



Homework 3
Due 5pm, Wednesday, December 21, 2022

Problem 1: In diffusion models, classifier-guided conditional generation uses

$$\nabla_x \log p_t(x) + \omega \nabla_x \log p_t(y | x)$$

with a classifier scale parameter $\omega \geq 1$. However, many papers in the literature add classifier guidance to an already conditional model via

$$\nabla_x \log p_t(x | y) + s \nabla_x \log p_t(y | x)$$

with $s \geq 0$. Show that the two are equivalent with $s = \omega - 1$.

Problem 2: Let $X_1, \dots, X_N \sim p$ be an IID sequence of random variables. Let $f: \mathcal{X} \rightarrow \mathbb{R}$. Let $g: \mathcal{X}^N \rightarrow \mathbb{R}$ be a permutation invariant function, i.e., $g(x_1, \dots, x_N) = g(x_{\sigma(1)}, \dots, x_{\sigma(N)})$ for any permutation σ . Show that

$$\mathbb{E} \left[\frac{f(X_1)}{g(X_1, \dots, X_N)} \right] = \frac{1}{N} \mathbb{E} \left[\frac{f(X_1) + \dots + f(X_N)}{g(X_1, \dots, X_N)} \right].$$

Remark. Exchangeability, i.e., that $(X_1, \dots, X_N) \stackrel{\mathcal{D}}{=} (X_{\sigma(1)}, \dots, X_{\sigma(N)})$ for any permutation σ , is actually sufficient to establish this result. IID implies exchangeability.

Problem 3: *Unitary change of coordinates of diffusion.* Consider the d -dimensional variance-exploding discrete diffusion

$$X_{k+1} = X_k + \sigma \varepsilon_k$$

for $k = 0, 1, \dots$ where $\sigma > 0$, $X_0 \sim p_0$, and $\varepsilon_1, \varepsilon_2, \dots$ is an IID sequence with distribution $\mathcal{N}(0, I)$. Let $U \in \mathcal{R}^{d \times d}$ be an orthogonal matrix and define

$$Z_k = U X_k$$

for $k = 0, 1, \dots$. Show that

$$Z_{k+1} \stackrel{\mathcal{D}}{=} Z_k + \sigma \tilde{\varepsilon}_k$$

for $k = 0, 1, \dots$, where $\tilde{\varepsilon}_1, \tilde{\varepsilon}_2, \dots$ is an IID sequence with distribution $\mathcal{N}(0, I)$.

Clarification. Of course, Z_0 does not have distribution p_0 .

Problem 4: DDPM considers the forward corruption process

$$X_t = \sqrt{\bar{\alpha}_t} X_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon_t$$

for $t = 0, \dots, T$, where $X_0 \sim p_0$, $\varepsilon_t \sim \mathcal{N}(0, I)$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. Sampling is done by

$$X_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(X_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \varepsilon_\theta(X_t, t) \right) + \sqrt{\tilde{\beta}_t} Z_t$$

for $t = T, T-1, \dots, 1$, where $Z_t \sim \mathcal{N}(0, I)$, $\tilde{\beta}_t := \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$ and $\alpha_t := 1 - \beta_t$. Define

$$\hat{X}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (X_t - \sqrt{1 - \bar{\alpha}_t} \varepsilon_\theta(X_t, t)).$$

Since ε_θ is trained so that $\varepsilon_\theta(X_t, t)$ estimates ε_t , we can view \hat{X}_0 is an estimator of X_0 . Show that

$$p_\theta(X_{t-1} | X_t) = q(X_{t-1} | X_t, X_0 = \hat{X}_0),$$

where p_θ is the learned reverse-conditional distribution and q is the true conditional distribution defined by the forward process.