Diffusion Distillation with Direct Preference Optimization for Efficient 3D LiDAR Scene Completion

An Zhao[1], Shengyuan Zhang[1], Zejian Li[2*], Ling Yang[3], Pei Chen[1], Jiale Wu[1], Haoran Xu[4], AnYang Wei[4], Perry Pengyun Gu[4], Lingyun Sun[1]

[1]College of Computer Science and Technology, Zhejiang University
[2]School of Software Technology, Zhejiang University
[3]Peking University
[4]Zhejiang Green Zhixing Technology co., ltd
{zhaoan040113,zhangshengyuan,zejianlee,chenpei,jialewu2022,sunly}@zju.edu.cn, yangling0818@163.com,
{Haoran.Xu5,weianyang,gupengyun}@geely.com

Abstract

The slow sampling speed of diffusion models hinders their application in 3D LiDAR scene completion. To address this, we propose Distillation-DPO, a novel framework that accelerates sampling through score distillation while simultaneously enhancing generation quality via preference alignment. Distillation-DPO follows a three-step procedure. First, the student model generates paired completion scenes with different initial noises. Second, using LiDAR scene evaluation metrics as preference, we construct winning and losing sample pairs. Third, as our core innovation, Distillation-DPO optimizes the student model by exploiting the difference in score functions between the teacher and student models on the paired completion scenes. This operation performs variational score distillation of the student model but simultaneously encourages the distilled student to prefer the winning samples over the losing ones. Extensive experiments demonstrate that Distillation-DPO achieves higher-quality scene completion than state-of-the-art diffusion models, while accelerating sampling by over 5-fold. To our knowledge, our work is the first to integrate the preference learning principle of DPO into the distillation of diffusion models, offering a new framework of preference-aligned distillation.

Code — https://github.com/happyw1nd/DistillationDPO

Introduction

In recent years, diffusion models have demonstrated exceptional performance in different generative tasks (Karras et al. 2022; Zhou et al. 2024; Wang et al. 2023), emerging as a key paradigm in visual content generation. As their capabilities continue to expand, diffusion models are increasingly being applied to more challenging 3D tasks, such as 3D LiDAR scene completion (Nunes et al. 2024; Ran, Guizilini, and Wang 2024). Due to occlusions, viewpoint limitations, and sensor sparsity, LiDAR point clouds are often sparse (Nunes et al. 2024; Nakashima and Kurazume 2024). Given the advantage in modeling complex data distributions, diffusion models have been regarded as a promising solution for high-quality point cloud completion. However, achieving high-fidelity completion typically requires more sampling steps, resulting in a slow inference process. This inefficiency in sampling limits their practicality in real-world applications.

Score distillation, as an effective distillation method for diffusion models, enables a student model to generate new samples in fewer steps than the teacher diffusion model requires (Luo et al. 2023; Yin et al. 2024a,b). It provides an effective pathway for accelerating LiDAR scene completion diffusion models. However, score distillation inevitably leads to information loss and a decline in quality in the completed scene during the sampling acceleration process.

Reward models provide a potential way to mitigate the performance degradation caused by distillation. The reward model learns human preferences to predict the rating of generated samples, while existing methods primarily enhance generation quality by maximizing the rating predicted by the reward model (Zhang et al. 2025; Xu et al. 2023). However, the application of the reward model in score distillation of LiDAR scene completion faces the following challenges. First, due to the complexity of LiDAR scenes, obtaining large-scale human-labeled data is challenging. With limited data, the reward model is easily over-optimized and faces the issue of reward hacking (Back, Piao, and Kim 2024). Second, existing methods often use differentiable rewards to optimize the model (Clark et al. 2024), but commonly evaluation metrics such as IoU (Song et al. 2017) and EMD (Fan, Su, and Guibas 2017) are non-differentiable and computationally expensive, difficult to use directly as rewards to optimize the diffusion model.

Compared to reward models, Diffusion-DPO (Wallace et al. 2024; Rafailov et al. 2023) directly optimizes the diffusion model using preference data pairs, eliminating the need for training an additional reward model and thus mitigating the issue of reward hacking. Thus, to tackle the above challenges, we incorporate score distillation with the post-training of DPO and propose a novel distillation framework dubbed Distillation-DPO for LiDAR scene completion diffusion models. Distillation-DPO includes an effective distillation strategy on the preference completed scene pairs for the first time. Specifically, based on the completed

*Corresponding author

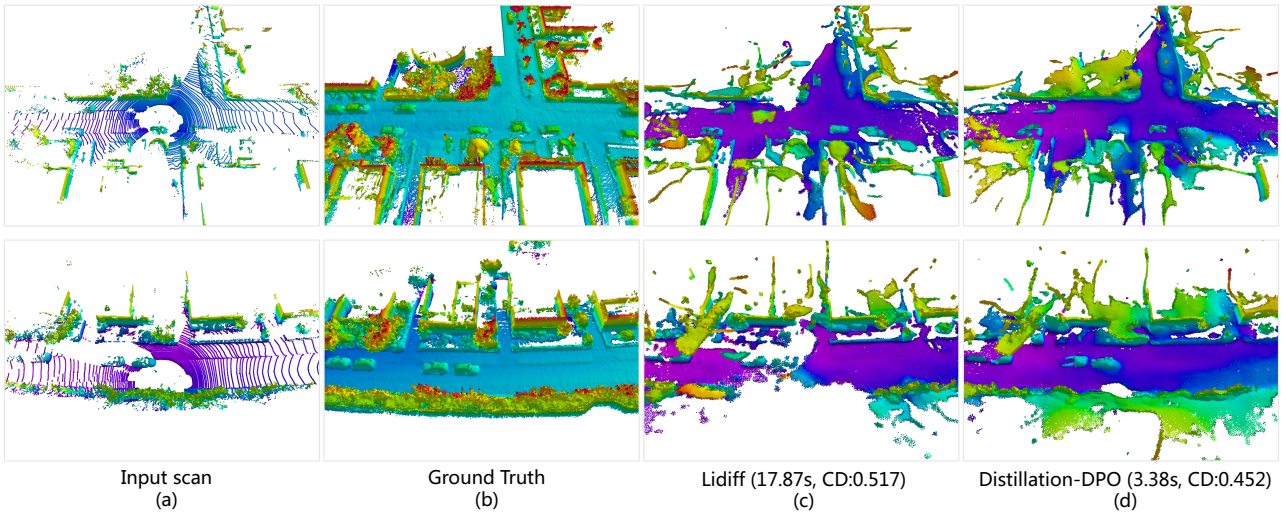

Figure 1: An example demonstration of Distillation-DPO for LiDAR scene completion on SemanticKITTI dataset. (a) The input sparse LiDAR scan. (b) The corresponding ground truth scene. (c) Completion results of the existing state-of-the-art (SOTA) model, LiDiff (Nunes et al. 2024). (d) Completion results of the proposed Distillation-DPO. Compared to LiDiff, Distillation-DPO can complete a scene more than 5 times faster while achieving higher completion quality (lower Chamfer Distance).

scene generated by the student model, we use LiDAR scene evaluation metrics as guidance to construct the win-lose preference pairs. Then, Distillation-DPO optimizes the student model by computing the score function on the student and teacher models with the pair-wise data. Compared with state-of-the-art LiDAR scene completion diffusion models, Distillation-DPO achieves five-fold acceleration for LiDAR completion while delivering higher-quality completion results, setting a new SOTA performance (Fig. 1).

