

# 000 001 002 003 004 005 FINE-TUNING DIFFUSION MODELS VIA INTERMEDIATE 006 DISTRIBUTION SHAPING 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029

## ABSTRACT

030 Diffusion models are widely used for generative tasks across domains. While  
031 pre-trained diffusion models effectively capture the training data distribution, it is  
032 often desirable to shape these distributions using reward functions to align with  
033 downstream applications. Policy gradient methods, such as Proximal Policy Opti-  
034 mization (PPO), are widely used in the context of autoregressive generation.  
035 However, the marginal likelihoods required for such methods are intractable for  
036 diffusion models, leading to alternative proposals and relaxations. In this context,  
037 we unify variants of Rejection sAmpling based Fine-Tuning (RAFT) as GRAFT,  
038 and show that this induces same solution as PPO objective with reshaped rewards.  
039 We then introduce P-GRAFT to shape distributions at intermediate noise levels  
040 and demonstrate empirically that this can lead to more effective fine-tuning. We  
041 mathematically explain this via a bias-variance tradeoff. Motivated by this, we  
042 propose inverse noise correction to improve flow models without leveraging explicit  
043 rewards. We empirically evaluate our methods on text-to-image(T2I) generation,  
044 layout generation, molecule generation and unconditional image generation. Not-  
045 ably, our framework, applied to Stable Diffusion v2, improves over policy gradient  
046 methods on popular T2I benchmarks in terms of VQAScore and shows an 8.81%  
047 relative improvement over the base model. For unconditional image generation,  
048 inverse noise correction improves FID of generated images at lower FLOPs/image.  
049

## 1 INTRODUCTION

050 Pre-trained generative models often require task-specific adaptations based on reward feedback  
051 - a standard strategy is to leverage RL algorithms, such as Proximal Policy Optimization (PPO)  
052 (Schulman et al., 2017). While such methods have found great success in the context of language  
053 modeling (Bai et al., 2022; Ouyang et al., 2022), their adoption to diffusion models is not straight-  
054 forward. In particular, unlike autoregressive (AR) models, marginal likelihoods required for the  
055 implementation of KL regularization in PPO are *intractable* for diffusion models. **Mathematically,**  
056 **suppose  $p(\cdot)$  is the distribution induced by a generative model. Given a sample (not necessarily**  
057 **generated by the model), we can compute the likelihood of  $x = (x^1, \dots, x^n)$  for AR models as**  
058  **$p(x^1)p(x^2|x^1)\dots p(x^n|x^1, \dots, x^{n-1})$ . This is because AR models predict the next token distribution**  
059 **conditioned on the previous tokens. This likelihood can directly be used for RL algorithms. Diffusion**  
060 **models, on the other hand, model the conditional likelihoods along the denoising time axis, i.e. we**  
061 **only have access to  $p(x_{t-1}|x_t)$ , the conditional distribution of the denoised state given the previous**  
062 **state. From these conditionals, it is intractable to compute the marginal likelihood  $p(x)$ . This makes**  
063 **the marginal likelihood unavailable to implement popular RL algorithms like PPO.** Hence, in practice,  
064 **KL regularization is ignored (Black et al., 2023) or relaxations such as trajectory KL regularization**  
065 **(Fan et al., 2023) is considered. However, ignoring the KL term results in unstable training in**  
066 **large-scale settings (Deng et al., 2024), whereas using the trajectory KL constraint gives subpar**  
067 **results (Black et al., 2023). Further, fine-tuning with trajectory KL also results in the initial value**  
068 **function bias problem (Domingo-Enrich et al., 2024; Uehara et al., 2024).**

069 Apart from policy gradient methods, recent research has also focused on fine-tuning methods based  
070 on *rejection sampling* such as RSO (Liu et al., 2023b), RAFT (Dong et al., 2023) and Reinforce-Rej  
071 (Xiong et al., 2025). Further, fine-tuning based on Best-of-N (BoN) sampling and its relation to  
072 policy gradient methods have also been explored, but in the context of autoregressive models (Amini  
073 et al., 2024; Gui et al., 2024). Given the intractability of **PPO with (marginal) KL regularization as**

054 **described above**, we explore conceptual connections between rejection sampling based fine-tuning  
 055 methods and PPO, specifically in the context of diffusion models. In particular, we make the following  
 056 contributions:

057 **(a)** We conceptualize a Generalized Rejection Sampling (GRS) framework which subsumes various  
 058 rejection sampling strategies including classical rejection sampling from MCMC literature and Best-  
 059 of-N. We show that GRS samples from the solution to PPO but with a reshaped reward - fine-tuning  
 060 using GRS, which we term Generalized Rejection sAmpling Fine-Tuning (GRAFT) enables implicit  
 061 *marginal KL regularization* for diffusion models, despite the *marginal likelihoods being intractable*.  
 062

063 **(b)** Leveraging properties of diffusion models, we propose Partial-GRAFT (P-GRAFT) a framework  
 064 which fine-tunes *only till an intermediate denoising step* by assigning the rewards of final generations  
 065 to partial, noisy generations. We show that this leads to better fine-tuning empirically and provide  
 066 a mathematical justification via a bias-variance tradeoff. Empirically, we demonstrate significant  
 067 quality gains across the tasks of text-to-image generation, layout generation and molecule generation.  
 068

069 **(c)** Motivated by P-GRAFT, we introduce Inverse Noise Correction - an adapter-based, parameter-  
 070 efficient method to improve flow models even *without explicit rewards*. We empirically demonstrate  
 071 improved quality as well as FLOPs for unconditional image generation.  
 072

073 **(d)** In particular, SDv2 fine-tuned using P-GRAFT demonstrates significant improvements in VQAS-  
 074 core over policy-gradient methods as well as SDXL-Base across datasets. The proposed Inverse  
 075 Noise Correction strategy provides significant FID improvement at reduced FLOPs/image.  
 076

077 A more comprehensive list of related work can be found in Appendix A.

## 078 2 PRELIMINARIES

079 **PPO for Generative Modeling:** Following (Stiennon et al., 2020), we introduce PPO in our setting:  
 080 Consider a state space  $\mathcal{X}$ , a reward function  $r : \mathcal{X} \rightarrow \mathbb{R}$  and a reference probability measure  $\bar{p}$  over  
 081  $\mathcal{X}$ . Let  $\mathcal{P}(\mathcal{X})$  be the set of probability measures over  $\mathcal{X}$  and  $\alpha \in (0, \infty)$ . Define  $R^{\text{reg}} : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$   
 082 by  $R^{\text{reg}}(p) = \mathbb{E}_{X \sim p}[r(X)] - \alpha \text{KL}(p \parallel \bar{p})$ , where  $\text{KL}(\cdot \parallel \cdot)$  is the KL divergence. PPO aims to obtain

$$p^{\text{ppo}} = \arg \sup_{p \in \mathcal{P}(\mathcal{X})} R^{\text{reg}}(p). \quad (1)$$

083 Using the method of Lagrangian Multipliers, we can show that  $p^{\text{ppo}}(x) \propto \exp(r(x)/\alpha)\bar{p}(x)$ . In  
 084 generative modeling literature,  $\bar{p}$  is often the law of generated samples from a pre-trained model -  
 085 fine-tuning is done on the model so as to sample from the tilted distribution  $p^{\text{ppo}}$ .  
 086

087 **PPO solution via Rejection Sampling:** Classical rejection sampling from the Monte Carlo literature  
 088 (Thompson, 2012) can be used to sample from  $p^{\text{ppo}}$ . We note this folklore result in our setting:

089 **Lemma 2.1.** *Let  $r(x) \leq r_{\max}$  for some  $r_{\max}$ . Given a sample  $Y \sim \bar{p}$ , we accept it with probability  
 090  $\mathbb{P}(\text{Accept}|Y) = \exp\left(\frac{r(Y) - r_{\max}}{\alpha}\right)$ . Then, conditioned on Accept,  $Y$  is a sample from  $p^{\text{ppo}}$ .*

091 Lemma 2.1 provides a way to obtain exact samples from  $p^{\text{ppo}}$ . A well known challenge with this  
 092 method is sample inefficiency - as often in practice,  $\alpha$  is small leading to small acceptance probability.  
 093 Thus, methods such as Best-of-N (BoN) which always accept a fixed fraction of samples are used.  
 094

095 We now introduce *Generalized Rejection sAmpling Fine Tuning* (GRAFT), a framework to unify  
 096 existing rejection sampling approaches such as BoN and RAFT. Going beyond, GRAFT also admits  
 097 novel rejection sampling strategies such as those involving de-duplication as we describe in the  
 098 next section. Theoretically, Lemma 3.2 shows that GRAFT samples from the solution to the PPO  
 099 objective, but with reshaped rewards. We then discuss its utility in the context of diffusion models.  
 100 Empirically, GRAFT outperforms policy-gradient methods across multiple tasks including text-to-  
 101 image generation.  
 102

## 103 3 GRAFT: GENERALIZED REJECTION SAMPLING FINE TUNING

104 Assume  $(X^{(i)})_{i \in [M]}$  are  $M$  i.i.d. samples with law  $\bar{p}$  over a space  $\mathcal{X}$ . Given reward function  $r : \mathcal{X} \rightarrow$   
 105  $\mathbb{R}$ , let the reward corresponding to  $X^{(i)}$  be  $R_i := r(X^{(i)})$ , the empirical distribution of  $(X^{(i)})_{i \in [M]}$

108 be  $\hat{P}_X(\cdot)$  and the empirical CDF of  $(R_i)_{i \in [M]}$  be  $\hat{F}_R(\cdot)$ . We introduce Generalized Rejection  
 109 Sampling (GRS) to accept a subset of high reward samples,  $\mathcal{A} := (Y^{(j)})_{j \in [M_s]} \subseteq (X^{(i)})_{i \in [M]}$ ,  
 110 where  $Y^{(j)}$  denotes the  $j^{\text{th}}$  accepted sample.  
 111

112 **Definition 3.1. Generalized Rejection Sampling (GRS):** Let the acceptance function  $A : \mathbb{R} \times$   
 113  $[0, 1] \times \mathcal{X} \times [0, 1] \rightarrow [0, 1]$  be such that  $A$  is co-ordinate wise increasing in the first two co-ordinates.  
 114 The acceptance probability of sample  $i$  is  $p_i := A(R_i, \hat{F}_R(R_i), X^{(i)}, \hat{P}_X)$ . Draw  $C_i \sim \text{Ber}(p_i) \forall i \in$   
 115  $\{1, \dots, M\}$ , not necessarily independent of each other. Then,  $X^{(i)} \in \mathcal{A}$  iff  $C_i = 1$ .  
 116

117 Definition 3.1 subsumes popular rejection sampling approaches such as RAFT and BoN. We now  
 118 show that GRS implicitly samples from the solution to PPO with the reshaped reward  $\hat{r}(\cdot)$ :  
 119

120 **Lemma 3.2.** *The law of accepted samples under GRS (Def 3.1) given by  $p(X^{(1)} = x | X^{(1)} \in \mathcal{A})$  is  
 121 the solution to the following Proximal Policy Optimization problem:*

$$122 \arg \max_{\hat{p}} [\mathbb{E}_{x \sim \hat{p}} \hat{r}(x) - \alpha \text{KL}(\hat{p} \| \bar{p})]; \quad \frac{\hat{r}(x)}{\alpha} := \log (\mathbb{E}[A(r(x), \hat{F}_R(r(x)), x, \hat{P}_X) | X^{(1)} = x])$$

123 Here, the expectation is with respect to the randomness in the empirical distributions  $\hat{F}_R$  and  $\hat{P}_X$ .  
 124

125  $\hat{r}(\cdot)$  is monotonically increasing with respect to the actual reward since  $A$  is an increasing function of  
 126 the reward and its empirical CDF. We now instantiate GRS with commonly used variants of  $A$ :  
 127

128 **Top – K Sampling:** Let the reward distribution be continuous with CDF  $F(\cdot)$ . We accept the top  $K$   
 129 samples out of the  $M$  samples based on their reward values.  
 130

$$131 \text{Corresponding Acceptance Function: } A(r, \hat{F}_R, x, \hat{P}_X) = \begin{cases} 0 & \text{if } \hat{F}_R(r) \leq 1 - \frac{K}{M} \\ 1 & \text{if } \hat{F}_R(r) > 1 - \frac{K}{M} \end{cases}$$

132 Lemma 3.2 shows that this acceptance function results in the reshaped reward  $\hat{r} : \frac{\hat{r}(x)}{\alpha} =$   
 133  $\log \left[ \sum_{k=0}^{K-1} \binom{M-1}{k} F(r(x))^{M-k-1} (1 - F(r(x)))^k \right]$ .  
 134

135 **Preference Rewards:** Setting  $M = 2$  and  $K = 1$  in the above formulation gives preference rewards,  
 136 i.e.,  $X^{(1)}$  is accepted and  $X^{(2)}$  is rejected if  $r(X^{(1)}) > r(X^{(2)})$  (and vice versa). This strategy  
 137 results in the reshaped reward  $\frac{\hat{r}(x)}{\alpha} = \log F(r(x))$ . Since  $F$  is an increasing function, the reward  
 138  $r(x)$  is monotonically reshaped to  $\log F(r(x))$ .  
 139

