


Class Embedding Patterns in Classifier-Free Diffusion

Background: Classifier-free DDPM

- Training: $\|\epsilon_{\theta}(\mathbf{z}_{\lambda}, \mathbf{c}) - \epsilon\|^2$ where $(\mathbf{x}, \mathbf{c}) \sim p(\mathbf{x}, \mathbf{c})$; $\mathbf{c} \leftarrow \emptyset$ with probability p_{uncond}
- Sampling: $\tilde{\epsilon}_t = (1 + w)\epsilon_{\theta}(\mathbf{z}_t, \mathbf{c}) - w\tilde{\epsilon}_{\theta}(\mathbf{z}_t)$

Novelty: Class Embedding Networks with various patterns



Class Embedding Patterns in Classifier-Free Diffusion

Experiments: Uniform Embedding, Pyramid Embedding and Bottleneck Embedding

Measures: Fréchet Inception Distance (FID), Inception Score (IS) and Negative Log Likelihood (NLL)

Class Embed Type	FID	IS	NLL
Uniform-1	3,253	3.575	3.084
Uniform-16	3.240	2.717	1.851
Uniform-256	3.240	4.671	0.392
Complete Bottleneck	3.223	7.416	0.003
Bottleneck	3.219	7.711	0.005
Pyramid	3.217	7.472	0.027



Class Embedding Patterns in Classifier-Free Diffusion

Improvement: Curriculum Learning with noisy labels

Intuition: (1) pick hard samples; (2) consistent with all-sample dynamics

Sampling training batch w.r.t the following score:

$$a_t(i) \triangleq \left\langle y_i - f(x_i; \theta_t), \frac{\partial f(x_i; \theta_t)}{\partial t} \Big|_D \right\rangle.$$

Residual: How hard the sample is

Linear Dynamics: Speed and direction how model changes