Our contributions are summarized as follows: (1) We propose Distillation-DPO, a novel distillation framework for LiDAR scene completion diffusion models, which is the first to perform distillation based on preference data pairs. (2) Compared to existing LiDAR scene completion models, Distillation-DPO achieves improvements in both completion quality and speed.

Related Work

LiDAR Scene Completion

LiDAR scene completion aims to reconstruct sparse LiDAR scans into dense and complete 3D point cloud scenes (Zhang et al. 2024). Traditional LiDAR scene completion methods recover dense depth maps from sparse point clouds (Xu et al. 2019), leveraging guidance from RGB images or bird's-eye view images to achieve high-quality completion (Chen et al. 2019). Some methods represent LiDAR scenes as voxels and utilize Signed Distance Fields (SDFs) to reconstruct complete point cloud scenes (Li et al. 2023). However, the completion quality of these methods is constrained by the voxel resolution (Nunes et al. 2024). Due to the high generative quality and strong training stability, many studies have recently leveraged diffusion models for high-quality LiDAR scene completion (Cao and Behnke 2024; Nakashima and Kurazume 2024). Some methods focus on reconstructing sparse LiDAR scans into dense scans, such as R2DM (Nakashima and Kurazume 2024), OLiDM (Yan et al. 2024), and LiDMs (Ran, Guizilini, and Wang 2024). Other approaches attempt to directly recover complete point cloud scenes from sparse LiDAR scans, including LiDiff (Nunes et al. 2024) and DiffSSC (Cao and Behnke 2024). To further accelerate completion speed, ScoreLiDAR introduces a distillation method based on structural loss, enabling fast and efficient scene completion (Zhang et al. 2024).

Diffusion Distillation

Distillation strategies for diffusion models leverage a teacher-student framework to compress the multi-step teacher model into student models capable of few-step or single-step sampling. For example, Knowledge Distillation (Luhman and Luhman 2021) compresses the complex generative trajectory of the teacher into the one-step output of the student while maintaining visual fidelity comparable. Building on this, Progressive Distillation (Salimans and Ho 2021) introduces a stage-wise compression strategy. Through multiple iterations, the sampling steps of the student model are progressively reduced, ultimately achieving high-quality outputs with few sampling steps. Recently, inspired by the variational score distillation paradigm proposed in ProlificDreamer (Wang et al. 2023), score distillation methods have advanced this line of work by directly aligning the multi-step generative distribution of the teacher with the few- or single-step distribution of the student. Score distillation minimizes the discrepancy between the score function predictions of the two models, effectively enhancing the student's generative capability (Wang et al. 2023; Luo et al. 2023; Yin et al. 2024b). While distillation methods for diffusion models have demonstrated impressive performance in image and video generation tasks, their application to LiDAR scene completion remains largely unexplored.

Preference Optimization for Diffusion Models

To generate results that better align with human preferences, some studies train models based on preference-optimization methods (Ouyang et al. 2022; Liang et al. 2024; Rafailov et al. 2023). ImageReward (Xu et al. 2023) proposes the first general human preference reward model for text-to-image tasks and directly optimizes the diffusion model based on human feedback. Subsequent studies leveraged more detailed annotation methods (Liang et al. 2024) and combined multiple open-source models (Zhang et al. 2025) to obtain richer human feedback datasets. Additionally, some works optimize reward feedback learning by integrating multiple reward models (Guo et al. 2024) or improving training policies (Zhang et al. 2025). Since obtaining large-scale human annotations is challenging, some methods have attempted to train reward models using semi-supervised learning with unlabeled data (He et al. 2024) or employing hybrid annotation strategies with AI and human (Mahan et al. 2024). Additionally, Diffusion-DPO (Wallace et al. 2024) is the first to extend Direct Preference Optimization (Rafailov et al. 2023) to diffusion models, directly optimizing the model based on image preferences to eliminate the complex reward modeling and improve training efficiency.

Preliminaries

LiDAR Scene Completion Diffusion Model

The goal of the LiDAR scene completion diffusion model $\boldsymbol{\epsilon}_\theta$ is to predict noise based on the given LiDAR sparse scan $\mathcal{P}$, enabling a step-by-step denoising process from an initial noisy sample $\mathcal{G}_T$ to obtain a dense scene reconstruction $\mathcal{G}_0$. In the existing SOTA model LiDiff (Nunes et al. 2024), the sampling step is often set to 50.

Given an input sparse scan $\mathcal{P} = \{\boldsymbol{p}^1, \boldsymbol{p}^2, ..., \boldsymbol{p}^N\}$ and the ground truth $\mathcal{G} = \{\boldsymbol{p}^1, \boldsymbol{p}^2, ..., \boldsymbol{p}^M\}$ (N $\ll$ M), the noisy point cloud $\mathcal{G}_t = \{\boldsymbol{p}_t^1, \boldsymbol{p}_t^2, ..., \boldsymbol{p}_t^M\}$ can be calculated in a point-wise manner (Nunes et al. 2024):

$$\boldsymbol{p}_t^m = \boldsymbol{p}^m + \left(\sqrt{\bar{\alpha}_t}\mathbf{0} + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon_t}\right) = \boldsymbol{p}^m + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon_t} \quad (1)$$

Here $p^m \in \mathbb{R}^3$ is the point cloud. This diffusion method is similar to the variance exploding stochastic differential equation (Song et al. 2021) with the maximum variance controlled. It is adopted because directly applying traditional noise injection methods like DDPM (Ho, Jain, and Abbeel 2020) would compress the LiDAR point cloud into a smaller scale than the original cloud, leading to loss of details.

Due to the local diffusion method in Eq. 1, $\mathcal{G}_T$ cannot be directly approximated by the Gaussian distribution. Given a sparse LiDAR scan $\mathcal{P}$, the point in $\mathcal{P}$ is first replicated K times to obtain a dense scan $\mathcal{P}^* = \{\boldsymbol{p}^{1*}, \boldsymbol{p}^{2*}, \ldots, \boldsymbol{p}^{M*}\}$. Then, the initial noisy point cloud $\mathcal{G}_T^* = \{\boldsymbol{p}_T^{1*}, \boldsymbol{p}_T^{2*}, \ldots, \boldsymbol{p}_T^{M*}\}$ is calculated by sampling a Gaussian noise for each $\boldsymbol{p}^{m*} \in \mathcal{P}^*$ based on Eq. 1. Finally, a step-by-step denoising process in Eq. 2 is conducted to generate the completed scene $\mathcal{G}_0$.

$$\mathcal{G}^{t-1} = \frac{1}{\sqrt{\alpha^t}}\left(\mathcal{G}^t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}^t}}\boldsymbol{\epsilon}_\theta\left(\mathcal{G}^t, \mathcal{P}, t\right)\right) + \sigma^t \mathbf{z} \quad (2)$$

Score Distillation

Score distillation aims to encourage the few-step distribution of the student model to be close to the multi-step distribution of the teacher model. Let p_η and p_θ be the distribution of the student model and the teacher model, respectively. Score Distillation aims to minimize the following KL divergence

$$\min_\eta D_{KL}\left(p_\eta\left(\boldsymbol{x}_0\right) \| p_\theta\left(\boldsymbol{x}_0\right)\right) \quad (3)$$