140 Varying  $K$  from 1 to  $M$ , varies the strength of the tilt in Top – K sampling. In particular,  $K = M$   
 141 corresponds to  $\frac{\hat{r}(x)}{\alpha} = 0$  (no tilt) and  $K = 1$  corresponds to  $\frac{\hat{r}(x)}{\alpha} = M \log F(r(x))$ .  
 142

143 **Binary Rewards with De-Duplication:** Suppose  $r(X) \in \{0, 1\}$  (for eg., corresponds to unstable/  
 144 stable molecules in molecule generation). De-duplication of the generated samples might be necessary  
 145 to maintain diversity. Given any structure function  $f$  (for eg., gets the molecule structure from a  
 146 configuration), let  $N_f(X, \hat{P}_X) = |\{i : f(X^{(i)}) = f(X)\}|$ , i.e, the number of copies of  $X$  in the data.  
 147

$$148 \text{Proposed Acceptance Function: } A(r, \hat{F}_R, x, \hat{P}_X) = \begin{cases} 0 & \text{if } r = 0 \\ \frac{1}{N_f(x, \hat{P}_X)} & \text{if } r = 1 \end{cases}$$

149 Draw  $C_i \sim \text{Ber}(p_i)$  without-replacement among the duplicate/similar samples (i.e, they are  
 150 marginally Bernoulli but are not independent). Thus, exactly one out of the duplicate molecules are  
 151 selected almost surely. Applying Lemma 3.2, we conclude that:  
 152

$$153 \frac{\hat{r}(x)}{\alpha} = \begin{cases} -\infty & \text{if } r(x) = 0 \\ \log \mathbb{E} \left[ \frac{1}{N_f(x, \hat{P}_X)} | X^{(1)} = x \right] & \text{if } r(x) = 1 \end{cases}$$

155 We see that the shaped reward increases with diversity and with the value of the original reward. We  
 156 use this in the molecule generation experiments to avoid mode collapse (Section 6.2).  
 157

158 **Implications for diffusion models:** While specialized versions of Lemma 3.2 are known in the  
 159 context of AR models (Amini et al., 2024), the result is particularly useful in the context of diffusion  
 160 models. Note that given a sample  $x$  along with a prompt  $y$ , the *marginal* likelihood  $\bar{p}(x|y)$  can be  
 161 easily computed for AR models. For diffusion models, we *only* have access to *conditional* likelihoods  
 162 along the denoising trajectory of the diffusion process whereas  $\text{KL}(p||\bar{p})$  is intractable. That is,

162 if the denoising process is run from  $t_N$  to  $t_0$ , we have access to  $\bar{p}(x_{t_i} | x_{t_{i+1}})$ . A commonly used  
 163 relaxation is the trajectory KL,  $\text{KL}(p(X_{0:T}) || \bar{p}(X_{0:T}))$ , which can be shown as an upper bound on  
 164 the marginal KL. As discussed in (Domingo-Enrich et al., 2024), this constraint can lead to the initial  
 165 value function bias problem since the KL regularization is with respect to the learned reverse process.  
 166 It becomes necessary to learn an appropriate tilt even at time  $T$ . In this context, Lemma 3.2 offers a  
 167 simple yet effective alternative to implicitly achieve marginal KL regularization.

168 Based on GRS, we propose **GRAFT: Generalized Rejection sAmpling Fine Tuning** (Algorithm 7)  
 169 - given a reference model  $\bar{p}$ , we generate samples and perform the GRS strategy proposed in 3.1. A  
 170 dataset is generated from the accepted samples and standard training is done on the generated dataset.  
 171

## 172 4 PARTIAL-GRAFT FOR DIFFUSION MODELS

173 Having established that GRAFT samples from the PPO solution, we now examine methods to further  
 174 improve the framework. Continuous diffusion models typically start with Gaussian noise  $X_T$  at time  
 175  $T$  and denoise it to the output  $X_0$  via a discretized continuous time SDE. With  $N$  denoising steps,  
 176 the model constructs a denoising trajectory  $X_{t_N} \rightarrow \dots X_{t_i} \rightarrow \dots \rightarrow X_{t_0}$  ( $t_N = T$  and  $t_0 = 0$ ),  
 177 denoted by  $X_{T:0}$ . We now consider the effect of *shaping the distribution of an intermediate state*  
 178  $X_t$ . To provide intuition, we consider the case of text-to-image generation. The higher noise levels  
 179 decide the “outline” of the image to be denoised and is sensitive to the conditioning. Additionally,  
 180 fine-tuning denoisers at lower noise levels require relatively more samples because learning is more  
 181 difficult. Thus, we introduce P-GRAFT to fine-tune only the decision-making higher noise levels. For  
 182 the rest of the denoising process, it suffices to revert to the original denoiser. Theoretically, we show  
 183 that although this results in higher reward variance (due to reward being used to select a partially  
 184 noisy state), learning becomes easier due to lower bias at higher noise levels. Empirically, this  
 185 results in consistent performance improvements over GRAFT (and other baselines) for text-to-image  
 186 benchmarks with Stable Diffusion v2.  
 187

188 For the rest of the discussion, we reserve  $n$  and  $N$  to refer to discrete timesteps, and  $t$  and  $T$  for  
 189 continuous time. For any  $t \in [0, T]$  denote the marginal density of  $X_t$  by  $\bar{p}_t(x)$ .

190 We first extend GRS to Partial Generalized Rejection Sampling (P-GRS). Let  $X_t^{(1)}, \dots, X_t^{(M)}$  be  
 191 partially denoised (denoised till time  $t$ ) samples. Let their corresponding completely denoised samples  
 192 be  $X_0^{(1)}, \dots, X_0^{(M)}$ . Rewards are computed using the completely denoised samples (i.e.  $R_i = r(X_0^{(i)})$   
 193 for the  $i^{\text{th}}$  sample). We denote the empirical distribution of  $\{X_0^{(1)}, \dots, X_0^{(M)}\}$  by  $\hat{P}_{X_0}(\cdot)$  and the  
 194 empirical CDF of  $\{R_1, \dots, R_M\}$  by  $\hat{F}_R(\cdot)$ .  
 195

196 **Definition 4.1. Partial Generalized Rejection Sampling (P-GRS):** Consider an acceptance function  
 197  $A : \mathbb{R} \times [0, 1] \times \mathcal{X} \times [0, 1] \rightarrow [0, 1]$  such that  $A$  is co-ordinate wise increasing in the first two  
 198 co-ordinates. The acceptance probability of sample  $i$  is  $p_i := A(R_i, \hat{F}_R(R_i), X_0^{(i)}, \hat{P}_{X_0})$ . Draw  
 199  $C_i \sim \text{Ber}(p_i) \forall i \in [M]$ , not necessarily independent of each other. Then,  $X_t^{(i)} \in \mathcal{A}$  iff  $C_i = 1$ .

200 **Lemma 4.2.** The law of the accepted samples under P-GRS (Def. 4.1) given by  $p_t(X_t^{(1)} = x | X_t^{(1)} \in \mathcal{A})$  is the solution to the following Proximal Policy Optimization problem:

$$201 \arg \max_{\hat{p}} [\mathbb{E}_{X \sim \hat{p}} \hat{r}(X) - \alpha \text{KL}(\hat{p} || \bar{p}_t)] ; \quad \frac{\hat{r}(x)}{\alpha} := \log (\mathbb{E}[A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})), X_0^{(1)}, \hat{P}_{X_0}) | X_t^{(1)} = x])$$

205 The key difference is that the reshaped reward now depends on the *expected value* of the acceptance  
 206 function given a partially denoised state  $X_t$ . This tilts  $\bar{p}_t$  instead of  $\bar{p}_0$ . It is straightforward to modify  
 207 the reshaped rewards corresponding to GRS to that of P-GRS. We illustrate this by instantiating  
 208 Lemma 4.2 for preference rewards, as done with GRS (Lemma 3.2) above.

209 **Preference rewards:** With P-GRS,  $p_t(X_t^{(1)} = x | X_t^{(1)} \in \mathcal{A}) \propto \bar{p}_t(x) \exp\left(\frac{\hat{r}(x)}{\alpha}\right)$  with  $\frac{\hat{r}(x)}{\alpha} =$   
 210  $\log \mathbb{E}[F(r(X_0)) | X_t = x]$ .

212 Based on Lemma 4.2, we introduce **P-GRAFT: Partial GRAFT** (Algorithms 1 and 2). Here, fine-  
 213 tuning is done on a (sampled) dataset of *partially denoised vectors* instead of fully denoised vectors.  
 214 The fine-tuned model is *only trained from times  $T$  to  $t$* , and is used for denoising from noise only till  
 215 time  $t$ . We switch to the reference model for further denoising. The resulting stitched distribution is  
 given in Appendix D.3.1. We will now give a mathematical justification for P-GRAFT.

216 4.1 A BIAS-VARIANCE TRADEOFF JUSTIFICATION FOR P-GRAFT  
217

218 We analyze P-GRAFT from a bias-variance tradeoff viewpoint. Let us associate reward  $r(X_0)$  with  
219  $X_t$ . As argued in Lemma 4.3, variance of  $r(X_0)$  conditioned on  $X_t$  increases with  $t$ . Consequently,  
220 P-GRAFT obtains noisy rewards, seemingly making it less effective than GRAFT. However, we  
221 subsequently show that the learning problem itself *becomes easier* when  $t$  is large since the score  
222 function becomes simpler (i.e, the bias reduces). Therefore, we can balance the trade-off between the  
223 two by choosing an “appropriate” intermediate time  $t$  for the distributional tilt.

224 **Lemma 4.3.** *The expected conditional variance  $\mathbb{E}[\text{Var}(r(X_0)|X_t)]$  is an increasing function of  $t$ .*

225 **Example:** Consider molecule generation, where molecules are generated by a pre-trained diffusion  
226 model. The generated molecule can be stable ( $r(X_0) = 1$ ) or unstable ( $r(X_0) = 0$ ). Intuitively,  $X_t$ ,  
227 for  $t < T$ , carries more information about  $r(X_0)$  than  $X_T$ . We reinforce this claim empirically by  
228 giving the following illustrative statistical test. Consider the two hypotheses:  
229

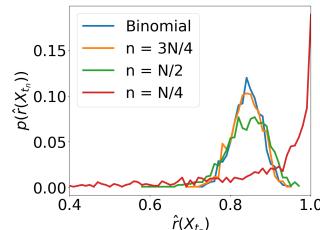
230  $H_0 : r(X_0)$  is independent of  $X_t$  ;       $H_1 : r(X_0)$  and  $X_t$  are dependent.  
231

232 **Algorithm 1** P-GRAFT: Training  
233

234 **Input:** Trainable model  $p_\theta$ , Reference model  $\bar{p}$ , Reward  
235 function  $r$ , Acceptance function  $A$ , Number of rounds  
236  $N_S$ , Intermediate timestep  $N_I$   
237  
1: Initialize empty set  $\mathcal{D}$   
2: **for**  $j = 1$  to  $N_S$  **do**  
3:    Generate  $M$  trajectories:  $X_{T:0}^{(i)} \sim \bar{p}_{T:0}$  ;     $i \in [M]$   
4:    Obtain rewards:  $r(X_0^{(i)})$  ;     $i \in [M]$   
5:    Perform P-GRS using acceptance function  $A$  on  
6:     $X_{t_{N_I}}^{(i)}$  ;     $i \in [M]$  to get accepted samples  $\mathcal{A}$   
7:    Perform  $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{A}$   
8: **end for**  
9: Train  $p_\theta$  on  $\mathcal{D}$  for  $t \in \{t_{N_I}, \dots, t_N\}$   
9: **return**  $p_\theta$

248 Given  $X_t$ , we obtain 100 roll outs  
249  $X_0^{(i)}|X_t$  for  $1 \leq i \leq 100$  and  
250 its empirical average  $\hat{r}(X_t) =$   
251  $\sum_{i=1}^{100} r(X_0^{(i)})/100$ . If  $r(X_0)$  is  
252 independent of  $X_t$  (under  $H_0$ ),  
253 the law of  $\hat{r}(X_t)$  is the binomial  
254 distribution  $\text{Bin}(100, \theta)$  with  $\theta =$   
255  $\mathbb{P}(r(X_0) = 1)$  being the marginal  
256 probability of observing a stable  
257 molecule. We perform 1000 rep-  
258 etitions for the experiment above  
259 for various values of  $t$  and plot the  
260 empirical distributions in Figure 1. For  $t = t_{3N/4}$  (when  $X_t$  close to  $\mathcal{N}(0, \mathbf{I})$ ), the distribution is  
261 close to the Binomial distribution and for  $t = t_{N/4}$  (when  $X_t$  is close to the target) it is far. That is,  
262  $X_{t_{N/4}}$  already carries a lot of information about  $r(X_0)$ . This is further supported by the expected  
263 conditional variances reported in Table 1.

264 **Bias reduces with increasing  $t$ :** We follow the Stochastic Differential Equation (SDE) framework  
265 from Song et al. (2020b) for our analysis. Let the target distribution  $q_0$  be the law of accepted samples  
266 under P-GRS. Diffusion models consider the forward process to be the Ornstein-Uhlenbeck Process  
267 given by  $dX_t^f = -X_t^f dt + \sqrt{2} dB_t$  where  $X_0^f \sim q_0$  is drawn from the target distribution over  $\mathbb{R}^d$   
268 and  $B_t$  is the standard Brownian motion in  $\mathbb{R}^d$ . It is well known that  $X_t^f \stackrel{d}{=} e^{-t} X_0^f + \sqrt{1 - e^{-2t}} Z$ ,  
269 where  $Z \sim \mathcal{N}(0, \mathbf{I})$  independent of  $X_0^f$ .

