
Reproducing Neural Discrete Representation Learning

Course Project for IFT 6135 - Representation Learning

Project Report link: [final_project.pdf](#)

Instructions

1. To train the VQVAE with default arguments as discussed in the report, execute:

```
1 python vqvae.py --data-folder /tmp/miniimagenet --output-folder models/  
   vqvae
```

2. To train the PixelCNN prior on the latents, execute:

```
1 python pixelcnn_prior.py --data-folder /tmp/miniimagenet --model models/  
   /vqvae --output-folder models/pixelcnn_prior
```

Datasets Tested

Image

1. MNIST
2. FashionMNIST
3. CIFAR10
4. Mini-ImageNet

Video

1. Atari 2600 - Boxing (OpenAI Gym) code

Reconstructions from VQ-VAE

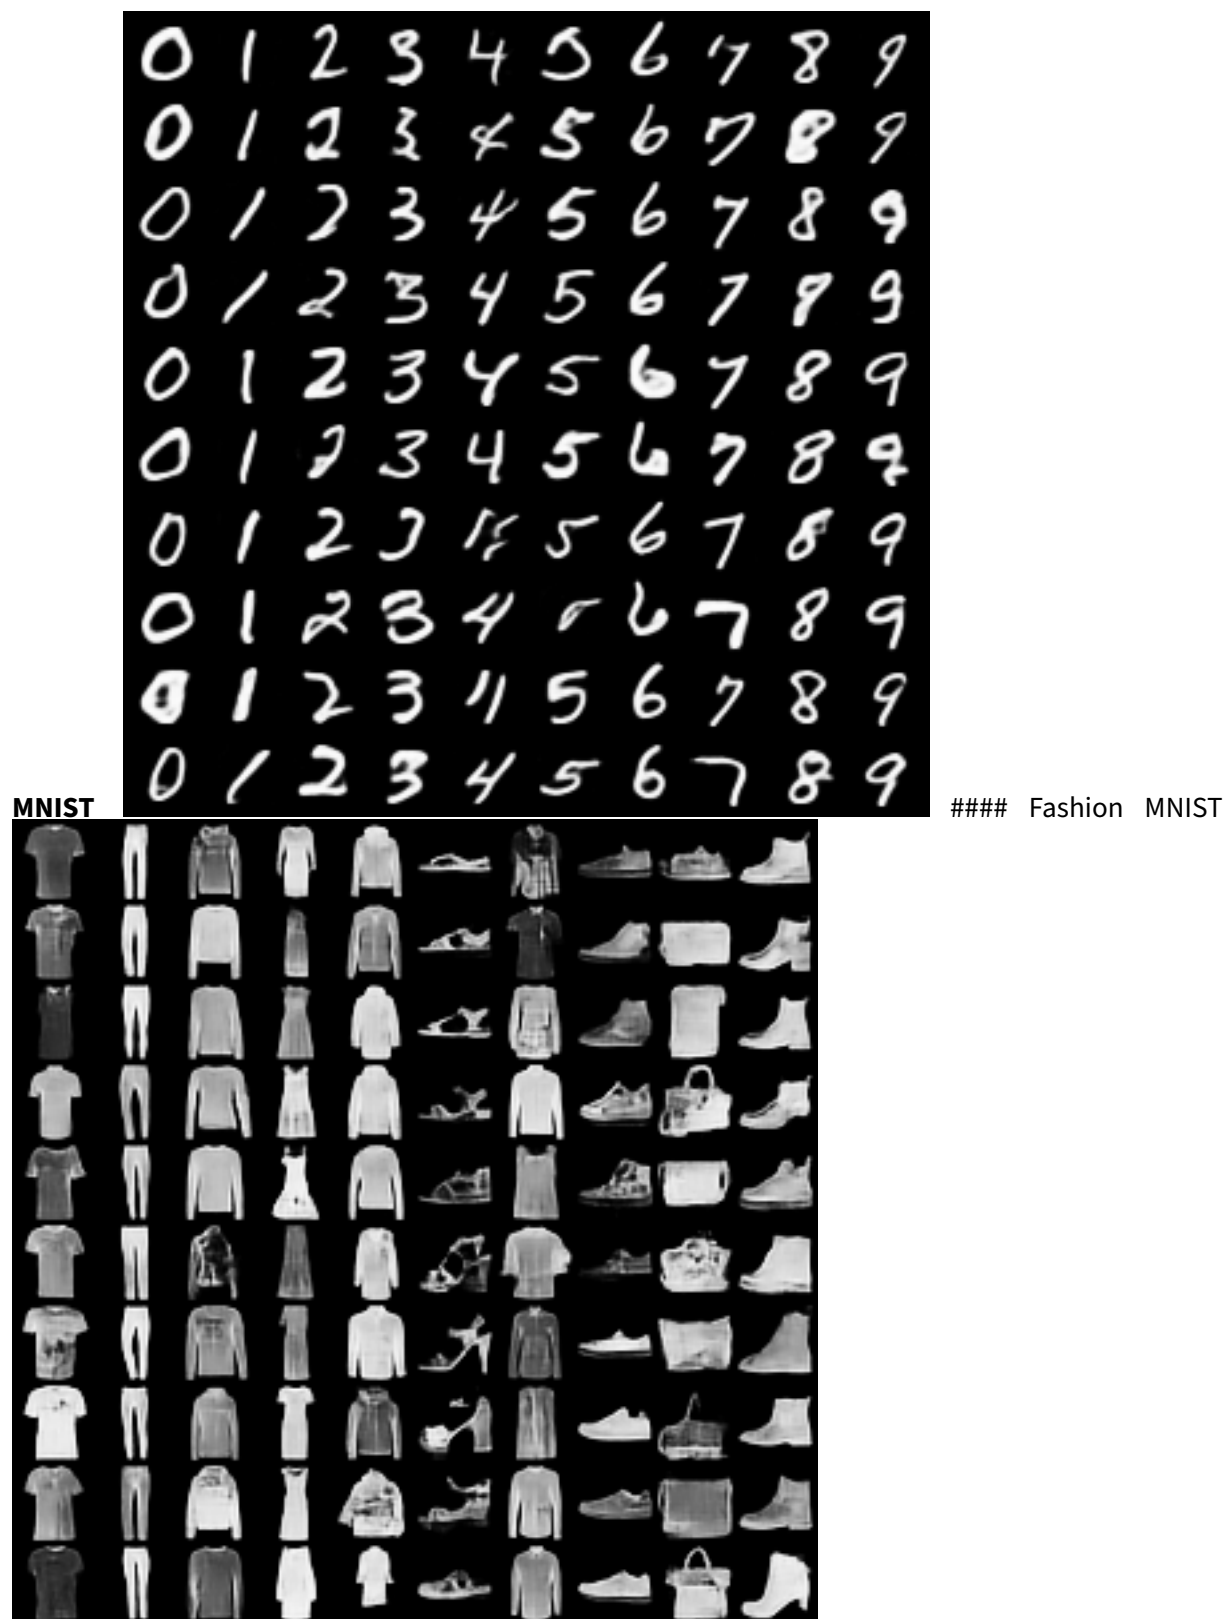


Top 4 rows are Original Images. Bottom 4 rows are Reconstructions. ##### MNIST



Fashion MNIST

Class-conditional samples from VQVAE with PixelCNN prior on the latents



Comments

1. We noticed that implementing our own VectorQuantization PyTorch function speeded-up training of VQ-VAE by nearly 3x. The slower, but simpler code is in this commit.
2. We added some basic tests for the vector quantization functions (based on [pytest](#)). To run these tests

```
1 py.test . -vv
```

Authors

1. Rithesh Kumar
2. Tristan Deleu
3. Evan Racah