Solving the optimization problem in Eq. 3 directly is difficult. According to Theorem 1 in Wang et al. (2023), Eq. 3 is equivalent to the optimization problems over the noisy distributions in different timesteps t

$$\min_\eta \mathbb{E}_{t,\boldsymbol{\epsilon}}\left[D_{KL}\left(p_{\eta,t}\left(\boldsymbol{x}_t\right) \| p_{\theta,t}\left(\boldsymbol{x}_t\right)\right)\right] \quad (4)$$

Thus, the gradient to the student parameterized by η is

$$\begin{aligned}&\nabla_\eta D_{KL}\left(p_{\eta,t}\left(\boldsymbol{x}_t\right) \| p_{\theta,t}\left(\boldsymbol{x}_t\right)\right)\\ &= \mathbb{E}_{t,\boldsymbol{\epsilon}}\left[\nabla_{\boldsymbol{x}_t}\log p_{\eta,t}\left(\boldsymbol{x}_t\right) - \nabla_{\boldsymbol{x}_t}\log p_{\theta,t}\left(\boldsymbol{x}_t\right)\right]\frac{\partial \boldsymbol{x}_t}{\partial \eta}\end{aligned} \quad (5)$$

Then, the score $\nabla_{\boldsymbol{x}_t}\log p_{\theta,t}\left(\boldsymbol{x}_t\right)$ can be approximated by the pre-trained diffusion model $\boldsymbol{\epsilon}_\theta$, and the score $\nabla_{\boldsymbol{x}_t}\log p_{\eta,t}\left(\boldsymbol{x}_t\right)$ can be approximated by an auxiliary model $\boldsymbol{\epsilon}_\phi$ which trained on the generative samples of the student model with standard diffusion loss (Wang et al. 2023). Thus, the gradient in Eq. 5 can be approximated by

$$\begin{aligned}&\nabla_\eta D_{KL}\left(p_{\eta,t}\left(\boldsymbol{x}_t\right) \| p_{\theta,t}\left(\boldsymbol{x}_t\right)\right)\\ &\approx \mathbb{E}_{t,\boldsymbol{\epsilon}}\left[\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t) - \boldsymbol{\epsilon}_\phi(\boldsymbol{x}_t, t)\right]\frac{\partial \boldsymbol{x}_t}{\partial \eta}\end{aligned} \quad (6)$$

During training, the student model and the auxiliary model $\boldsymbol{\epsilon}_\phi$ are optimized alternately.

A Brief Introduction of Diffusion-DPO

This part reviews the Direct Preference Optimization in diffusion models (Diffusion-DPO) (Wallace et al. 2024). Let $\mathcal{D} = \{(\boldsymbol{c}, \boldsymbol{x}_0^w, \boldsymbol{x}_0^l)\}$ is a dataset, where each data sample consists of a prompt $\boldsymbol{c}$ and a pair of images $\boldsymbol{x}_0^w$ and $\boldsymbol{x}_0^l$ with preference $\boldsymbol{x}_0^w \succ \boldsymbol{x}_0^l$. The image $\boldsymbol{x}_0^w$ and $\boldsymbol{x}_0^l$ are both sampled from a reference distribution p_{ref}. To obtain the reward on the whole diffusion path, $r(\boldsymbol{c}, \boldsymbol{x}_0)$ is defined as:

$$r\left(\boldsymbol{c}, \boldsymbol{x}_0\right) = \mathbb{E}_{p_\eta(\boldsymbol{x}_{1:T}|\boldsymbol{x}_0, \boldsymbol{c})}\left[R\left(\boldsymbol{c}, \boldsymbol{x}_{0:T}\right)\right] \quad (7)$$

Here p_η is a diffusion model trained to align with the preferences. Then, p_η can be optimized by maximizing the following objective

$$\begin{aligned}&\max_{p_\eta} \mathbb{E}_{\boldsymbol{c}\sim\mathcal{D}_c, \boldsymbol{x}_{0:T}\sim p_\eta(\boldsymbol{x}_{0:T}|\boldsymbol{c})}\left[r\left(\boldsymbol{c}, \boldsymbol{x}_0\right)\right]\\ &-\beta\mathbb{D}_{\text{KL}}\left[p_\eta\left(\boldsymbol{x}_{0;T} \mid \boldsymbol{c}\right) \| p_{\text{ref}}\left(\boldsymbol{x}_{0:T} \mid \boldsymbol{c}\right)\right]\end{aligned} \quad (8)$$

Compared to traditional DPO (Rafailov et al. 2023), the objective function in Eq. 8 is defined throughout the diffusion path $\boldsymbol{x}_{0:T}$, which aims to maximize the reward $r\left(\boldsymbol{c}, \boldsymbol{x}_0\right)$ while ensuring that the distributions of p_η and p_{ref} remain close. Following Wallace et al. (2024), the objective in Eq. 8 can be further transformed into the following objective:

$$\begin{aligned}&L_{\text{DPO}-\text{Diffusion}}(\eta) = -\mathbb{E}_{\left(\boldsymbol{x}_0^w, \boldsymbol{x}_0^t\right)\sim\mathcal{D}}\log\sigma\\ &\left(\beta\mathbb{E}_{\substack{\boldsymbol{x}_{1:T}^w\sim p_\eta\left(\boldsymbol{x}_{1:T}^w|\boldsymbol{x}_0^w\right)\\ \boldsymbol{x}_{1:T}^L\sim p_\eta\left(\boldsymbol{x}_{1:T}^l|\boldsymbol{x}_0^l\right)}}\left[\log\frac{p_\eta\left(\boldsymbol{x}_{0:T}^w\right)}{p_{\text{ref}}\left(\boldsymbol{x}_{0:T}^\omega\right)} - \log\frac{p_\eta\left(\boldsymbol{x}_{0:T}^l\right)}{p_{\text{ref}}\left(\boldsymbol{x}_{0:T}^l\right)}\right]\right)\end{aligned} \quad (9)$$