Figure 1: Law of  $\hat{r}(X_t)$ 232 **Algorithm 2** P-GRAFT: Inference  
233

234 **Input:** Fine tuned model  $\hat{p}$ , Reference  
235 model  $\bar{p}$ , Intermediate timestep  $N_I$ ,  
236 Per-step denoiser DEN  
237  
1: Sample  $X_T \sim \mathcal{N}(0, \mathbf{I})$   
2: **for**  $n = N - 1$  to  $N_I$  **do**  
3:     $X_{t_n} \leftarrow \text{DEN}(\hat{p}, X_{t_{n+1}}, t_{n+1})$   
4: **end for**  
5: **for**  $n = N_I - 1$  to 0 **do**  
6:     $X_{t_n} \leftarrow \text{DEN}(\bar{p}, X_{t_{n+1}}, t_{n+1})$   
7: **end for**  
8: **return**  $X_{t_0}$

Table 1: Conditional variance.

$n$	$\mathbb{E}[\text{Var}(r(X_0) X_{t_n})]$
$N$	0.1341
$3N/4$	0.1327
$N/2$	0.1312
$N/4$	0.0848

Let  $q_t$  be the density of the law of  $X_t^f$ . Diffusion models learn the score function  $[0, T] \times \mathbb{R}^d \ni (t, X) \rightarrow \nabla \log q_t(X)$  via score matching (see Appendix A for literature review on score matching). P-GRAFT, in contrast, attempts to learn  $\nabla \log q_s$  between  $s \in [t, T]$ . At time  $T$ ,  $\nabla \log q_T(X) \approx -X$ , the score of the standard Gaussian distribution, which is easy to learn. When  $t = 0$ , the score  $\nabla \log q_0(X)$  corresponds to the data distribution which can be very complicated. Diffusion models use Denoising Score Matching, based on Tweedie's formula introduced by (Vincent, 2011). We show via Bakry-Emery theory (Bakry et al., 2013) that the score function  $\nabla \log q_t(X)$  converges to  $q_\infty(X)$  exponentially in  $t$ , potentially making the learning easier. Consider  $s_\theta(X, t) : \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}^d$  to be a neural network with parameters  $\theta$ , then score matching objective is given by:

$$\mathcal{L}(\theta) = \mathbb{E} \int_0^T dt \left\| \frac{X_t^f - e^{-t} X_0^f}{1 - e^{-2t}} + s_\theta(X_t^f, t) \right\|^2.$$

In practice,  $\mathcal{L}(\theta)$  is approximated with samples. By Tweedie's formula, we have:  $\mathbb{E} \left[ \frac{X_t^f - e^{-t} X_0^f}{1 - e^{-2t}} \mid X_t^f \right] = -\nabla \log q_t(X_t^f)$ . Thus, for some constant  $C$ , independent of  $\theta$ :

$$\mathcal{L}(\theta) + C = \mathbb{E} \int_0^T dt \left\| \nabla \log q_t(X_t^f) - s_\theta(X_t^f, t) \right\|^2 = \int_0^T dt \int_{\mathbb{R}^d} dX q_t(X) \left\| \nabla \log q_t(X) - s_\theta(X, t) \right\|^2.$$

As shown by (Benton et al., 2023),  $\mathcal{L}(\theta)$  directly controls the quality of generations. Note that  $q_\infty$  is the density of  $\mathcal{N}(0, \mathbf{I})$  and  $\nabla \log q_\infty(X) = -X$ . The theorem below is proved in Appendix D.5.

**Theorem 4.4.** Define  $H_t^s$  for  $s \leq t$ :  $H_t^s = \int_s^t dt \int_{\mathbb{R}^d} dX q_s(X) \left\| \nabla \log q_s(X) - \nabla \log q_\infty(X) \right\|^2$ . Then,

$$H_t^T \leq \frac{e^{-2t}}{1 - e^{-2t}} H_0^t$$

Therefore, the score functions between time  $(t, T)$  are exponentially closer to the simple Gaussian score function compared to the score functions between times  $(0, t)$  in the precise sense given in Theorem 4.4. This means that the score functions at later times should be **easier** to learn.

## 5 INVERSE NOISE CORRECTION FOR FLOW MODELS

In the analysis so far, we have established bias-variance tradeoffs for diffusion models - models which use SDEs to sample from a target distribution. We now extend this analysis to flow models, which use ODEs to sample. Flow models follow a deterministic ODE starting from an initial (random) noise. The bias-variance results from the previous section indicate that, conditioned on the initial noise vector, the variance of reward should be zero, making the learning process potentially easier. Another property of flow models is that because of the deterministic mapping, they admit reversal - this property has been utilized extensively in the literature to map images to 'noise' for image editing Rout et al. (2024); Garibi et al. (2024) and as part of the 2-rectification Liu et al. (2022) to achieve straighter flows. We will combine these two ideas to develop a framework for improving flow models *even without explicit rewards*. We now develop this idea from first principles.

We restrict our attention to flow models with optimal transport based interpolation (Lipman et al., 2022; Liu et al., 2022), which learn a velocity field  $v(x, t) : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$  such that the following ODE's solution at time  $t = 1$  has the target distribution  $p^{\text{data}}$ :

$$\frac{dX_t}{dt} = v(X_t, t), \quad X_0 \sim \mathcal{N}(0, \mathbf{I}). \quad (2)$$

Note that *in the literature for flow models (unlike diffusion models)*,  $t = 0$  corresponds to noise and  $t = 1$  corresponds to the target, a convention we follow in this section.

324 **The errors in learned model:** Suppose we have a pre-trained vector-field, corresponding to parameter  
 325  $\theta$  and solve the ODE equation 2 with  $v(x, t) = v_\theta(x, t)$ . Then,  $\text{Law}(X_1) \neq p^{\text{data}}$  due to:

326 **a)** Discretization error of the ODE and **b)** Statistical error due to imperfect learning.

328 Despite these two errors, the trained ODE is still invertible. We will leverage reversibility to arrive at  
 329 our algorithm. To this end, consider the time reversal of equation 2:

$$330 \quad \frac{dx_t^{\text{rev}}}{dt} = -v_\theta(x_t^{\text{rev}}, 1-t), \quad x^{\text{rev}}(0) \sim p^{\text{data}}. \quad (3)$$

---

333 **Algorithm 3** Inverse Noise Correction: Training

---

334 **Input:** Dataset  $\mathcal{D} := \{X^{(1)}, X^{(2)}, \dots, X^{(M)}\} \sim$   
 335  $p^{\text{data}}$ , step-size  $\eta$ , backward Euler steps  $N_b$

337 1:  $v_\theta = \text{TRAIN\_FLOW}(\mathcal{N}(0, \mathbf{I}), \mathcal{D})$ .

338 2: **for**  $i = 1$  to  $M$  **do**

339 3:  $X_1^{\text{rev},(i)} \leftarrow \text{BWD\_Euler}(v_\theta, \eta, X^{(i)}, N_b)$

340 4: **end for**

341 5: Dataset  $\mathcal{D}^{\text{rev}} \leftarrow \{X_1^{\text{rev},(1)}, \dots, X_1^{\text{rev},(M)}\} \sim p_1^{\text{rev}}$

342 6:  $v_{\theta^{\text{rev}}} = \text{TRAIN\_FLOW}(\mathcal{N}(0, \mathbf{I}), \mathcal{D}^{\text{rev}})$

343 7: **return**  $v_\theta, v_{\theta'}$

---



---

344 **Algorithm 4** Inference

---

345 **Input:** Flow models  $v_\theta, v_{\theta^{\text{rev}}}$ , step-size  $\eta$ ,  
 346 Initial point  $X_0 \sim \mathcal{N}(0, \mathbf{I})$

347 1:  $X_1^{\text{rev}} \leftarrow \text{FWD\_Euler}(v_{\theta^{\text{rev}}}, \eta, X_0)$

348 2:  $X_1 \leftarrow \text{FWD\_Euler}(v_\theta, \eta, X_1^{\text{rev}})$

349 3: **return**  $X_1$

---

350 **The Inverse Noise:** Consider the forward Euler discretization of equation 2 with step-size  $\eta$ :

$$351 \quad \hat{x}_{(i+1)\eta} \leftarrow \hat{x}_{i\eta} + \eta v_\theta(\hat{x}_{i\eta}, i\eta). \quad (4)$$

352 Let  $T_{\theta, \eta}$  be the function which maps  $\hat{x}_0$  to  $\hat{x}_1$  i.e.,  $\hat{x}_1 = T_{\theta, \eta}(\hat{x}_0)$ . The forward Euler approximation  
 353  $T_{\theta, \eta}^{-1}(\hat{x}_1) \approx \hat{y}_1$  where  $\hat{y}_{i\eta} \leftarrow \hat{y}_{(i-1)\eta} - \eta v_\theta(\hat{y}_{(i-1)\eta}, 1 - (i-1)\eta)$  with  $\hat{y}_0 = \hat{x}_1$  is not good enough  
 354 as noted in the image inversion/ editing literature Rout et al. (2024); Wang et al. (2024); Garibi et al.  
 355 (2024). This is mitigated via numerical and control theoretic techniques. We utilize the ‘backward  
 356 Euler discretization’ (equation 3, as used in Garibi et al. (2024)) to exactly invert equation 4.

$$357 \quad \hat{x}_{\eta i}^{\text{rev}} \leftarrow \hat{x}_{\eta(i-1)}^{\text{rev}} - \eta v_\theta(\hat{x}_{\eta(i-1)}^{\text{rev}}, 1 - \eta(i-1)) \quad (5)$$

358 This is an implicit equation since  $\hat{x}_{\eta i}^{\text{rev}}$  being calculated in the LHS also appears in the RHS. It is not  
 359 apriori clear that this can be solved. Lemma 5.1 addresses this issue:

360 **Lemma 5.1.** Suppose  $v_\theta$  is  $L$  Lipschitz in  $x$  under  $\ell_2$ -norm and  $\eta L < 1$ . Then,

361 (1)  $\hat{x}_{\eta i}^{\text{rev}}$  in equation 5 has a unique solution which can be obtained by a fixed point method.

362 (2)  $T_{\theta, \eta}$  is invertible and  $T_{\theta, \eta}^{-1}(x_0^{\text{rev}}) = x_1^{\text{rev}}$ .

363 That is, the mapping from noise to data given by the learned, discretized model is invertible. We show  
 364 some important consequences of this in Lemma 5.2. Define the following probability distributions.  
 365 Let  $p^{\text{data}} = \text{Law}(\text{Data})$  (i.e, target data distribution).

$$366 \quad \boxed{p_0 = \text{Law}(\hat{x}_0) = \mathcal{N}(0, \mathbf{I}) \quad p_1 = \text{Law}(\hat{x}_1) \quad p_0^{\text{rev}} = \text{Law}(\hat{x}_0^{\text{rev}}) = p^{\text{data}} \quad p_1^{\text{rev}} = \text{Law}(\hat{x}_1^{\text{rev}})}$$

367 We call  $p_1^{\text{rev}}$  the inverse noise distribution. With perfect training and 0 discretization error,  $p_1^{\text{rev}} =$   
 368  $\mathcal{N}(0, \mathbf{I})$ . However, due to these errors  $p_1^{\text{rev}} \neq \mathcal{N}(0, \mathbf{I})$ .

369 **Lemma 5.2.** Under the assumption of Lemma 5.1,  $p_1^{\text{rev}}, p_1, p^{\text{data}}$  and  $p_0 = \mathcal{N}(0, \mathbf{I})$  satisfy:

370 1.  $(T_{\theta, \eta})_# p_1^{\text{rev}} = p^{\text{data}}$  ; 2.  $\text{TV}(p_1^{\text{rev}}, p_0) = \text{TV}(p_1, p^{\text{data}})$  ; 3.  $\text{KL}(p_0 || p_1^{\text{rev}}) = \text{KL}(p_1 || p^{\text{data}})$ .

371 That is, the distance between the inverse noise and the true noise is the same as the distance between  
 372 the generated distribution and the true target distribution. Item 1 shows that if we can sample from the  
 373 inverse noise distribution  $p_1^{\text{rev}}$ , then we can use the pre-trained model  $v_\theta(\cdot, \cdot)$  with discretization and  
 374 obtain samples from the true target  $p^{\text{data}}$ . In Kim et al. (2024), the authors note that even 2-rectification  
 375 suffers when the inverse noise  $p_1^{\text{rev}}$  is far from  $\mathcal{N}(0, \mathbf{I})$ . While 2-rectification aims to improve improve  
 376 the computational complexity while *maintaining quality* by aiming to obtain straight flows, we  
 377 introduce inverse noise correction to *improve quality* of generations in a sample efficient way.

