_default_tensor_type()`).
- **layout** (`torch.layout`, optional) – the desired layout of returned Tensor. Default: `torch.strided`.
- **device** (`torch.device`, optional) – the desired device of returned tensor. Default: if `None`, uses the current device for the default tensor type (see `torch.set_default_tensor_type()`). `device` will be the CPU for CPU tensor types and the current CUDA device for CUDA tensor types.
- **requires_grad** (*bool, optional*) – If autograd should record operations on the returned tensor. Default: `False`.

What does the following code snippet return?

```
import torch

x1 = torch.full((2, 3), 3)
x2 = torch.ones(2,3) * 3

print(x1 == x2)
```

Output:

```
tensor([[True, True, True], [True, True, True]])
```

Helpful methods and functions for your homework

torch.full

TORCH.FULL

```
torch.full(size, fill_value, *, out=None, dtype=None, layout=torch.strided, device=None, requires_grad=False) → Tensor
```

Creates a tensor of size `size` filled with `fill_value`. The tensor's dtype is inferred from `fill_value`.

Parameters

- **size** (*int...*) – a list, tuple, or `torch.Size` of integers defining the shape of the output tensor.
- **fill_value** (*Scalar*) – the value to fill the output tensor with.

Keyword Arguments

- **out** (*Tensor, optional*) – the output tensor.
- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. Default: if `None`, uses a global default (see `torch.set_default_tensor_type()`).
- **layout** (`torch.layout`, optional) – the desired layout of returned Tensor. Default: `torch.strided`.
- **device** (`torch.device`, optional) – the desired device of returned tensor. Default: if `None`, uses the current device for the default tensor type (see `torch.set_default_tensor_type()`). `device` will be the CPU for CPU tensor types and the current CUDA device for CUDA tensor types.
- **requires_grad** (*bool, optional*) – If autograd should record operations on the returned tensor. Default: `False`.

What does the following code snippet return?

```
import torch

x1 = torch.full((2,3), 3)
x2 = torch.ones(2,3) * 3

print((x1 == x2).all())
```

Helpful methods and functions for your homework

torch.full

TORCH.FULL

```
torch.full(size, fill_value, *, out=None, dtype=None, layout=torch.strided, device=None, requires_grad=False) → Tensor
```

Creates a tensor of size `size` filled with `fill_value`. The tensor's dtype is inferred from `fill_value`.

Parameters

- **size** (*int...*) – a list, tuple, or `torch.Size` of integers defining the shape of the output tensor.
- **fill_value** (*Scalar*) – the value to fill the output tensor with.

Keyword Arguments

- **out** (*Tensor, optional*) – the output tensor.
- **dtype** (`torch.dtype`, optional) – the desired data type of returned tensor. Default: if `None`, uses a global default (see `torch.set_default_tensor_type()`).
- **layout** (`torch.layout`, optional) – the desired layout of returned Tensor. Default: `torch.strided`.
- **device** (`torch.device`, optional) – the desired device of returned tensor. Default: if `None`, uses the current device for the default tensor type (see `torch.set_default_tensor_type()`). `device` will be the CPU for CPU tensor types and the current CUDA device for CUDA tensor types.
- **requires_grad** (*bool, optional*) – If autograd should record operations on the returned tensor. Default: `False`.

What does the following code snippet return?

```
import torch

x1 = torch.full((2,3), 3)
x2 = torch.ones(2,3) * 3

print((x1 == x2).all())
```

Output:
tensor(True)

You after finishing HW2