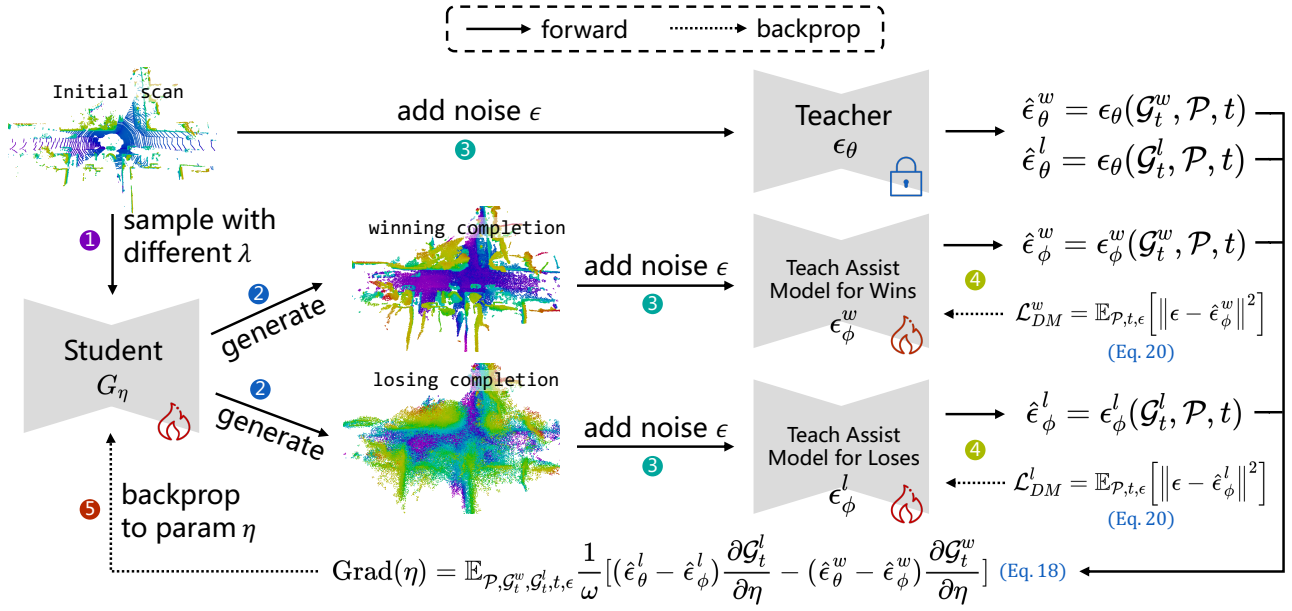

Figure 2: The overall structure of Distillation-DPO. ①The student model generates the completed scenes with different initial noise levels λ based on the sparse scan. ②Choosing the winning sample $\mathcal{G}^w_t$ and losing samples $\mathcal{G}^l_t$. ③The sparse scan $\mathcal{P}$, the noisy completed scenes $\mathcal{G}^w_t$ and $\mathcal{G}^l_t$ are input to $\boldsymbol{\epsilon}_\theta$, $\boldsymbol{\epsilon}^w_\phi$ and $\boldsymbol{\epsilon}^l_\phi$, respectively, to get the predicted noise. ④The model $\boldsymbol{\epsilon}^w_\theta$ and $\boldsymbol{\epsilon}^l_\theta$ are optimized on $\mathcal{G}^w_t$ and $\mathcal{G}^l_t$, separately. ⑤The student model is optimized by the DPO gradient.

Here prompt $\boldsymbol{c}$ is omitted for compactness. By approximating the reverse process $p_\eta(\boldsymbol{x}_{1:T}|\boldsymbol{x}_0)$ with the forward process $q(\boldsymbol{x}_{1:T}|\boldsymbol{x}_0)$, with some simplification, we have:

$$\begin{aligned} L(\eta) = &- \mathbb{E}_{(\boldsymbol{x}^w_0,\boldsymbol{x}^l_0),t,\boldsymbol{x}^w_t,\boldsymbol{x}^l_t} \log \sigma(-\beta T \omega(\lambda_t) \\ &(\|\boldsymbol{\epsilon}^w - \boldsymbol{\epsilon}_\eta(\boldsymbol{x}^w_t, t)\|^2_2 - \|\boldsymbol{\epsilon}^w - \boldsymbol{\epsilon}_{\mathrm{ref}}(\boldsymbol{x}^w_t, t)\|^2_2 \\ &- (\|\boldsymbol{\epsilon}^l - \boldsymbol{\epsilon}_\eta(\boldsymbol{x}^l_t, t)\|^2_2 - \|\boldsymbol{\epsilon}^l - \boldsymbol{\epsilon}_{\mathrm{ref}}(\boldsymbol{x}^l_t, t)\|^2_2))) \end{aligned} \tag{10}$$

Here $\boldsymbol{x}^i_t = \alpha_t \boldsymbol{x}^i_0 + \sigma_t \boldsymbol{\epsilon}^i \quad (i \in \{w, l\})$, $\lambda_t = \frac{\alpha^2_t}{\sigma^2_t}$ is the signal-noise ratio, and $\omega(\lambda_t)$ is the weighted function.

Method

In this section, we introduce the proposed Distillation-DPO. Distillation-DPO aims to use pair-wise winning and losing samples to distill a pre-trained LiDAR scene completion diffusion model into a student model with significantly fewer sampling steps. This allows the student model to perform fast scene completion while encouraging its outputs to resemble the winning samples and diverge from the losing ones, thereby enhancing the overall completion quality. The structure of Distillation-DPO is shown in Fig. 2.

Given a spase scan $\mathcal{P} = \{\boldsymbol{p}^1, \boldsymbol{p}^2, ..., \boldsymbol{p}^N\}$ and the completed scene $\mathcal{G}_0 = \{\boldsymbol{p}^1_0, \boldsymbol{p}^2_0, ..., \boldsymbol{p}^M_0\}$, we rewrite the optimization in Eq. 3 as

$$\min_\eta D_{KL}(p_\eta(\mathcal{G}_0|\mathcal{P}))||p_\theta(\mathcal{G}_0|\mathcal{P})) - \omega \mathbb{E}_{\mathcal{P},\mathcal{G}_0}\left[r(\mathcal{G}_0, \mathcal{P})\right] \tag{11}$$

Here p_θ is the pre-trained distribution of the teacher model parameterized by θ, and p_η is the generative distribution of G_{stu} parameterized by η. The completed LiDAR scene $\mathcal{G}_0$ is generated by G_{stu} with fewer inference steps based on the sparse LiDAR scan $\mathcal{P}$. Eq. 11 performs the distillation and simultaneously encourages the distillated student model to maximize the reward function, resulting in improved efficiency and quality. However, directly optimizing Eq. 11 is challenging. This is because the high-density regions of $p_\eta(\mathcal{G}_0|\mathcal{P})$ are sparse in high-dimensional spaces (Wang et al. 2023) and the computation is intractable. According to Theorem 1 in Wang et al. (2023), we extend Eq. 11 into a noisy optimization over different time steps,

$$\min_\eta D_{KL}(p_{\eta,t}(\mathcal{G}_t|\mathcal{P}))||p_{\theta,t}(\mathcal{G}_t|\mathcal{P})) - \omega \mathbb{E}_{\mathcal{P},t,\boldsymbol{\epsilon}}\left[r(\mathcal{G}_t, \mathcal{P})\right] \tag{12}$$

Here $\boldsymbol{\epsilon}$ is a random noise. $p_{\eta,t}$ and $p_{\theta,t}$ are the noisy distributions of the student model and the teacher model at timestep t, respectively. The parameter ω is the weight to control preference learning. Noisy completed LiDAR scene $\mathcal{G}_t = \{\boldsymbol{p}^1_t, \boldsymbol{p}^2_t, ..., \boldsymbol{p}^M_t\}$ is obtained with the point-level noise addition operation in Eq. 1. Using Eq. 12 and some algebra, the optimization problem can be written as

$$\min_\eta \mathbb{E}_{\mathcal{P},\mathcal{G}_0,\boldsymbol{\epsilon},t}[\log \frac{p_{\eta,t}(\mathcal{G}_t|\mathcal{P})}{p_{\theta,t}(\mathcal{G}_t|\mathcal{P})} - \omega r(\mathcal{G}_t, \mathcal{P})] \tag{13}$$