378 **Inverse Noise Correction:** Inverse Noise Correction is given in Algorithms 3 and 4, and illustrated  
 379 in Figure 2. Given samples from the target distribution,  $\mathcal{D}$ , TRAIN\_FLOW( $\mathcal{N}(0, \mathbf{I})$ ,  $\mathcal{D}$ ) trains a  
 380 rectified flow model between  $\mathcal{N}(0, \mathbf{I})$  to the target distribution Liu et al. (2022). Now, suppose we  
 381 are given a dataset  $\{X^{(1)}, \dots, X^{(M)}\} \sim p^{\text{data}}$  and a trained flow model  $v_\theta$  which generates  $\hat{x}_1 \sim p_1$   
 382 using equation 4 starting with  $\hat{x}_0 \sim p_0$ . We obtain samples  $X_1^{\text{rev},(i)} \sim p_1^{\text{rev}}$  by backward Euler  
 383 iteration in equation 5. Thereafter, we train another flow model  $v_{\theta^{\text{rev}}}$  which learns to sample from  $p_1^{\text{rev}}$   
 384 starting from  $\mathcal{N}(0, \mathbf{I})$ .  
 385

386 During inference, we sample a point from  $X_0 \sim \mathcal{N}(0, \mathbf{I})$  and obtain a sample  $X_1^{\text{rev}} \sim p_1^{\text{rev}}$  using  
 387  $v_{\theta^{\text{rev}}}$ . Once we have the corrected noise sample, we generate images using the original flow model  
 388  $v_\theta$  which now starts from  $X_1^{\text{rev}}$  instead of  $X_0$ . FWD\_Euler( $v_\theta, \eta, \hat{x}_0$ ) obtains  $\hat{x}_1$  via Euler iteration  
 389 (equation 4). Similarly, BWD\_Euler( $v_\theta, \eta, \hat{x}_0^{\text{rev}}, N_b$ ) obtains  $x_1^{\text{rev}}$  by approximately solving backward  
 390 Euler iteration (equation 5). They are formally described as Algorithms 5 and 6 in Appendix B.  
 391 **Theoretical Justification** along the lines of Section 4.1 is given in Appendix D.8.  
 392

## 6 EXPERIMENTS

394 We use the notation P-GRAFT( $N_I$ ) to denote P-GRAFT with intermediate timestep  $N_I$  as described  
 395 in Algorithms 1 and 2. For instance, P-GRAFT( $0.75N$ ) would denote instantiating P-GRAFT with  
 396  $N_I = 0.75N$ , where  $N$  is the total number of denoising steps. Recall that  $t_N$  corresponds to pure  
 397 noise and  $t_0$  corresponds to a completely denoised sample.  
 398

### 6.1 TEXT-TO-IMAGE GENERATION

401 **Setup:** The objective is to fine-tune a pre-trained model so that generated images better align with  
 402 prompts. We consider Stable Diffusion v2 (SDv2) (Rombach et al., 2022) as the pre-trained model.  
 403 The reward model used is VQAScore (Lin et al., 2024) - a prompt-image alignment score between 0  
 404 to 1, with higher scores denoting better prompt-alignment. We fine-tune (separately) on GenAI-Bench  
 405 (Li et al., 2024a) as well as the train split of T2ICompBench++ (Huang et al., 2025). Evaluations are  
 406 done on GenAI-Bench, validation split of T2ICompBench++ and GenEval (Ghosh et al., 2023). We  
 407 use LoRA (Hu et al., 2021) for compute-efficient fine-tuning. Top – K sampling (Section 3) is used  
 408 for both GRAFT and P-GRAFT. Since LoRA fine-tuning is used, the model switching in 2 can be  
 409 done by simply turning off the LoRA adapter. More implementation details are given in Appendix E.  
 410

411 **Results:** are reported in Table 2 - for fine-tuning on GenAI-Bench, we use Top – 10 of 100 samples  
 412 and on T2ICompBench++, we use Top – 1 of 4 samples. First, note that **both GRAFT and P-**  
 413 **GRAFT outperform** base SDv2, SDXL-Base and DDPO. The **best performance is obtained for**  
 414 **P-GRAFT with  $N_I = 0.25N$**  across all evaluations - this clearly shows the *bias-variance tradeoff*  
 in action. Further, both GRAFT and P-GRAFT also **generalize to unseen prompts**.  
 415

416 In particular, DDPO did not improve over the baseline even when trained with more samples and  
 417 FLOPs as compared to GRAFT/P-GRAFT. Experiments with different sets of hyperparameters as  
 418 well as adding other features such as KL regularization and a per-prompt advantage estimator on  
 419 top of DDPO also did not show any significant improvements over SDv2 (see Appendix E.3). We  
 420 also conduct ablations to further verify the effectiveness of the proposed methods - these include  
 421 experiments on different values of  $(M, K)$  in Top – K of  $M$  sampling, different LoRA ranks for  
 422 fine-tuning as well as a reverse P-GRAFT strategy (where the fine-tuned model is used in the later  
 423 denoising steps instead of initial steps). We find that P-GRAFT remains effective across different  
 424  $(M, K)$  and that performance is insensitive to the LoRA rank. Further, P-GRAFT significantly  
 425 outperforms reverse P-GRAFT. More details on ablations can be found in Appendix E.1.  
 426

### 6.2 LAYOUT AND MOLECULE GENERATION

427 **Setup:** All experiments are done on pre-trained models trained using IGD (Anil et al., 2025), a  
 428 discrete-continuous diffusion framework capable of handling both layout generation and molecule  
 429 generation. For layouts, we experiment with improving the alignment of elements in the generated  
 430 layout as measured by the alignment metric - note that the reward is taken as 1 - alignment since lower  
 431 values for the metric indicate better alignment. For molecules, the objective is to generate a larger  
 fraction of stable molecules - molecules which are deemed stable are assigned a reward of 1 whereas

432 Table 2: **Text-to-Image Generation fine-tuning on SDv2:** VQAScore (normalized to 100) reported  
 433 on GenAI-Bench, T2ICompBench++ - Val (denoted as T2I - Val) and GenEval.  
 434

435 Model	436 Fine-Tuned on GenAI-Bench			437 Fine-Tuned on T2ICompBench++ - Train		
	438 GenAI	439 T2I - Val	440 GenEval	441 GenAI	442 T2I - Val	443 GenEval
SD v2	66.87 $\pm$ 0.14	69.20 $\pm$ 0.17	73.49 $\pm$ 0.41	66.87 $\pm$ 0.14	69.20 $\pm$ 0.17	73.49 $\pm$ 0.41
SDXL-Base	69.69 $\pm$ 0.17	72.98 $\pm$ 0.16	73.90 $\pm$ 0.40	69.69 $\pm$ 0.17	72.98 $\pm$ 0.16	73.90 $\pm$ 0.40
DDPO	65.70 $\pm$ 0.17	68.03 $\pm$ 0.16	72.13 $\pm$ 0.37	64.65 $\pm$ 0.17	69.05 $\pm$ 0.15	69.60 $\pm$ 0.37
GRAFT	70.51 $\pm$ 0.15	75.69 $\pm$ 0.13	79.85 $\pm$ 0.31	70.97 $\pm$ 0.14	75.88 $\pm$ 0.13	79.57 $\pm$ 0.30
P-GRAFT(0.75N)	69.46 $\pm$ 0.15	74.51 $\pm$ 0.14	79.44 $\pm$ 0.33	69.51 $\pm$ 0.15	74.30 $\pm$ 0.13	78.50 $\pm$ 0.33
P-GRAFT(0.5N)	71.00 $\pm$ 0.14	75.45 $\pm$ 0.14	80.60 $\pm$ 0.31	70.73 $\pm$ 0.14	75.37 $\pm$ 0.12	79.25 $\pm$ 0.30
P-GRAFT(0.25N)	71.94 $\pm$ 0.14	<b>76.12</b> $\pm$ 0.13	<b>80.96</b> $\pm$ 0.29	<b>71.42</b> $\pm$ 0.14	<b>76.15</b> $\pm$ 0.13	<b>80.29</b> $\pm$ 0.30

444 Table 3: **Layout Generation:** Fine-tuning  
 445 results for unconditional and category-  
 446 conditional generation on PubLayNet.  
 447

448 Model	449 Unconditional		450 Class-conditional	
	451 Alignment	452 FID	453 Alignment	454 FID
Baseline	0.094	8.32	0.088	4.08
GRAFT	0.064	10.68	0.068	5.04
P-GRAFT(0.5N)	0.071	9.24	0.072	4.55
P-GRAFT(0.25N)	<b>0.053</b>	9.91	<b>0.064</b>	4.67

455 unstable molecules are assigned a reward of 0. For molecule generation, we use the *de-duplication*  
 456 instantiation of GRAFT/P-GRAFT (Section 3) to ensure diversity of generated molecules - we  
 457 use RDKit to determine whether two molecules are identical or not. We use PubLayNet (Zhong  
 458 et al., 2019) for layout generation, and QM9 (Ramakrishnan et al., 2014) for molecule generation.  
 459 To the best of our knowledge, this is the first work which addresses fine-tuning in the context of  
 460 discrete-continuous diffusion models. Ablations and experimental details are given in Appendix F.  
 461

462 **Results:** for layout generation are given in Table 3. Both P-GRAFT and GRAFT uniformly improve  
 463 performance across both unconditional and class-conditional generation, with P-GRAFT:0.25N  
 464 giving the best performance. We also report FID scores computed between the generated samples and  
 465 the test set of PubLayNet - this is a measure of how close the generated samples are to the pre-training  
 466 distribution. As expected, the baseline has the lowest FID. Note that the FID score for P-GRAFT is  
 467 smaller than GRAFT, indicating that *P-GRAFT aligns more closely* to the pre-training distribution.  
 468 For molecule generation, results are given in Table 4. Again, the best performance is with P-GRAFT  
 469 at 0.25N. Note that improvement with GRAFT is marginal, despite being trained on 9 $\times$  the number  
 470 of samples used for P-GRAFT - this points to the learning difficulty in later denoising steps.  
 471

### 472 6.3 IMAGE GENERATION WITH INVERSE NOISE CORRECTION

473 **Setup:** We consider unconditional image generation on CelebA-HQ (Karras et al., 2017) and LSUN-  
 474 Church (Yu et al., 2015) at  $256 \times 256$  resolution. We first train pixel-space flow models from scratch.  
 475 A training corpus of inverted noise is then generated by running the trained flow models in reverse,  
 476 employing the backward Euler method, on all samples in the dataset. A second flow model, which  
 477 we refer to as the Noise Corrector model, is then trained to generate this inverse noise. Once the  
 478 Noise Corrector is trained, this model is first used to transform standard Gaussian noise to the inverse  
 479 noise. The pre-trained model then generates samples *starting from the inverse noise*. FID with 50000  
 480 generated samples with respect to the dataset is used to measure the performance. We emphasize  
 481 that the our goal is not to compete with state-of-the-art (SOTA) models rather to demonstrate that  
 482 our procedure can be used to improve the performance of a given flow model by simply learning the  
 483 distributional shift of noise at  $t = 0$ . SOTA models are larger (Rombach et al. (2022) has  $\approx 300$ M  
 484 parameters) and are more sophisticated - we do not seek to match their performance.  
 485

486 **Results:** Table 5, shows that the Noise Corrector **significantly improves FID** scores across both  
 487 datasets. Apart from quality gains, Noise Corrector also allows for **faster generation** - running the  
 488 Noise Corrector for 100 steps and then running the pre-trained model for 100 steps can *outperforms*

486 Table 5: **Image Generation**: Results for inverse noise correction on CelebA-HQ and LSUN-Church.  
 487 The noise corrector samples the inverse noise starting from  $\mathcal{N}(0, \mathbf{I})$  for ‘Sampling Steps’, and the  
 488 pre-trained model samples the image starting from the inverse noise.

490	Sampling Steps		FID		FLOPs/image ( $\times 10^{12}$ )	Inference Time (in seconds)
	491 Noise Corrector (16M parameters)	492 Pre-Trained Model (65M parameters)	493 CelebA-HQ (256 $\times$ 256)	494 LSUN-Church (256 $\times$ 256)		
495	-	1000	11.93	8.40	6.869	482.85
496	-	200	13.39	8.63	1.374	96.64
497	100	100	8.94	7.90	0.903	67.85
498	200	200	<b>8.02</b>	<b>7.26</b>	1.806	135.42

499 the pre-trained model with 1000 steps. The Noise Corrector only has  $0.25 \times$  the number of parameters,  
 500 leading to further latency gains as evidenced by FLOPs counts.

## 501 7 CONCLUSION

502 We establish GRAFT, a framework for provably sampling from PPO solution with marginal KL for  
 503 diffusion models through rejection sampling. We then introduce P-GRAFT, a principled framework  
 504 for intermediate distribution shaping of diffusion models and provide a mathematical justification for  
 505 this framework. Both GRAFT and P-GRAFT perform well empirically, outperforming policy gradient  
 506 methods on the text-to-image generation task. Further, both frameworks also extend seamlessly to  
 507 discrete-continuous diffusion models. Finally, we introduce Inverse Noise Correction, a strategy to  
 508 improve flow models even without explicit rewards and demonstrate significant quality gains even  
 509 with lower FLOPs/image.

