

TTIC 31230 Fundamentals of Deep Learning, Fall 2022
Quiz 3

Problem 1. Consider a DDPM (denoising diffusion probabilistic model) for modeling $P(y)$ with $y \in \mathbb{R}^d$ where the noise model is defined by

$$z_0 = y$$

$$z_\ell = \alpha z_{\ell-1} + \sqrt{1 - \alpha^2} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$$

For technical simplicity we take α to be constant for all ℓ and allow $\ell \geq 1$ to be arbitrarily large.

(a) Write an equation for directly sampling z_ℓ from z_0 .

Solution: $z_\ell = \alpha^\ell z_0 + \sqrt{1 - \alpha^{2\ell}} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$

(b) Write an equation for sampling $z_{(\ell+k)|\ell}$ from $P(z_{(\ell+k)}|z_\ell)$.

Solution: $z_{(\ell+k)|\ell} = \alpha^k z_\ell + \sqrt{1 - \alpha^{2k}} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$

(c) Repeat (a) and (b) with $\alpha = e^{-\frac{1}{N}}$.

Solution:

$$z_\ell = e^{-\frac{\ell}{N}} z_0 + \sqrt{1 - e^{-\frac{2\ell}{N}}} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$$

$$z_{(\ell+k)|\ell} = e^{-\frac{k}{N}} z_\ell + \sqrt{1 - e^{-\frac{2k}{N}}} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$$

(d) Now let t range over fractions of the form k/N for $k \geq 0$. Rewrite your solution to (c) in terms of t rather than ℓ using $\ell = Nt$ and writing z_{Nt} as $w(t)$ and $z_{(Nt+N\Delta t)|Nt}$ as $w(t+\Delta t|t)$. It should be possible to interpret your solution for any real numbers $t \geq 0$ and $\Delta t \geq 0$.

Solution:

$$w(t) = e^{-t} w(0) + \sqrt{1 - e^{-2t}} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$$

$$w((t+\Delta t)|t) = e^{-\Delta t} w(t) + \sqrt{1 - e^{-2\Delta t}} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$$

(e) For small ϵ we have $e^{-\epsilon} \approx 1 - \epsilon$. Use this approximation to simplify your expression in (d) for $w((t+\Delta t)|t)$. Your solution should have the form of

a stochastic differential equation $\Delta w = \dot{w}\Delta t + \sqrt{\Delta t} \delta$ with δ drawn from a multivariate Gaussian independent of Δt .

Solution:

$$w((t + \Delta t)|t) \approx w(t) - w(t)\Delta t + \sqrt{2\Delta t} \epsilon$$

$$\Delta w \approx -w\Delta t + \sqrt{\Delta t} \delta \quad \delta \sim \mathcal{N}(0, 2I)$$

Problem 2. Again consider the noise process from problem 1.

$$z_0 = y$$

$$z_\ell = \alpha z_{\ell-1} + \sqrt{1 - \alpha^2} \epsilon \quad \epsilon \sim \mathcal{N}(0, I)$$

Again we take α constant for all ℓ with $0 < \alpha < 1$. However we now take the levels ℓ to be bounded with $1 \leq \ell \leq L$. We assume that α^L is sufficiently small that z_L is independent of y .

Now consider training a deterministic decoder (aka denoiser) $\hat{y}_\Phi(z_\ell, \ell)$ for recovering the image y from z_ℓ using the following loss.

$$\Phi^* = \operatorname{argmin}_\Phi E_{y, \ell, z_\ell} \|y - z_\Phi(z_\ell, \ell)\|^2 \quad (1)$$

This is done in training the decoder in some successful diffusion models.

Although the model $z_\Phi(z_\ell, \ell)$ is trained to predict y it is used to generate $z_{\ell-1}$ from z_ℓ in generation. The decoding process is sometimes made to be deterministic with

$$z_{\ell-1} = z_\Phi(z_\ell, \ell) \quad (2)$$

(a) Assuming universality for Φ , and assuming that z_L is distributed as $\mathcal{N}(0, I)$ independent of y , write $z_{\Phi^*}(z_L, L)$ as an expression not involving optimization (not involving argmin).

Solution: Since z_L is independent of y , the minimizer is just the mean image.

$$\hat{y}_{\Phi^*}(z_L, L) = E y$$

(b) Does your answer to (a) have implications for the diversity of generated images in a model that uses both (1) and (2). Explain your answer.

Solution: Part (a) gives that z_{L-1} is deterministically the mean image independent of the noise z_L . If all other stages in the decoding are deterministic then the sampler should always get the same image.

(c) Why might an image generator based on (1) and (2) be diverse in practice?

Solution: The universality assumption does not hold in practice and hence we should not expect $z_{\ell-1}(z_L, L)$ to be truly constant independent of z_L . This variation can be amplified through the decoding process.

Problem 3. In a progressive VAE we are interested in modeling the distribution $p_{\Phi}(z_{\ell-1}|z_{\ell}, \ell)$. Current diffusion models use

$$z_{\ell-1} = z_{\Phi}(z_{\ell}, \ell) + \sigma \delta \quad \delta \sim \mathcal{N}(0, I)$$

However, if want to reduce the number of layers we have more noise between layers and the true conditional distribution $p(z_{\ell-1}|z_{\ell}, \ell)$ will not be Gaussian.

This problem asks you to formulate a conditional Gaussian VAE model for the conditional distribution $p_{\Phi}(z_{\ell-1}|z_{\ell}, \ell)$. Here z_{ℓ} and ℓ are given and you are to introduce a latent variable with an encoder, decoder and prior, for modeling $p_{\Phi}(z_{\ell-1}|z_{\ell}, \ell)$. In a Gaussian VAE we have that, without loss of generality, the prior can be taken to be $\mathcal{N}(0, I)$. In a Gaussian VAE for images the latent variable typically has smaller dimension than the images.

Letting ϵ denote the latent variable of the Gaussian VAE, and taking the prior to be $\mathcal{N}(0, 1)$, sampling from the prior, followed by sampling from the decoder, can be written as

$$z_{\ell-1} = z_{\Phi}(\epsilon, z_{\ell}, \ell) + \sigma_{\Phi}(\epsilon, z_{\ell}, \ell) \odot \delta \quad \epsilon \sim \mathcal{N}(0, I), \quad \delta \sim \mathcal{N}(0, I)$$

Here \odot denotes Hadamard product, or dimensionwise product, with $(x \odot y)[i] = x[i]y[i]$.

Write a similar equation for the Gaussian VAE encoder generating the latent variable ϵ from z_{ℓ} and ℓ and give the objective function for jointly training the encoder and the decoder.

Solution: For encoding the latent variable ϵ we introduce a model

$$\epsilon = \epsilon_{\Psi}(z_{\ell-1}, z_{\ell}, \ell) + \sigma_{\Psi}(z_{\ell-1}, z_{\ell}, \ell) \odot \gamma \quad \gamma \sim \mathcal{N}(0, I)$$

The training objective is then

$$\Phi^*, \Psi^* = \underset{\Phi, \Psi}{\operatorname{argmin}} E_{y, \ell, z_{\ell-1}, z_{\ell}, \epsilon} KL(p_{\Psi}(\epsilon|z_{\ell-1}, z_{\ell}, \ell), \mathcal{N}(0, I)) - \ln p_{\Phi}(z_{\ell-1}|\epsilon, z_{\ell}, \ell)$$

There are various equivalent ways of writing this which also get full credit.