For Eq. 13, the global optimal solution p^*_η is

$$\begin{aligned} p^*_{\eta,t}(\mathcal{G}_t|\mathcal{P}) &= \frac{p_{\theta,t}(\mathcal{G}_t|\mathcal{P}) \exp(\omega r(\mathcal{G}_t, \mathcal{P}))}{Z(\mathcal{P})} \\ Z(\mathcal{P}) &= \mathbb{E}_{\mathcal{G}_0,t,\epsilon} p_{\theta,t}(\mathcal{G}_t|\mathcal{P}) \exp(\omega r(\mathcal{G}_t, \mathcal{P})) \end{aligned} \tag{14}$$

Therefore, the reward function takes the form:

$$r(\mathcal{G}_0, \mathcal{P}) = \frac{1}{\omega} \log \frac{p_{\eta,t}(\mathcal{G}_t|\mathcal{P})}{p_{\theta,t}(\mathcal{G}_t|\mathcal{P})} + \frac{1}{\omega} \log Z(\mathcal{P}) \tag{15}$$

Then, based on the Bradley-Terry model (Bradley and Terry 1952; Wallace et al. 2024), the objective of Distillation-DPO is

$$\min_{\eta} \mathbb{E}_{\mathcal{P},\mathcal{G}_t^w,\mathcal{G}_t^l,t,\boldsymbol{\epsilon}} \frac{1}{\omega}[\log \frac{p_{\eta,t}(\mathcal{G}_t^l|\mathcal{P})}{p_{\theta,t}(\mathcal{G}_t^l|\mathcal{P})} - \log \frac{p_{\eta,t}(\mathcal{G}_t^w|\mathcal{P})}{p_{\theta,t}(\mathcal{G}_t^w|\mathcal{P})}] \tag{16}$$

Similarly, $\mathcal{G}_t^w$ and $\mathcal{G}_t^l$ represent the completed scenes by the student model G_{stu} with completion quality $\mathcal{G}_t^w \succ \mathcal{G}_t^l$. The gradient of G_{stu} can be calculated as

$$\begin{aligned}\text{Grad}(\eta) = \mathbb{E}_{\mathcal{P},\mathcal{G}_t^w,\mathcal{G}_t^l,t,\boldsymbol{\epsilon}} \\ \frac{1}{\omega}[(\nabla_{\mathcal{G}_t^l} \log p_{\eta,t}(\mathcal{G}_t^l|\mathcal{P}) - \nabla_{\mathcal{G}_t^l} \log p_{\theta,t}(\mathcal{G}_t^l|\mathcal{P}))\frac{\partial \mathcal{G}_t^l}{\partial \eta} - \\ (\nabla_{\mathcal{G}_t^w} \log p_{\eta,t}(\mathcal{G}_t^w|\mathcal{P}) - \nabla_{\mathcal{G}_t^w} \log p_{\theta,t}(\mathcal{G}_t^w|\mathcal{P}))\frac{\partial \mathcal{G}_t^w}{\partial \eta}]\end{aligned} \tag{17}$$

The score $\nabla_{\mathcal{G}_t^l} \log p_{\theta,t}(\mathcal{G}_t^w|\mathcal{P})$ and $\nabla_{\mathcal{G}_t^l} \log p_{\theta,t}(\mathcal{G}_t^l|\mathcal{P})$ are approximated by the teacher diffusion model $\boldsymbol{\epsilon}_\theta$. Differently, the score $\nabla_{\mathcal{G}_t^w} \log p_{\eta,t}(\mathcal{G}_t^w|\mathcal{P})$ and $\nabla_{\mathcal{G}_t^l} \log p_{\eta,t}(\mathcal{G}_t^w|\mathcal{P})$ are approximated by two auxiliary models $\boldsymbol{\epsilon}_\phi^w$ and $\boldsymbol{\epsilon}_\phi^l$. $\frac{\partial \mathcal{G}_t^l}{\partial \eta}$ and $\frac{\partial \mathcal{G}_t^w}{\partial \eta}$ are calculated with automatic differentiation on G_{stu}. Therefore, the approximated gradient of G_{stu} is

$$\begin{aligned}&\text{Grad}(\eta) \\ &\propto \mathbb{E}_{\mathcal{P},\mathcal{G}_t^w,\mathcal{G}_t^l,t,\boldsymbol{\epsilon}} \frac{1}{\omega}[(\boldsymbol{\epsilon}_\theta(\mathcal{G}_t^l,t,\mathcal{P}) - \boldsymbol{\epsilon}_\phi^l(\mathcal{G}_t^l,t,\mathcal{P}))\frac{\partial \mathcal{G}_t^l}{\partial \eta} \\ &-(\boldsymbol{\epsilon}_\theta(\mathcal{G}_t^w,t,\mathcal{P}) - \boldsymbol{\epsilon}_\phi^w(\mathcal{G}_t^w,t,\mathcal{P}))\frac{\partial \mathcal{G}_t^w}{\partial \eta}]\end{aligned} \tag{18}$$

To generate winning and losing scenes $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$, we first introduce a parameter λ when computing $\mathcal{G}_T$, which controls the initial noise scale by

$$\boldsymbol{p}_T^m = \boldsymbol{p}^m + \lambda\sqrt{1-\bar{\alpha}_T}\boldsymbol{\epsilon}_{\boldsymbol{T}} \tag{19}$$

By default, $\lambda = 1$ as Eq. 1. To generate completed scene $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$ separately based on the same sparse scan $\mathcal{P}$, we obtain different completion results by using different noise and adjusting different values of λ. We set $\lambda = 1.1$ to obtain a $\mathcal{G}_T'$ different from $\mathcal{G}_T$, which is then used to generate $\mathcal{G}_0'$ different from $\mathcal{G}_0$. Then, according to the completion quality metrics such as CD or JSD, we assign the sample with the higher quality as $\mathcal{G}_0^w$ and another as $\mathcal{G}_0^l$.

During the training process, the student model G_{stu} and two auxiliary models $\boldsymbol{\epsilon}_\phi^w$ and $\boldsymbol{\epsilon}_\phi^l$ are optimized alternately. The auxiliary models $\boldsymbol{\epsilon}_\phi^w$ and $\boldsymbol{\epsilon}_\phi^l$ are trained on the completed scene generated by G_{stu} with the standard diffusion objective (Ho, Jain, and Abbeel 2020)

$$\mathcal{L}_{DM}^i = \mathbb{E}_{\mathcal{P},t,\boldsymbol{\epsilon}}\left[\left\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\phi^i\left(\mathcal{G}_t^i,\mathcal{P},t\right)\right\|^2\right] \quad i \in \{w,l\} \tag{20}$$

Experiments

Model and datasets We use the SOTA 3D LiDAR scene completion diffusion model LiDiff (Nunes et al. 2024) as the teacher and train a few-step student model with Eq. 18. LiDiff can complete a scene with 50 sampling steps based on the sparse LiDAR scan. The student model G_{stu} and the auxiliary models $\boldsymbol{\epsilon}_\phi^w$ and $\boldsymbol{\epsilon}_\phi^l$ are initialized with the pre-trained LiDiff model, but the student model performs scene completion with fewer sampling steps. The experiments are conducted on SemanticKITTI (Behley et al. 2019) and KITTI-360 (Liao, Xie, and Geiger 2022) datasets.