511  
 512  
 513  
 514  
 515  
 516  
 517  
 518  
 519  
 520  
 521  
 522  
 523  
 524  
 525  
 526  
 527  
 528  
 529  
 530  
 531  
 532  
 533  
 534  
 535  
 536  
 537  
 538  
 539

540 **8 ETHICS STATEMENT**  
541542 The proposed method fine-tunes a pre-trained diffusion model based on rewards. Potentially, fine-  
543 tuning towards undesirable goals is possible by using specialized rewards. Practitioners are suggested  
544 to exercise caution in this regard.  
545546 **9 REPRODUCIBILITY STATEMENT**  
547548 Algorithms 1, 2, 3, 4, 5, 6 and 7 provide algorithmic descriptions of the proposed methods. The  
549 experimental setup used for experiments, including hyperparameters, are described in Section 6 as  
550 well as Appendices E, F and G. Proofs for the theoretical claims are given in Appendix D.  
551

552

553

554

555

556

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594 REFERENCES  
595

596 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,  
597 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning  
598 from human feedback in llms. [arXiv preprint arXiv:2402.14740](https://arxiv.org/abs/2402.14740), 2024.

599 Afra Amini, Tim Vieira, Elliott Ash, and Ryan Cotterell. Variational best-of-n alignment. [arXiv  
600 preprint arXiv:2407.06057](https://arxiv.org/abs/2407.06057), 2024.

601 Gautham Govind Anil, Sachin Yadav, Dheeraj Nagaraj, Karthikeyan Shanmugam, and Prateek Jain.  
602 Interleaved gibbs diffusion for constrained generation. [arXiv preprint arXiv:2502.13450](https://arxiv.org/abs/2502.13450), 2025.

603 Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain,  
604 Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with  
605 reinforcement learning from human feedback. [arXiv preprint arXiv:2204.05862](https://arxiv.org/abs/2204.05862), 2022.

606 Dominique Bakry, Ivan Gentil, and Michel Ledoux. [Analysis and geometry of Markov diffusion  
607 operators](https://link.springer.com/chapter/10.1007/978-3-319-16018-1_1), volume 348. Springer Science & Business Media, 2013.

608 Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly  $d$ -linear conver-  
609 gence bounds for diffusion models via stochastic localization. [arXiv preprint arXiv:2308.03686](https://arxiv.org/abs/2308.03686),  
610 2023.

611 Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models  
612 with reinforcement learning. [arXiv preprint arXiv:2305.13301](https://arxiv.org/abs/2305.13301), 2023.

613 Adam Block, Youssef Mroueh, and Alexander Rakhlin. Generative modeling with denoising auto-  
614 encoders and langevin sampling. [arXiv preprint arXiv:2002.00107](https://arxiv.org/abs/2002.00107), 2020.

615 John Charles Butcher. [Numerical methods for ordinary differential equations](https://www.wiley.com/doi/10.1002/9781118629906). John Wiley & Sons,  
616 2016.

617 Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and  
618 Arnaud Doucet. A continuous time framework for discrete denoising models. [Advances in Neural  
619 Information Processing Systems](https://proceedings.neurips.cc/paper/2022/file/28266-28279.pdf), 35:28266–28279, 2022.

620 Minshuo Chen, Kaixuan Huang, Tuo Zhao, and Mengdi Wang. Score approximation, estimation and  
621 distribution recovery of diffusion models on low-dimensional data. In [International Conference  
622 on Machine Learning](https://proceedings.mlr.press/v137/chen23a.html), pp. 4672–4712. PMLR, 2023.

623 Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly fine-tuning diffusion models on  
624 differentiable rewards. [arXiv preprint arXiv:2309.17400](https://arxiv.org/abs/2309.17400), 2023.

625 Valentin De Bortoli, Michael Hutchinson, Peter Wirnsberger, and Arnaud Doucet. Target score  
626 matching. [arXiv preprint arXiv:2402.08667](https://arxiv.org/abs/2402.08667), 2024.

627 Fei Deng, Qifei Wang, Wei Wei, Tingbo Hou, and Matthias Grundmann. Prdp: Proximal reward  
628 difference prediction for large-scale reward finetuning of diffusion models. In [Proceedings of the  
629 IEEE/CVF Conference on Computer Vision and Pattern Recognition](https://openaccess.thecvf.com/content/ICCV2024/papers/Fei_Deng_Proximal_Reward_Difference_Prediction_for_Large-Scale_Reward_Finetuning_of_Diffusion_Models_ICCV_2024.pdf), pp. 7423–7433, 2024.

630 Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky TQ Chen. Adjoint matching:  
631 Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control.  
632 [ICLR 2025, arXiv:2409.08861](https://arxiv.org/abs/2409.08861), 2024.

633 Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao,  
634 Jipeng Zhang, Kashun Shum, and Tong Zhang. Raft: Reward ranked finetuning for generative  
635 foundation model alignment. [arXiv preprint arXiv:2304.06767](https://arxiv.org/abs/2304.06767), 2023.

636 Bradley Efron. Tweedie’s formula and selection bias. [Journal of the American Statistical Association](https://doi.org/10.1214/jasa.1166570177),  
637 106(496):1602–1614, 2011.

638 Luca Eyring, Shyamgopal Karthik, Alexey Dosovitskiy, Nataniel Ruiz, and Zeynep Akata. Noise hy-  
639 pernetworks: Amortizing test-time compute in diffusion models. [arXiv preprint arXiv:2508.09968](https://arxiv.org/abs/2508.09968),  
640 2025.

648 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,  
 649 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Dpok: Reinforcement learning for  
 650 fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*,  
 651 36:79858–79885, 2023.

652 Daniel Garibi, Or Patashnik, Andrey Voynov, Hadar Averbuch-Elor, and Daniel Cohen-Or. Renoise:  
 653 Real image inversion through iterative noising. In *European Conference on Computer Vision*, pp.  
 654 395–413. Springer, 2024.

655 Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and  
 656 Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:  
 657 133345–133385, 2024.

658 Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework  
 659 for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:  
 660 52132–52152, 2023.

661 Leonard Gross. Logarithmic sobolev inequalities. *American Journal of Mathematics*, 97(4):1061–  
 662 1083, 1975.

663 Lin Gui, Cristina Gârbacea, and Victor Veitch. Bonbon alignment for large language models and  
 664 the sweetness of best-of-n sampling. *Advances in Neural Information Processing Systems*, 37:  
 665 2851–2885, 2024.

666 Shivam Gupta, Aditya Parulekar, Eric Price, and Zhiyang Xun. Improved sample complexity  
 667 bounds for diffusion model training. *Advances in Neural Information Processing Systems*, 37:  
 668 40976–41012, 2024.

669 Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-  
 670 to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022.

671 Seongmin Hong, Kyeonghyun Lee, Suh Yoon Jeon, Hyewon Bae, and Se Young Chun. On exact  
 672 inversion of dpm-solvers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*  
 673 *Pattern Recognition*, pp. 7069–7078, 2024.

674 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 675 and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL <https://arxiv.org/abs/2106.09685>.

676 Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. Reinforce++: An efficient rlhf algorithm with  
 677 robustness to both prompt and reward models. *arXiv preprint arXiv:2501.03262*, 2025.

678 Kaiyi Huang, Chengqi Duan, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench++:  
 679 An enhanced and comprehensive benchmark for compositional text-to-image generation, 2025.  
 680 URL <https://arxiv.org/abs/2307.06350>.

681 Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching.  
 682 *Journal of Machine Learning Research*, 6(4), 2005.

683 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child,  
 684 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models,  
 685 2020. URL <https://arxiv.org/abs/2001.08361>.

686 Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for  
 687 improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.

688 Beomsu Kim, Yu-Guan Hsieh, Michal Klein, Marco Cuturi, Jong Chul Ye, Bahjat Kawar, and  
 689 James Thornton. Simple reflow: Improved techniques for fast flow models. *arXiv preprint*  
 690 *arXiv:2410.07815*, 2024.

691 Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models  
 692 for robust image manipulation. In *Proceedings of the IEEE/CVF conference on computer vision*  
 693 *and pattern recognition*, pp. 2426–2435, 2022.

702 Durk P Kingma and Yann Cun. Regularized estimation of image statistics by score matching.  
 703 *Advances in neural information processing systems*, 23, 2010.  
 704

705 Syamantak Kumar, Dheeraj Nagaraj, and Purnamrita Sarkar. Dimension-free score matching and  
 706 time bootstrapping for diffusion models. *arXiv preprint arXiv:2502.10354*, 2025.

707 Sangyun Lee, Zinan Lin, and Giulia Fanti. Improving the training of rectified flows. *Advances in*  
 708 *neural information processing systems*, 37:63082–63109, 2024.  
 709

710 Jose Lezama, Tim Salimans, Lu Jiang, Huiwen Chang, Jonathan Ho, and Irfan Essa. Discrete  
 711 predictor-corrector diffusion models for image synthesis. In *The Eleventh International Conference*  
 712 *on Learning Representations*, 2022.

713 Baiqi Li, Zhiqiu Lin, Deepak Pathak, Jiayao Li, Yixin Fei, Kewen Wu, Tiffany Ling, Xide Xia,  
 714 Pengchuan Zhang, Graham Neubig, and Deva Ramanan. Genai-bench: Evaluating and improving  
 715 compositional text-to-visual generation, 2024a. URL <https://arxiv.org/abs/2406.13743>.  
 716

717 Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Yusuke Kato, and Kazuki Kozuka. Aligning dif-  
 718 fusion models by optimizing human utility. *Advances in Neural Information Processing Systems*,  
 719 37:24897–24925, 2024b.  
 720

721 Ziniu Li, Tian Xu, Yushun Zhang, Zhihang Lin, Yang Yu, Ruoyu Sun, and Zhi-Quan Luo. Remax: A  
 722 simple, effective, and efficient reinforcement learning method for aligning large language models.  
 723 *arXiv preprint arXiv:2310.10505*, 2023.  
 724

725 Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and  
 726 Deva Ramanan. Evaluating text-to-visual generation with image-to-text generation, 2024. URL  
 727 <https://arxiv.org/abs/2404.01291>.  
 728

729 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching  
 730 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.  
 731

732 Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. Chain of hindsight aligns language models with  
 733 feedback. *arXiv preprint arXiv:2302.02676*, 2023a.  
 734

735 Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, and Jialu Liu.  
 736 Statistical rejection sampling improves preference optimization. *arXiv preprint arXiv:2309.06657*,  
 737 2023b.  
 738

739 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and  
 740 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.  
 741

742 Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al. Instaflow: One step is enough for  
 743 high-quality diffusion-based text-to-image generation. In *The Twelfth International Conference*  
 744 *on Learning Representations*, 2023c.  
 745

746 Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a reference-  
 747 free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235, 2024.  
 748

749 Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for  
 750 editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF conference*  
 751 *on computer vision and pattern recognition*, pp. 6038–6047, 2023.  
 752

753 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
 754 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow  
 755 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–  
 27744, 2022.  
 756

757 Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning*. Cambridge  
 758 University Press, 2025.  
 759

760 Mihir Prabhudesai, Anirudh Goyal, Deepak Pathak, and Katerina Fragkiadaki. Aligning text-to-image  
 761 diffusion models with reward backpropagation. 2023.  
 762

756 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea  
 757 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*  
 758 in neural information processing systems

759, 36:53728–53741, 2023.

760 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum  
 761 chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.

762 Allen Z Ren, Justin Lidard, Lars L Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar,  
 763 Benjamin Burchfiel, Hongkai Dai, and Max Simchowitz. Diffusion policy policy optimization.  
 764 *arXiv preprint arXiv:2409.00588*, 2024.

765 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 766 resolution image synthesis with latent diffusion models, 2022. URL <https://arxiv.org/abs/2112.10752>.

767

768 Litu Rout, Yujia Chen, Nataniel Ruiz, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng  
 769 Chu. Semantic image inversion and editing using rectified stochastic differential equations. *arXiv*  
 770 *preprint arXiv:2410.10792*, 2024.

771

772 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy  
 773 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

774

775 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Huawei Zhang,  
 776 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical  
 777 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

778

779 Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach  
 780 to density and score estimation. In *Uncertainty in artificial intelligence*, pp. 574–584. PMLR,  
 2020a.

781

782 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 783 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*  
 784 *arXiv:2011.13456*, 2020b.

785

786 Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,  
 787 Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in*  
 788 *neural information processing systems*, 33:3008–3021, 2020.

789

790 Nick T Thomopoulos. *Essentials of Monte Carlo simulation: Statistical methods for building*  
 791 *simulation models*. Springer Science & Business Media, 2012.

792

793 Masatoshi Uehara, Yulai Zhao, Kevin Black, Ehsan Hajiramezanali, Gabriele Scalia, Nathaniel Lee  
 794 Diamant, Alex M Tseng, Tommaso Biancalani, and Sergey Levine. Fine-tuning of continuous-time  
 795 diffusion models as entropy-regularized control. *arXiv preprint arXiv:2402.15194*, 2024.

796

797 Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted langevin algorithm:  
 798 Isoperimetry suffices. *Advances in neural information processing systems*, 32, 2019.

799

800 Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural*  
 801 *computation*, 23(7):1661–1674, 2011.

802

803 Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul,  
 804 Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas  
 805 Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022.

806

807 Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam,  
 808 Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using  
 809 direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision*  
 810 and *Pattern Recognition*, pp. 8228–8238, 2024.

811

812 Jiangshan Wang, Junfu Pu, Zhongang Qi, Jiayi Guo, Yue Ma, Nisha Huang, Yuxin Chen, Xiu Li,  
 813 and Ying Shan. Taming rectified flow for inversion and editing. *arXiv preprint arXiv:2411.04746*,  
 814 2024.

810 Ronald J Williams and Jing Peng. Function optimization using connectionist reinforcement learning  
811 algorithms. *Connection Science*, 3(3):241–268, 1991.  
812

813 Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong  
814 Zhang, Caiming Xiong, et al. A minimalist approach to llm reasoning: from rejection sampling to  
815 reinforce. [arXiv preprint arXiv:2504.11343](https://arxiv.org/abs/2504.11343), 2025.

816 Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun:  
817 Construction of a large-scale image dataset using deep learning with humans in the loop. [arXiv  
818 preprint arXiv:1506.03365](https://arxiv.org/abs/1506.03365), 2015.

819 Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. Slic-hf:  
820 Sequence likelihood calibration with human feedback. [arXiv preprint arXiv:2305.10425](https://arxiv.org/abs/2305.10425), 2023.

821

822 Yixiu Zhao, Jiaxin Shi, Feng Chen, Shaul Druckmann, Lester Mackey, and Scott Linderman. Informed  
823 correctors for discrete diffusion models. [arXiv preprint arXiv:2407.21243](https://arxiv.org/abs/2407.21243), 2024.

824

825 Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. Publaynet: largest dataset ever for document lay-  
826 out analysis. In [2019 International Conference on Document Analysis and Recognition \(ICDAR\)](https://www.iaprtcod.org/2019/),  
827 2019. doi: 10.1109/ICDAR.2019.00166.

828

829 Yuanzhi Zhu, Xingchao Liu, and Qiang Liu. Slimflow: Training smaller one-step diffusion models  
830 with rectified flow. In [European Conference on Computer Vision](https://www.springer.com/10626/100000000/100000000), pp. 342–359. Springer, 2024.

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

## APPENDIX

## A RELATED WORK

868

869

870

871

872

873

874

**Policy Gradient Methods:** Majority of the existing literature on policy gradient methods in the context of generative modeling draw inspiration from Proximal Policy Optimization(PPO) (Schulman et al., 2017) and REINFORCE (Williams & Peng, 1991). PPO based methods in the context of language modeling include Bai et al. (2022); Ouyang et al. (2022); Liu et al. (2023a); Stiennon et al. (2020), whereas frameworks based on REINFORCE include (Li et al., 2023; Ahmadian et al., 2024; Shao et al., 2024; Hu et al., 2025). Policy gradient methods have also been studied in the context of fine-tuning diffusion models (Black et al., 2023; Fan et al., 2023; Ren et al., 2024).

875

876

877

878

879

**Offline Fine-Tuning Methods:** Algorithms which utilize offline preference datasets for fine-tuning generative models have also been widely studied. In the context of language modeling, these include methods like SLiC (Zhao et al., 2023), DPO (Rafailov et al., 2023) and SimPO (Meng et al., 2024). Such methods have also been explored in the context of diffusion models as well - these include methods like Diffusion-DPO (Wallace et al., 2024) and Diffusion-KTO (Li et al., 2024b).

880

881

882

883

**Rejection Sampling Methods:** Recently, many works have explored rejection sampling methods in the context of autoregressive models - these include RSO (Liu et al., 2023b), RAFT (Dong et al., 2023) and Reinforce-Rej (Xiong et al., 2025). In particular, Reinforce-Rej demonstrated that rejection sampling methods can match or even outperform policy gradient methods.

884

885

886

887

888

889

890

891

892

893

**Fine-Tuning Diffusion Models:** Apart from the policy gradient methods discussed already, a host of other methods have also been proposed for fine-tuning diffusion models. Direct reward backpropagation methods include DRaFT (Clark et al., 2023) and AlignProp (Prabhudesai et al., 2023). Note that these methods assume access to a differentiable reward. Uehara et al. (2024) approaches the problem from the lens of entropy-regularized control - however, the method is computationally heavy and requires gradient checkpointing as well as optimizing an additional neural SDE. Domingo-Enrich et al. (2024) proposes a memoryless forward process to overcome the initial value function bias problem for the case of ODEs. PRDP Deng et al. (2024) formulates a supervised learning objective whose optimum matches with the solution to PPO, but with trajectory KL constraint - the supervised objective, with clipping, was found to make the training stable as compared to DDPO.

894

895

896

897

898

899

900

**Score Matching:** Score matching for distribution estimation was first introduced in (Hyvärinen & Dayan, 2005). The algorithm used in this case is called Implicit Score Matching. Diffusion models primarily use Denoising Score Matching (DSM), which is based on Tweedie's formula (Vincent, 2011; Kingma & Cun, 2010). The sample complexity of DSM has been extensively studied in the literature (Kumar et al., 2025; Block et al., 2020; Gupta et al., 2024; Chen et al., 2023). Many alternative procedures such as Sliced Score Matching (Song et al., 2020a) and Target Score Matching (De Bortoli et al., 2024) have been proposed.

901

902

903

904

905

**ODE Reversal in Flow Models:** A prominent use case of ODE reversal in flow models is that of image editing (Hertz et al., 2022; Kim et al., 2022; Hong et al., 2024; Mokady et al., 2023; Rout et al., 2024; Garibi et al., 2024). The reverse ODE has also been used to achieve straighter flows, allowing for faster generation, through 2-rectification/reflow algorithm (Liu et al., 2022; Lee et al., 2024; Zhu et al., 2024; Liu et al., 2023c). Notably, concurrent work Eyring et al. (2025) also proposes a strategy for aligning distilled models by fine-tuning at the noise level.

906

907

908

909

910

911

912

913

914

915

916

917

918 **B ODE SOLVER ALGORITHMS**  
919920 In the backward Euler Algorithm 6, at each time instant  $j$  in the reverse procedure, we solve a fixed  
921 point equation to obtain high precision solution of Eq. equation 5. The step-size  $\eta$  is tuned empirically  
922 so that the recursion does not blow up. Once the step-size is carefully tuned, the iteration converges  
923 to the solution at an exponential rate. In practice, we observed that  $N_b = 10$  is sufficient to obtain  
924 satisfactory results.  
925926 **Algorithm 5** Forward Euler (FWD\_Euler)  
927928 **Input:** Flow model  $v_\theta$ , step-size  $\eta$ , Initial  
929 point  $X_0$   
930  
1: **for**  $j = 0$  to  $\lfloor 1/\eta \rfloor - 1$  **do**  
2:    $X_{j+1} \leftarrow X_j + \eta v_\theta(X_j, \eta j)$   
3: **end for**  
4: **return**  $X_{\lfloor 1/\eta \rfloor}$ 

---

**Algorithm 6** Backward Euler (BWD\_Euler)

---

**Input:** Flow model  $v_\theta$ , step size  $\eta$ , sample  $X^{(i)}$  from  
the dataset, Number of fixed point iterations  $N_b$ 


---

1:  $X_1^{\text{rev}} = X^{(i)}$   
2: **for**  $j = 0$  to  $\lfloor 1/\eta \rfloor - 1$  **do**  
3:    $\hat{X}_0^{\text{rev}} = X_j^{\text{rev}}$   
4:   **for**  $k = 0$  to  $N_b - 1$  **do**  
5:      $\hat{X}_{k+1}^{\text{rev}} \leftarrow X_j^{\text{rev}} - \eta v_\theta(\hat{X}_k^{\text{rev}}, 1 - \eta(j+1))$   
6:   **end for**  
7:    $X_{j+1}^{\text{rev}} \leftarrow \hat{X}_{N_b}^{\text{rev}}$   
8: **end for**  
9: **return**  $X_{\lfloor 1/\eta \rfloor}^{\text{rev}}$ 


---

940 **C GRAFT: ALGORITHM**  
941942 While instantiations of GRAFT are well-known in the literature and are straightforward to implement,  
943 we provide the exact algorithm here for the sake of completeness.  
944

---

**Algorithm 7** GRAFT: Training

---

**Input:** Trainable  $p_\theta$ , Reference  $\bar{p}$ , Reward function  $r$ , Acceptance function  $A$ , Number of sampling  
945 rounds  $N_S$   
946

---

1: Initialize empty set  $\mathcal{D}$   
2: **for**  $i = 0$  to  $N_S$  **do**  
3:   Get  $M$  samples:  $\{X^{(1)}, \dots, X^{(M)}\} \sim \bar{p}$   
4:   Obtain rewards:  $r(X^{(i)})$ ;  $i \in [M]$   
5:   Perform GRS using acceptance function  $A$  to get accepted samples  $\mathcal{A}$   
6:   Perform  $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{A}$   
7: **end for**  
8: Train  $p_\theta$  on  $\mathcal{D}$   
9: **return**  $p_\theta$ 


---

959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

972 **D PROOFS**  
 973

974 **D.1 LEMMA 3.2**  
 975

976 *Proof.* Let  $B$  be any measurable set. Consider the following probability measure:  
 977

978  $\mathbb{P}(X^{(1)} \in B | X^{(1)} \in \mathcal{A}).$   
 979

980 Using Bayes' rule, this measure can be rewritten as:  
 981

982 
$$\mathbb{P}(X^{(1)} \in B | X^{(1)} \in \mathcal{A}) = \frac{\mathbb{P}(X^{(1)} \in B, X^{(1)} \in \mathcal{A})}{\mathbb{P}(X^{(1)} \in \mathcal{A})}.$$
  
 983

984 Recall that  $X^{(1)}$  is drawn from the distribution  $\bar{p}$ . Then, from the definition of  $\mathbb{P}(X^{(1)} \in B, X^{(1)} \in \mathcal{A})$ , we have:  
 985

986 
$$\mathbb{P}(X^{(1)} \in B, X^{(1)} \in \mathcal{A}) = \int_B \mathbb{P}(X^{(1)} \in \mathcal{A} | X^{(1)} = x) d\bar{p}(x).$$
  
 987

988 From Definition 3.1, we know that  $X^{(1)} \in \mathcal{A}$  iff  $C_1 = 1$ . Therefore:  
 989

990 
$$\begin{aligned} \mathbb{P}(X^{(1)} \in B, X^{(1)} \in \mathcal{A}) &= \int_B \mathbb{P}(C_1 = 1 | X^{(1)} = x) d\bar{p}(x) \\ 991 &= \int_B \mathbb{E} [\mathbb{1}(C_1 = 1) | X^{(1)} = x] d\bar{p}(x) \end{aligned}$$
  
 992

993 where  $\mathbb{1}(\cdot)$  denotes the indicator function. Using the tower property of expectations, this can be  
 994 rewritten as:  
 995

996 
$$\begin{aligned} \mathbb{P}(X^{(1)} \in B, X^{(1)} \in \mathcal{A}) &= \int_B \mathbb{E} [\mathbb{E} [\mathbb{1}(C_1 = 1) | X^{(1)} = x, X^{(2)}, \dots, X^{(M)}] | X^{(1)} = x] d\bar{p}(x) \\ 997 &= \int_B \mathbb{E} [\mathbb{P} (C_1 = 1 | X^{(1)} = x, X^{(2)}, \dots, X^{(M)}) | X^{(1)} = x] d\bar{p}(x). \end{aligned}$$
  
 998

999 Note that in the conditional expectation here,  $X^{(1)}, \dots, X^{(M)}$ , are distributed according to  $\bar{p}_0$  since  
 1000  $\{X^{(j)}\}_{j=1}^M$  are i.i.d samples. Again, from Definition 3.1, we know that  
 1001

1002 
$$\mathbb{P}(C_1 = 1 | X^{(1)} = x, \{X^{(j)}\}_{j=2}^M) = A(r(x), \hat{F}_R(r(x)), x, \hat{P}_X)$$
  
 1003

1004 where  $\hat{F}_R$  and  $\hat{P}_X$  are computed using the samples  $\{X^{(j)}\}_{j=1}^n$ . From definition of Radon-Nikodym  
 1005 derivative, the distribution of the accepted samples can therefore be written as:  
 1006

1007 
$$\bar{p}^a(x) = Z_1 \mathbb{E} [A(r(x), \hat{F}_R(r(x)), x, \hat{P}_X) | X^{(1)} = x] \bar{p}(x) \quad (6)$$
  
 1008

1009 where  $Z_1 = 1 / \mathbb{P}(X^{(1)} \in \mathcal{A})$  is a normalizing constant independent of  $x$ . Now, from the method of  
 1010 Lagrangian Multipliers, as mentioned in Section 2, the solution to the PPO optimization objective  
 1011 with reward function  $\hat{r}(\cdot)$  is given by:  
 1012

1013 
$$p^{\text{ppo}}(x) = Z_2 \exp \left( \frac{\hat{r}(x)}{\alpha} \right) \bar{p}(x) \quad (7)$$
  
 1014

1015 where  $Z_2$  is the normalization constant. Comparing equation 6 and equation 7,  $\bar{p}^a = p^{\text{ppo}}$  whenever:  
 1016

1017 
$$\frac{\hat{r}(x)}{\alpha} = \log \left( \mathbb{E} [A(r(x), \hat{F}_R(r(x)), x, \hat{P}_X) | X^{(1)} = x] \right)$$
  
 1018

1019

1020  $\square$

1026 D.2 INSTANTIATIONS OF GRAFT  
10271028 D.2.1 TOP- $K$  OUT OF  $M$  SAMPLING  
10291030 Substituting  $A(\cdot)$  in:

1031 
$$\log \left( \mathbb{E}_{\{X^{(j)}\}_{j=2}^n} \left[ A(r(x), \hat{F}_R(r(x)), x, \hat{P}_X) | X^{(1)} = x, \{X^{(j)}\}_{j=2}^M \right] \right)$$
  