Model	CD ↓	JSD ↓	EMD ↓	Times (s) ↓
LMSCNet	0.641	0.431	-	**0.31**
LODE	1.029	0.451	-	<u>0.58</u>
MID	0.503	0.470	-	4.94
PVD	1.256	0.498	-	145.39
LiDiff	0.434	0.444	**22.15**	17.75
LiDiff (Refined)	<u>0.375</u>	<u>0.416</u>	<u>23.16</u>	17.87
Distillation-DPO	0.414	0.419	23.29	3.28
Distillation-DPO (Refined)	**0.354**	**0.387**	23.66	3.38

Table 1: The completion performances of Distillation-DPO with existing models on SemanticKITTI dataset. Bold and underlined represent the optimal and suboptimal performance, respectively. The completion time is calculated based on the official implementation and released checkpoints. Here LiDiff takes 50 NFEs while ours takes 8 only.

Model	CD ↓	JSD ↓	EMD ↓	Times (s) ↓
LMSCNet	0.979	0.496	-	0.29
LODE	1.565	0.483	-	0.55
MID	0.637	0.476	-	4.61
LiDiff	0.564	0.459	**21.98**	17.01
LiDiff (Refined)	<u>0.517</u>	0.446	<u>22.96</u>	17.12
Distillation-DPO	0.533	<u>0.434</u>	22.99	2.97
Distillation-DPO (Refined)	**0.497**	**0.422**	23.45	3.09

Table 2: Completion performances on KITTI-360. All presentation meaning are the same as Table 1.

Baselines and metrics Except for the existing SOTA LiDAR scene completion diffusion model LiDiff (Nunes et al. 2024), we also choose LMSCNet (Roldao, de Charette, and Verroust-Blondet 2020), LODE (Li et al. 2023), MID (Vizzo et al. 2022) and PVD (Zhou, Du, and Wu 2021) as the baselines. We evaluate the performance of the proposed Distillation-DPO on Chamfer Distance (CD) (Butt and Maragos 1998), Jensen-Shannon Divergence (JSD) (Menéndez et al. 1997) and Earth Mover's Distance (EMD) (Fan, Su, and Guibas 2017). These three metrics can provide a comprehensive evaluation of the completed LiDAR scene quality from different perspectives.

Evaluation on LiDAR Scene Completion

We first compared the performance of the proposed Distillation-DPO and existing models in LiDAR scene completion on the SemanticKITTI and KITTI-360 datasets. According to different settings, Distillation-DPO can perform sampling with different inference steps. To achieve an optimal trade-off between sampling quality and speed, we chose the result with 8 sampling steps as the completion output of Distillation-DPO for comparison with existing models. In Sec. , we further compare the performance of Distillation-

Model	NFE ↓	CD ↓	JSD ↓	EMD ↓	Time (s) ↓
LiDiff	50	0.434	0.444	**22.15**	17.75
LiDiff (Refined)	50	0.375	0.416	<u>23.16</u>	17.87
LiDiff	8	0.447	0.432	24.90	3.35
LiDiff (Refined)	8	0.411	0.406	25.74	3.48
Distillation-DPO (Refined)	8	**0.354**	**0.387**	23.66	3.38
Distillation-DPO (Refined)	4	<u>0.369</u>	<u>0.389</u>	24.56	1.97
Distillation-DPO (Refined)	2	0.401	0.390	26.01	<u>1.22</u>
Distillation-DPO (Refined)	1	0.601	0.437	28.01	**0.82**

Table 3: Comparison results of different inference steps on SemanticKITTI dataset.

Model	CD ↓	JSD ↓	EMD ↓
$\lambda = 1.1$ (ours)	**0.354**	**0.387**	**23.66**
$\lambda = 1.05$	<u>0.358</u>	<u>0.391</u>	23.74
$\lambda = 1.2$	0.359	0.404	23.93
$\lambda = 1.5$	0.409	0.422	<u>23.71</u>
$\lambda = 2.0$	0.382	0.421	24.10

Table 4: Comparison results of different λ value on SemanticKITTI dataset. All results have been refined.

DPO under different sampling steps.

The comparison results of Distillation-DPO on the SemanticKITTI dataset are shown in Tab. 1. Distillation-DPO achieves the optimal completion quality except in EMD. Compared with the SOTA LiDAR scene completion method LiDiff (Nunes et al. 2024), Distillation-DPO accelerates the completion speed by over 5 times while achieving improvements of 6% and 7% in CD and JSD. As for EMD, Distillation-DPO incurs less than a 5% performance loss in EMD compared to LiDiff without refinement, and the loss is further reduced to under 2% with refinement. This is because the optimization objective of DPO encourages the model to focus more on generating locally accurate points, potentially at the expense of global point cloud uniformity. However, in autonomous driving scenarios, the reconstruction of fine-grained objects on the road is often more critical than preserving coarse global structures. Moreover, although LMSCNet (Roldao, de Charette, and Verroust-Blondet 2020) and LODE (Li et al. 2023) offer faster completion speeds, their completion quality is significantly lower, which is unacceptable for safety-critical autonomous driving applications. The results on KITTI-360 dataset are shown in Tab. 2. Distillation-DPO also outperforms LiDiff by achieving the best scene completion (4% improvement in CD and 6% in JSD) quality with a 5× speedup.

Ablation Study

We conduct a series of ablation studies on the SemanticKITTI dataset to prove the effectiveness of Distillation-DPO. We first show the performance of Distillation-DPO with different inference steps in Tab. 3. As the number of inference steps decreases, the completion time of Distillation-DPO is further reduced. With just one sampling step, it only takes 0.82 seconds to complete a scene. However, the reduction in inference steps leads to a decline in completion quality. The speed improvement gained from sampling step reduction is not enough to compensate for the loss in quality. Therefore, choosing 8 steps by default is a good balance of speed and quality.

Model	CD ↓	JSD ↓	EMD ↓
LiDiff	0.375	0.416	23.16
LiDiff*	0.368	0.401	**22.69**
Distillation-DPO	0.354	0.387	23.66
Distillation-DPO*	**0.343**	**0.385**	23.53

Table 5: Comparison results of using different teacher models. LiDiff* represents the LiDiff model refined with Diffusion-DPO with boosted performance. Distillation-DPO* represents the Distillation-DPO trained with LiDiff*.

Model	CD ↓	JSD ↓	EMD ↓
Distillation-DPO (CD)	**0.354**	**0.387**	**23.66**
Distillation-DPO (JSD)	0.382	0.405	25.29

Table 6: Comparison results of using different metrics to determine $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$. All results have been refined.

Then, we further compare the completion quality of different values of λ ($\lambda = 1.1$ by default). As shown in Tab. 4, when decreasing or increasing λ, the completion performance deteriorates. When λ is small, the difference between $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$ is minimal, making the gradients of the student model in Eq. 18 small, which leads to unstable training. Conversely, when λ is large, the quality of $\mathcal{G}_0'$ generated from $\mathcal{G}_T'$ degrades significantly, causing it to fall outside the distribution learned by the pre-trained teacher model $\boldsymbol{\epsilon}_\theta$. This mismatch leads to inaccurate predictions from $\boldsymbol{\epsilon}_\theta$ (Zhang et al. 2024), resulting in incorrect gradients for the student model and ultimately lowering the completion quality.