1032

1033 we get:

1034  
1035 
$$\log \left( \int_{\{X^{(j)}\}_{j=2}^n} \mathbb{1}(r(x) \in \text{Top} - \mathsf{K}(r(x), r(x^{(2)}), \dots, r(x^{(M)}))) d\bar{p}(x^{(2)}) \dots d\bar{p}(x^{(M)}) \right)$$
  
1036

1037 where  $\text{Top} - \mathsf{K}(r(x), r(x^{(2)}), \dots, r(x^{(M)}))$  denotes the top- $K$  samples in  
1038  $\{r(x), r(x^{(2)}), \dots, r(x^{(M)})\}$ . Let  $U_K$  denote the event where  $X^{(1)} = x$  ranks in top- $K$   
1039 among the  $M$  samples, where the other  $M - 1$  samples are i.i.d from  $\bar{p}$ . This event can be  
1040 decomposed as:  
1041

1042 
$$U_K = \bigcup_{k=1}^K E_k$$
  
1043

1044 where  $E_k$  denotes the event where  $r(x)$  is the  $k^{\text{th}}$  in the ranked (descending) ordering of rewards.  
1045 Further, note that  $\{E_k\}$  are *mutually exclusive* events. Therefore:  
1046

1047 
$$\mathbb{P}(U_K) = \sum_{k=1}^K \mathbb{P}(E_k)$$
  
1048  
1049

1050 **Computing  $\mathbb{P}(E_k)$ :** If  $x$  ranks  $k^{\text{th}}$  when ranked in terms of rewards, there are  $k - 1$  samples  
1051 which have *higher* rewards than  $x$  and  $M - k$  samples which have *lower* rewards than  $x$ . Thus, the  
1052 required probability can be computed by finding the probability of having  $K - 1$  samples having  
1053 higher rewards and the rest having lower rewards. Note that the ordering within the  $K - 1$  group or  
1054  $M - K$  group doesn't matter. The probability of any one sample having a higher reward than  $r(x)$   
1055 is  $1 - F(r(x))$  and having a lower reward is  $F(r(x))$ . Therefore, the required probability can be  
1056 computed as:  
1057

1058 
$$\mathbb{P}(E_k) = \binom{M-1}{k-1} (1 - F(r(x)))^{k-1} (F(r(x)))^{M-k}$$
  
1059

1060 And hence:  
1061

1062 
$$\mathbb{P}(U_K) = \sum_{k=0}^{K-1} \binom{M-1}{k} (1 - F(r(x)))^k (F(r(x)))^{M-k-1}$$
  
1063  
1064

1065 Therefore:

1066 
$$\frac{\hat{r}(x)}{\alpha} = \log \left( \sum_{k=0}^{K-1} \binom{M-1}{k} (1 - F(r(x)))^k (F(r(x)))^{M-k-1} \right)$$
  
1067  
1068

1069 It is straightforward to check that this is an increasing function in  $r$ .  
10701072 D.2.2 PREFERENCE REWARDS  
10731074 Substituting  $A(\cdot)$  in:

1075 
$$\log \left( \mathbb{E}_{X^{(2)}} A(r(x), \hat{F}_R(r(x)), x, \hat{P}_X) | X^{(1)} = x, X^{(2)} \right)$$
  
1076

1077 we get:  
1078

1079 
$$\log \left( \int_{X^{(2)}} \mathbb{1}(r(x^{(2)}) \leq r(x)) d\bar{p}(x^{(2)}) \right) = \log F(r(x))$$

1080 D.3 LEMMA 4.2  
10811082 *Proof.* Let  $B$  be any measurable set. Consider the following probability measure:  
1083

1084 
$$\mathbb{P}(X_t^{(1)} \in B | X_t^{(1)} \in \mathcal{A}).$$
  
1085

1086 Using Bayes' rule, this measure can be rewritten as:  
1087

1088 
$$\mathbb{P}(X_t^{(1)} \in B | X_t^{(1)} \in \mathcal{A}) = \frac{\mathbb{P}(X_t^{(1)} \in B, X_t^{(1)} \in \mathcal{A})}{\mathbb{P}(X_t^{(1)} \in \mathcal{A})}.$$
  
1089

1090 Recall that  $X_t^{(1)}$  is drawn from the distribution  $\bar{p}_t$ . Then, from the definition of  $\mathbb{P}(X_t^{(1)} \in B, X_t^{(1)} \in \mathcal{A})$ , we have:  
1091

1092 
$$\mathbb{P}(X_t^{(1)} \in B, X_t^{(1)} \in \mathcal{A}) = \int_B \mathbb{P}(X_t^{(1)} \in \mathcal{A} | X_t^{(1)} = x) d\bar{p}_t(x).$$
  
1093

1094 From Definition 4.1, we know that  $X_t^{(1)} \in \mathcal{A}$  iff  $C_1 = 1$ . Therefore:  
1095

1096 
$$\begin{aligned} \mathbb{P}(X_t^{(1)} \in B, X_t^{(1)} \in \mathcal{A}) &= \int_B \mathbb{P}(C_1 = 1 | X_t^{(1)} = x) d\bar{p}_t(x) \\ &= \int_B \mathbb{E} [\mathbb{1}(C_1 = 1) | X_t^{(1)} = x] d\bar{p}_t(x) \end{aligned}$$
  
1097

1098 where  $\mathbb{1}(\cdot)$  denotes the indicator function. Using the tower property of expectations, this can be  
1099 rewritten as:  
1100

1101 
$$\begin{aligned} \mathbb{P}(X_t^{(1)} \in B, X_t^{(1)} \in \mathcal{A}) &= \int_B \mathbb{E} [\mathbb{E} [(\mathbb{1}(C_1 = 1) | X_t^{(1)} = x, X_0^{(1)}, X_0^{(2)}, \dots, X_0^{(M)})] | X_t^{(1)} = x] d\bar{p}_t(x) \\ &= \int_B \mathbb{E} [\mathbb{P} (C_1 = 1 | X_t^{(1)} = x, X_0^{(1)}, X_0^{(2)}, \dots, X_0^{(M)}) | X_t^{(1)} = x] d\bar{p}_t(x). \end{aligned}$$
  
1102

1103 Note that in the conditional expectation here,  $X_0^{(2)}, \dots, X_0^{(M)}$ , are distributed according to  $\bar{p}_0$   
1104 since  $\{X_0^{(j)}\}_{j=1}^n$  are i.i.d samples. However,  $X_0^{(1)}$  is distributed according to  $\bar{p}_{0|t}$  because of the  
1105 conditioning on  $X_t^{(1)}$ . Again, from Definition 4.1, we know that

1106 
$$\mathbb{P}(C_1 = 1 | X_t^{(1)} = x, \{X_0^{(j)}\}_{j=1}^M) = A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})), X_0^{(1)}, \hat{P}_X)$$
  
1107

1108 where  $\hat{F}_R$  and  $\hat{P}_X$  are computed using the samples  $\{X^{(j)}\}_{j=1}^M$ . From the definition of Radon-  
1109 Nikodym derivative, the density of the accepted samples can therefore be written as:  
1110

1111 
$$\bar{p}_t^a(x) = Z_1 \mathbb{E} [A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})), X_0^{(1)}, \hat{P}_X | X_t^{(1)} = x)] \bar{p}_t(x) \quad (8)$$
  
1112

1113 where  $Z_1 = 1 / \mathbb{P}(X_t^{(1)} \in \mathcal{A})$  is a normalizing constant independent of  $x$ . Now, from the method of  
1114 Lagrangian Multipliers, as mentioned in Section 2, the solution to the PPO optimization objective  
1115 (with reward function  $\hat{r}(\cdot)$ ) is (where  $Z_2$  is the normalization constant):  
1116

1117 
$$p^{\text{ppo}}(x) = Z_2 \exp \left( \frac{\hat{r}(x)}{\alpha} \right) \bar{p}_t(x). \quad (9)$$
  
1118

1119 Comparing equation 8 and equation 9,  $\bar{p}_t^a = p^{\text{ppo}}$  whenever:  
1120

1121 
$$\frac{\hat{r}(x)}{\alpha} = \log \left( \mathbb{E} [A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})), X_0^{(1)}, \hat{P}_X | X_t^{(1)} = x)] \right) \quad (10)$$
  
1122

1123  $\square$   
1124

1134 D.3.1 DISTRIBUTION INDUCED BY THE STITCHED MODEL  
1135

1136 Let us denote the distribution the stitched model samples from as  $p^s$ . From the discussion above, the  
1137 distribution the P-GRAFT fine-tuned model samples from at time  $t$  is  $\bar{p}_t^a$ . Further, let  $\bar{p}_{0|t}$  denote the  
1138 distribution of samples at time 0, given a sample at time  $t$  under the base model. Then clearly:

$$1139 \quad p^s(x_0) = \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t^a(x_t) dx_t \\ 1140 \\ 1141 \quad = Z \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) \exp\left(\frac{\hat{r}(x_t)}{\alpha}\right) dx_t \\ 1142 \\ 1143$$

1144 where  $Z$  is a normalization constant independent of  $x_0$ . In general, an explicit solution to this integral  
1145 cannot be computed.

1146 To get some intuition about the tilt, we analyze two edge cases. Assume that the acceptance probability  
1147 depends only on  $r(X_0)$  and  $\hat{F}_R(r(X_0))$ .

1148 **Case 1:** The reward of a sample at time 0 is independent of the latent at time  $t$   $x_t$ .

1149 The acceptance probability is also independent of the latent at time  $t$   $x_t$ .

1150 From equation 10,

1151

$$1152 \quad \frac{\hat{r}(x)}{\alpha} = \log\left(\mathbb{E}\left[A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})))|X_t^{(1)} = x\right]\right) \\ 1153 \\ 1154 \quad = \log\left(\mathbb{E}\left[A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})))\right]\right) \\ 1155 \\ 1156$$

1157 since acceptance probability is independent of  $x_t$ . Note that the expectation is over  $X_0^{(1)}, \dots, X_0^{(M)}$ ,  
1158 as discussed above. Hence, because of the independence assumption,  $\log(\mathbb{E}[A(\cdot)])$  is independent of  
1159  $X_t^{(1)}$  and hence  $\frac{\hat{r}(x)}{\alpha}$  is independent of  $x$ . Let us denote this quantity as  $Z_1$ . Then:

1160

$$1161 \quad p^s(x_0) = Z \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) \exp(Z_1) dx_t \\ 1162 \\ 1163 \quad = Z \exp(Z_1) \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) dx_t \\ 1164 \\ 1165 \quad = Z \exp(Z_1) \bar{p}(x_0) \\ 1166 \\ 1167 \quad = \bar{p}(x_0) \\ 1168$$

1169 where we have used the fact that  $\bar{p}(x_0)$  is the normalized reference distribution. Hence, if acceptance  
1170 probability of a sample at time 0 is independent of the latent at time  $t$   $x_t$ , *the stitched model results in*  
1171 *no tilt whatsoever*.

1172 **Case 2:** The reward of a sample at time 0 is completely determined by the latent at time  $t$   $x_t$ .

1173

1174 A unique mapping  $r_t$  exists such that  $r(X_0) = r_t(X_t)$ , where  $X_0 \sim p_{0:t}(\cdot|X_t)$ . Therefore, from  
1175 equation 10,

1176

$$1177 \quad \frac{\hat{r}(x)}{\alpha} = \log\left(\mathbb{E}\left[A(r(X_0^{(1)}), \hat{F}_R(r(X_0^{(1)})))|X_t^{(1)} = x\right]\right) \\ 1178 \\ 1179 \quad = \log\left(\mathbb{E}\left[A(r_t(x), \hat{F}_R(r_t(x)))\right]\right) \\ 1180$$

1181 Note that the expectation here is only over  $X_0^{(2)}, \dots, X_0^{(M)}$  because of the conditioning. And hence:

1182

$$1183 \quad p^s(x_0) = Z \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) \exp\left(\frac{\hat{r}(x_t)}{\alpha}\right) dx_t \\ 1184 \\ 1185 \quad = Z \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) \mathbb{E}\left[A(r_t(x_t), \hat{F}_R(r_t(x_t)))\right] dx_t \\ 1186 \\ 1187$$

1188 Using the fact that  $r_t(x_t) = r(x_0)$ , we have:

$$\begin{aligned}
 1190 \quad p^s(x_0) &= Z \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) \mathbb{E} \left[ A(r(x_0), \hat{F}_R(r(x_0))) \right] dx_t \\
 1191 &= Z \mathbb{E} \left[ A(r(x_0), \hat{F}_R(r(x_0))) \right] \int \bar{p}_{0|t}(x_0|x_t) \bar{p}_t(x_t) dx_t \\
 1192 &= Z \bar{p}(x_0) \mathbb{E} \left[ A(r(x_0), \hat{F}_R(r(x_0))) \right]
 \end{aligned}$$

1196 Comparing  $p^s(x_0)$  with Lemma 3.2, we see that in this case, P-GRAFT results in sampling from the  
1197 exact same distribution as GRAFT.

1198 For the cases in between case 1 and case 2, P-GRAFT interpolates between the reference distribution  
1199 and the distribution induced by GRAFT.