We also explored the impact of teacher model performances on the effectiveness of Distillation-DPO. Theoretically, the final performance of the student model is constrained by the teacher model. The better the performance of the teacher model is, the better the final performance of the student model would be. Thus, we first fine-tuned LiDiff (Nunes et al. 2024) using DiffusionDPO (Wallace et al. 2024) to enhance its performance. Then, we retrained Distillation-DPO using the fine-tuned model. Results in Tab. 5 display that as the performance of the teacher model improves, the student model performance also improves.

Moreover, we conduct experiments by changing the evaluation metric for determining $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$ to JSD. The results in Tab. 6 show that the performance significantly deteriorates when using JSD. Since JSD measures the similarity of point cloud distributions, it requires a large number of samples to estimate the probability density distribution accurately. However, when comparing and determining whether a sample is $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$, the metric is computed using only a single generated sample and its corresponding ground truth. In this case, JSD becomes inaccurate and may even lose its practical significance, leading to a performance decline.

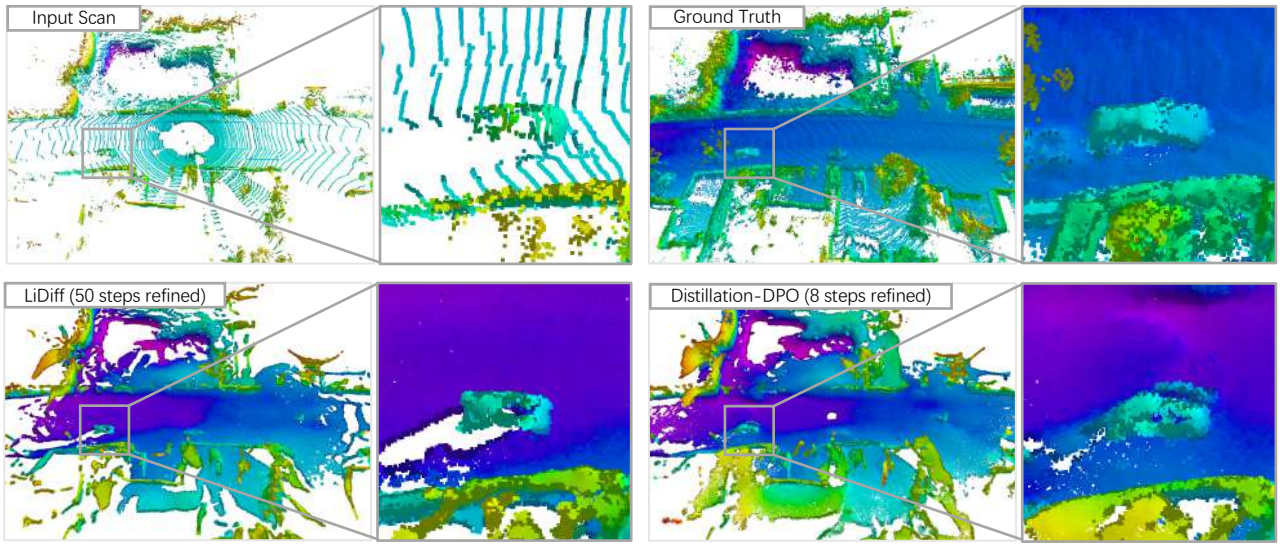

Figure 3: Qualitative results on SemanticKITTI dataset.

Model	CD ↓	JSD ↓	EMD ↓
Distillation-DPO	**0.354**	**0.387**	**23.66**
One auxiliary diffusion model	0.372	0.418	24.35

Table 7: Ablation study of the auxiliary diffusion models. All results have been refined.

Model	NFE ↓	CD ↓	JSD ↓	EMD ↓
LiDiff	50	0.375	0.416	**23.16**
Score Distillation	8	0.419	0.430	24.61
Distillation-DPO	8	**0.354**	**0.387**	23.66

Table 8: Comparison between Distillation-DPO and traditional score distillation. All results have been refined.

EMD and IoU can also be used as the evaluation metrics. However, their computation is expensive, which negatively impacts training efficiency.

In the default setting, Distillation-DPO employs two auxiliary diffusion models ϵ_ϕ^w and ϵ_ϕ^l to independently approximate the distributions of $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$. Tab. 7 compares this setup with a variant that uses a single auxiliary diffusion model to approximate both distributions simultaneously. The results show that using only one auxiliary model leads to degraded performance of the student model. This is primarily because $\mathcal{G}_0^w$ and $\mathcal{G}_0^l$ follow different distributions, and a single model fails to accurately capture both, resulting in inaccurate gradient signals being propagated to the student model and ultimately impairing its performance.

Finally, we compare Distillation-DPO with traditional score distillation methods as in Eq. 6 to validate the effectiveness of the proposed distillation framework. Tab. 8 shows that the results obtained using score distillation are even inferior to those of the original teacher model LiDiff (Nunes et al. 2024). This is consistent with the statement in Sec. that directly employing score distillation can accelerate the sampling speed while inevitably decreasing the performance. In contrast, the proposed Distillation-DPO distillation framework incorporates guidance from pair-wise winning and losing samples, which not only accelerates sampling but also further enhances completion quality, thereby achieving efficient and high-quality scene completion.

More ablation study results on KITTI-360 dataset are shown in Sec.5 of Supplementary Materials (Sec.S5).

Qualitative Comparison

We visualized the scene completion results of Distillation-DPO and compared them with those of the SOTA model LiDiff (Nunes et al. 2024), as shown in Fig. 3. Compared to LiDiff, Distillation-DPO achieves higher scene completion quality with only 8 sampling steps, surpassing LiDiff's results even with 50 sampling steps. Moreover, Distillation-DPO provides more complete reconstructions of fine details as shown in the zoomed-in region of Fig. 3.

Conclusion

This paper proposes a novel LiDAR scene completion diffusion model distillation framework, Distillation-DPO. Distillation-DPO introduces the constraints of pair-wise winning and losing samples into the score distillation strategy, enabling effective distillation of LiDAR scene completion diffusion models. Compared to existing models, Distillation-DPO achieves new SOTA completion performance while improving completion speed more than five times over existing SOTA models. To our best knowledge, we are the first to integrate distillation and post-training with preference and provide insight into preference-aligned diffusion distillation for both areas of LiDAR scene completion and visual generation.

Acknowledgements

This paper is supported by Provincial Key Research and Development Plan of Zhejiang Province under No. 2024C01250(SD2), National Natural Science Foundation of China (Grant No. 62576306) and Dream Set Off - Kunpeng & Ascend Seed Program.

References

Back, K.; Piao, X.; and Kim, J.-K. 2024. Enhancing Reinforcement Learning Finetuned Text-to-Image Generative Model Using Reward Ensemble. In *International Conference on Intelligent Tutoring Systems*, 213–224.

Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C.; and Gall, J. 2019. Semantickitti: A Dataset For Semantic Scene Understanding Of Lidar Sequences. In *ICCV*, 9297–9307.

Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4): 324–345.

Butt, M. A.; and Maragos, P. 1998. Optimum Design of Chamfer Distance Transforms. *IEEE Transactions on Image Processing*, 7(10): 1477–1484.

Cao, H.; and Behnke, S. 2024. DiffSSC: Semantic LiDAR Scan Completion Using Denoising Diffusion Probabilistic Models. *arXiv preprint arXiv:2409.18092*.

Chen, Y.; Yang, B.; Liang, M.; and Urtasun, R. 2019. Learning Joint 2d-3d Representations for Depth Completion. In *ICCV*, 10023–10032.

Clark, K.; Vicol, P.; Swersky, K.; and Fleet, D. J. 2024. Directly Fine-Tuning Diffusion Models on Differentiable Rewards. In *ICLR*.

Fan, H.; Su, H.; and Guibas, L. J. 2017. A Point Set Generation Network for 3d Object Reconstruction From a Single Image. In *CVPR*, 605–613.

Guo, J.; Chai, W.; Deng, J.; Huang, H.-W.; Ye, T.; Xu, Y.; Zhang, J.; Hwang, J.-N.; and Wang, G. 2024. Versat2i: Improving Text-to-Image Models With Versatile Reward. *arXiv preprint arXiv:2403.18493*.

He, Y.; Wang, H.; Jiang, Z.; Papangelis, A.; and Zhao, H. 2024. Semi-Supervised Reward Modeling Via Iterative Self-Training. *arXiv preprint arXiv:2409.06903*.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *NeurIPS*, volume 33, 6840–6851.

Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *NeurIPS*, 35: 26565–26577.

Li, P.; Zhao, R.; Shi, Y.; Zhao, H.; Yuan, J.; Zhou, G.; and Zhang, Y.-Q. 2023. Lode: Locally Conditioned Eikonal Implicit Scene Completion From Sparse Lidar. In *IEEE Intl. Conf. on Robotics and Automation*, 8269–8276.

Liang, Y.; He, J.; Li, G.; Li, P.; Klimovskiy, A.; Carolan, N.; Sun, J.; Pont-Tuset, J.; Young, S.; Yang, F.; et al. 2024. Rich Human Feedback for Text-to-Image Generation. In *CVPR*, 19401–19411.

Liao, Y.; Xie, J.; and Geiger, A. 2022. Kitti-360: A Novel Dataset And Benchmarks For Urban Scene understanding in 2d And 3d. *IEEE TPAMI*, 45(3): 3292–3310.

Luhman, E.; and Luhman, T. 2021. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*.

Luo, W.; Hu, T.; Zhang, S.; Sun, J.; Li, Z.; and Zhang, Z. 2023. Diff-Instruct: A Universal Approach for Transferring Knowledge From Pre-trained Diffusion Models. In *NeurIPS*, volume 36, 76525–76546.

Mahan, D.; Van Phung, D.; Rafailov, R.; Blagden, C.; Lile, N.; Castricato, L.; Fränken, J.-P.; Finn, C.; and Albalak, A. 2024. Generative Reward Models. *arXiv preprint arXiv:2410.12832*.

Menéndez, M. L.; Pardo, J.; Pardo, L.; and Pardo, M. 1997. The Jensen-Shannon Divergence. *Journal of the Franklin Institute*, 334(2): 307–318.

Nakashima, K.; and Kurazume, R. 2024. LiDAR Data Synthesis With Denoising Diffusion Probabilistic Models. In *IEEE Intl. Conf. on Robotics and Automation*, 14724–14731. IEEE.

Nunes, L.; Marcuzzi, R.; Mersch, B.; Behley, J.; and Stachniss, C. 2024. Scaling Diffusion Models To Real-World 3D LiDAR Scene Completion. In *CVPR*, 14770–14780.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training Language Models To Follow Instructions with Human Feedback. *NeurIPS*, 35: 27730–27744.

Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. *NeurIPS*, 36: 53728–53741.

Ran, H.; Guizilini, V.; and Wang, Y. 2024. Towards Realistic Scene Generation With LiDAR Diffusion Models. In *CVPR*, 14738–14748.

Roldao, L.; de Charette, R.; and Verroust-Blondet, A. 2020. Lmscnet: Lightweight Multiscale 3d Semantic Completion. In *2020 International Conference on 3D Vision (3DV)*, 111–119.

Salimans, T.; and Ho, J. 2021. Progressive Distillation for Fast Sampling of Diffusion Models. In *ICLR*.

Song, S.; Yu, F.; Zeng, A.; Chang, A. X.; Savva, M.; and Funkhouser, T. 2017. Semantic Scene Completion From A Single Depth Image. In *CVPR*, 1746–1754.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-based generative modeling through stochastic differential equations. In *ICLR*.

Vizzo, I.; Mersch, B.; Marcuzzi, R.; Wiesmann, L.; Behley, J.; and Stachniss, C. 2022. Make It Dense: Self-Supervised Geometric Scan Completion of Sparse 3d Lidar Scans In Large Outdoor Environments. *IEEE Robotics and Automation Letters*, 7(3): 8534–8541.

Wallace, B.; Dang, M.; Rafailov, R.; Zhou, L.; Lou, A.; Purushwalkam, S.; Ermon, S.; Xiong, C.; Joty, S.; and Naik, N. 2024. Diffusion Model Alignment Using Direct Preference Optimization. In *CVPR*, 8228–8238.

Wang, Z.; Lu, C.; Wang, Y.; Bao, F.; Li, C.; Su, H.; and Zhu, J. 2023. ProlificDreamer: High-Fidelity And Diverse Text-to-3D Generation With Variational Score Distillation. In *NeurIPS*, volume 36, 8406–8441.

Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2023. Imagereward: Learning And Evaluating Human Preferences For Text-to-Image Generation. *NeurIPS*, 36: 15903–15935.

Xu, Y.; Zhu, X.; Shi, J.; Zhang, G.; Bao, H.; and Li, H. 2019. Depth Completion from Sparse Lidar Data With Depth-Normal Constraints. In *ICCV*, 2811–2820.

Yan, T.; Yin, J.; Lang, X.; Yang, R.; Xu, C.-Z.; and Shen, J. 2024. OLiDM: Object-Aware LiDAR Diffusion Models for Autonomous Driving. *arXiv preprint arXiv:2412.17226*.

Yin, T.; Gharbi, M.; Park, T.; Zhang, R.; Shechtman, E.; Durand, F.; and Freeman, W. T. 2024a. Improved Distribution Matching Distillation for Fast Image Synthesis. *arXiv preprint arXiv:2405.14867*.

Yin, T.; Gharbi, M.; Zhang, R.; Shechtman, E.; Durand, F.; Freeman, W. T.; and Park, T. 2024b. One-step Diffusion With Distribution Matching Distillation. In *CVPR*, 6613–6623.

Zhang, S.; Zhao, A.; Yang, L.; Li, Z.; Meng, C.; Xu, H.; Chen, T.; Wei, A.; GU, P. P.; and Sun, L. 2024. Distilling Diffusion Models to Efficient 3D LiDAR Scene Completion. *arXiv preprint arXiv:2412.03515*.

Zhang, X.; Yang, L.; Li, G.; Cai, Y.; Xie, J.; Tang, Y.; Yang, Y.; Wang, M.; and Cui, B. 2025. Itercomp: Iterative Composition-aware Feedback Learning From Model Gallery For Text-to-Image Generation. In *ICLR*.

Zhou, L.; Du, Y.; and Wu, J. 2021. 3d Shape Generation And Completion Through Point-voxel Diffusion. In *ICCV*, 5826–5835.

Zhou, S.; Yang, P.; Wang, J.; Luo, Y.; and Loy, C. C. 2024. Upscale-a-Video: Temporal-Consistent Diffusion Model For Real-World Video Super-Resolution. In *CVPR*, 2535–2545.