### 1201 D.3.2 JUSTIFICATION FOR KL REGULARIZATION AT INTERMEDIATE Timestep

1203 Intuitively, KL regularization at an intermediate time-step ensures that the generated latents at this  
1204 time-step remain close (in distribution) to the latents of the reference model. Since the rest of  
1205 denoising happens with the reference model, this also ensures that the distribution of final generated  
1206 samples remains close to the sample distribution of the reference model.

1207 More formally, suppose the fine-tuned model obeys  $KL(p_t || \bar{p}_t) \leq \alpha$ , where  $p_t$  is the fine-tuned  
1208 model and  $\bar{p}_t$  is the reference model. Let the final distribution induced by the stitched model be  $p_0^s$ .  
1209 Due to the contraction of KL divergences under Markov Chains, a.k.a Data Processing Inequality  
1210 (Theorem 7.4 in Polyanskiy & Wu (2025)),  $KL(p_0^s || \bar{p}_0) \leq \alpha$  because we follow the same Markov  
1211 Chain from time  $t$  to 0 during denoising due to stitching.

### 1213 D.4 PROOF OF LEMMA 4.3

1215 *Proof.* Let  $s > t$ . Note that  $X_s \rightarrow X_t \rightarrow X_0$  forms a Markov chain. By the law of total variance,  
1216 we have for any random variables  $Y, Z$ :

$$\begin{aligned}
 1217 \quad \text{Var}(Z) &= \mathbb{E} \text{Var}(Z|Y) + \text{Var}(\mathbb{E}[Z|Y]) \\
 1218 &\geq \mathbb{E} \text{Var}(Z|Y)
 \end{aligned} \tag{11}$$

1220 Given  $X_s$ , Suppose  $Z, Y$  be jointly distributed as the law of  $(r(X_0), X_t)$ . Then, we have  $X_s$  almost  
1221 surely:

$$1224 \quad \text{Var}(r(X_0)|X_s) \geq \mathbb{E}[\text{Var}(r(X_0)|X_s, X_t)|X_s] = \mathbb{E}[\text{Var}(r(X_0)|X_t)|X_s] \tag{12}$$

1226 In the last line, we have used the Markov property to show that the law of  $r(X_0)|X_s, X_t$  is the same  
1227 as the law of  $r(X_0)|X_t$  almost surely. We conclude the result by taking expectation over both the  
1228 sides.  $\square$

### 1229 D.5 PROOF OF THEOREM 4.4

1231 *Proof.* We will follow the exposition in Vempala & Wibisono (2019) for our proofs.  $q_t$  converges  
1232 to  $q_\infty$  as  $t \rightarrow \infty$ . By (Vempala & Wibisono, 2019, Lemma 2) applied to the forward process, we  
1233 conclude that:

$$\begin{aligned}
 1235 \quad \frac{d}{dt} \text{KL}(q_t || q_\infty) &= - \int_{\mathbb{R}^d} dX q_t(X) \|\nabla \log q_t(X) - \nabla \log q_\infty(X)\|^2 \\
 1236 \\
 1238 &\implies \int_t^T dt \int_{\mathbb{R}^d} dX q_s(X) \|\nabla \log q_s(X) - \nabla \log q_\infty(X)\|^2 = \text{KL}(q_t || q_\infty) - \text{KL}(q_T || q_\infty) \tag{13}
 \end{aligned}$$

1240 For brevity, we call the LHS to be  $H_t^T$ . Clearly,

$$1241 \quad H_t^T - e^{-2t} H_0^T = \text{KL}(q_t || q_\infty) - e^{-2t} \text{KL}(q_0 || q_\infty) + \text{KL}(q_T || q_\infty)(e^{-2t} - 1).$$

Notice that  $q_\infty$  is the density of the standard Gaussian random variable. Therefore, it satisfies the Gaussian Logarithmic Sobolev inequality Gross (1975). Thus, we can apply (Vempala & Wibisono, 2019, Theorem 4) to conclude that for every  $s \geq 0$ ,  $\text{KL}(q_s || q_\infty) \leq e^{-2s} \text{KL}(q_0 || q_\infty)$ . Thus,

$$H_t^T \leq \frac{e^{-2t}}{1 - e^{-2t}} H_0^t$$

□

## D.6 PROOF OF LEMMA 5.1

The uniqueness and the convergence of fixed point iteration for implicit Euler methods have been established under great generality in Butcher (2016). However, we give a simpler proof for our specialized setting here.

1. Consider the update for the backward Euler iteration at each time step  $t = \eta i$

$$\hat{x}_{\eta i}^{\text{rev}} \rightarrow \hat{x}_{\eta(i-1)}^{\text{rev}} - \eta v_\theta(\hat{x}_{\eta i}^{\text{rev}}, 1 - \eta(i-1))$$

Let us define an operator  $T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  such that

$$T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}(x) = \hat{x}_{\eta(i-1)}^{\text{rev}} - \eta v_\theta(x, 1 - \eta(i-1))$$

First, we will show that  $T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}$  as defined above is a contractive operator under the condition  $\eta L < 1$ . Then, one can use Banach fixed point theorem to establish uniqueness of the solution and obtain the solution through fixed point iteration. To this end, consider two point  $x_1$  and  $x_2$  in  $\mathbb{R}^d$  and apply  $T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}$  to them

$$\begin{aligned} \|T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}(x_1) - T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}(x_2)\|_2 &= \eta \|v_\theta(x_1, 1 - \eta(i-1)) - v_\theta(x_2, 1 - \eta(i-1))\|_2 \\ &\leq \eta L \|x_1 - x_2\|_2. \end{aligned}$$

Since  $\eta L < 1$ , we conclude that  $T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}$  is a contractive operator. Thus, by Banach fixed point theorem, the fixed point equation  $T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}(x) = x$  has a unique solution for each step  $t = \eta i$ . To obtain the solution to the backward Euler update, we use the Banach fixed point method, i.e., start with  $x_{(0)} = \hat{x}_{\eta(i-1)}^{\text{rev}}$  (or any arbitrary point in  $\mathbb{R}^d$ ) and run the iteration  $x_{(k+1)} = T_{\theta, \eta}^{\hat{x}_{\eta(i-1)}^{\text{rev}}}(x_{(k)})$ . Then,  $\lim_{k \rightarrow \infty} x_{(k)} = \hat{x}_{\eta i}^{\text{rev}}$ .

2. The invertibility of the operator  $T_{\theta, \eta}$  follows directly from the previous part. Since the solution for the backward Euler method is unique at each time step  $t = \eta i$ , it implies that there exists a one-to-one mapping between sample points  $x_0^{\text{rev}}$  and  $x_1^{\text{rev}}$ .

## D.7 PROOF OF LEMMA 5.2

Before starting the proof of this lemma, we will state the following well-known theorem from information theory.

**Theorem D.1.** [Data Processing Inequality] Let  $\mathcal{X}$  and  $\mathcal{Y}$  be two sample spaces. Denote  $\mathcal{P}(\mathcal{X})$  and  $\mathcal{P}(\mathcal{Y})$  as the set of all possible probability distributions on  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively. Let  $P_X, Q_X \in \mathcal{P}(\mathcal{X})$  and  $P_{Y|X}$  be a transition kernel. Denote  $P_Y$  and  $Q_Y$  to be the push through, i.e.,  $P_Y(B) = \int_{\mathcal{X}} P_{Y|X}(B|X=x) dP_X(x)$ . Then, for any  $f$ -divergence we have

$$D_f(P_X || Q_X) \geq D_f(P_Y || Q_Y) \tag{14}$$

1. By part 3 of Lemma 5.1, we have that  $x_1^{\text{rev}} = T_{\theta, \eta}^{-1}(x_0^{\text{rev}}) \implies T_{\theta, \eta}(x_1^{\text{rev}}) = x_0^{\text{rev}}$ . Suppose  $x_0^{\text{rev}} \sim p^*$ , then by definition,  $x_1^{\text{rev}} \sim p_1^{\text{rev}}$ . This concludes the result.

1296     2. Recall that TV-norm is an  $f$ -divergence. Furthermore,  $T_{\theta,\eta}$  is the push forward function  
 1297     from  $p_0$  and  $p_1^{\text{rev}}$  to  $p_1$  and  $p^*$ , respectively. Thus, using DPI D.1, we have  
 1298

$$1299 \quad \text{TV}(p_1^{\text{rev}}, p_0) \geq \text{TV}(p^*, p_1).$$

1300  
 1301     Additionally,  $T_{\theta,\eta}$  is an invertible mapping. Hence,  $T_{\theta,\eta}^{-1}$  can also be viewed as the push  
 1302     forward function from  $p_1$  and  $p^*$  to  $p_0$  and  $p_1^{\text{rev}}$ , respectively. Thus, again using DPI D.1, we  
 1303     get

$$1304 \quad \text{TV}(p^*, p_1) \geq \text{TV}(p_1^{\text{rev}}, p_0).$$

1305  
 1306     Combining both the bounds, we get the desired claim.

1307  
 1308     3. KL divergence is also a valid  $f$ -divergence. Thus, repeating the arguments from the previous  
 1309     part, one gets the desired equality.

## 1311     D.8 THEORETICAL JUSTIFICATION FOR INVERSE NOISE CORRECTION

1312  
 1313     In this section, our goal is to provide a theoretical justification for inverse noise correction in the  
 1314     context of flow models. Specifically, we will argue that if  $\text{KL}(p^X \parallel \mathcal{N}(0, \mathbf{I}))$  is small, then it is  
 1315     less challenging to learn the score function corresponding to  $p_t$  and thereby the velocity field  $v_t^X$   
 1316     governing the rectified flow. To this end, let  $X$  be a sample from a distribution  $p^X$  and  $Z, Y$  be  
 1317     standard normal random variables all independent of each other. Consider the following two linear  
 1318     interpolations:

$$1319 \quad X_t = tX + (1-t)Z \quad (15a)$$

$$1320 \quad Y_t = tY + (1-t)Z. \quad (15b)$$

1321  
 1322     Denote  $p_t$  and  $q_t$  as the distribution of  $X_t$  and  $Y_t$ , respectively. Then, it is easy to verify that they  
 1323     satisfy the following continuity equations:

$$1325 \quad \dot{p}_t + \nabla \cdot (v_t^X p_t) = 0 \quad (16a)$$

$$1326 \quad \dot{q}_t + \nabla \cdot (v_t^Y q_t) = 0 \quad (16b)$$

1327  
 1328     where  $v_t^X(x) = \mathbb{E}[X - Z | X_t = x]$  and  $v_t^Y(x) = \mathbb{E}[Y - Z | Y_t = x]$ . Then, we have the following  
 1329     theorem which establishes the relation between KL-divergence of  $p_1$  and  $q_1$  in terms of the velocities  
 1330      $v_t^X$  and  $v_t^Y$ . The proof for the theorem is provided in Section D.9.

1331     **Theorem D.2.** *Let  $p_t$  and  $q_t$  be the distribution of  $X_t$  and  $Y_t$  defined in equation 15. Then, the  
 1332     KL-divergence between  $p_1$  and  $q_1$  satisfy the following relation*

$$1334 \quad \text{KL}(p_1 \parallel q_1) = \text{KL}(p^X \parallel \mathcal{N}(0, \mathbf{I})) = \int_0^1 \frac{t}{1-t} \int_{\mathbb{R}^d} p_t(x) \|v_t^X(x) - v_t^Y(x)\|^2 dx dt. \quad (17)$$

1337  
 1338     Now, consider the distribution of the inverse noise  $p_1^{\text{rev}}$  obtained by iterating equation 5 and substitute  
 1339     it with  $p^X$  in the theorem above. Suppose that the flow model is trained such that  $\text{KL}(p^{\text{data}} \parallel p_1) \leq \epsilon$ .  
 1340     Then, by Lemma 5.2 it follows that  $\text{KL}(p_1^{\text{rev}} \parallel p_0) \leq \epsilon$ . Combining this observation with equation 17,  
 1341     it is easy to see that the velocities  $v_t^X(x)$  and  $v_t^Y(x)$  should be close to each other. Additionally, since  
 1342      $q_t$  simply corresponds to learning a flow model from standard Gaussian to itself, we can explicitly  
 1343     compute  $v_t^Y$  as follows:

$$1344 \quad v_t^Y(x) = \frac{x}{t} + \frac{1-t}{t} \frac{-x}{(1-t)^2 + t^2}$$

$$1345 \quad = \frac{x(2t-1)}{(1-t)^2 + t^2}.$$

1346  
 1347     Thus,  $v_t^Y(x)$  is a linear function of  $x$  and a rational function of  $t$ . Because  $\text{KL}(p_1^{\text{rev}} \parallel p_0) \leq \epsilon$ , Theorem  
 1348     D.2 suggests that learning  $v_t^X$  from data should be relatively easier as it is close to  $v_t^Y$ .

1350 D.9 PROOF OF THEOREM D.2  
13511352 Then, the time derivative of the KL-divergence between  $p_t$  and  $q_t$  is given by  
1353

1354 
$$\frac{d\text{KL}(p_t||q_t)}{dt} = \int_{\mathbb{R}^d} \frac{d}{dt} \left( p_t(x) \log \left( \frac{p_t(x)}{q_t(x)} \right) \right) dx$$
  
1355 
$$= \int_{\mathbb{R}^d} \left( \dot{p}_t(x) \log \left( \frac{p_t(x)}{q_t(x)} \right) + p_t(x) \frac{d}{dt} \log(p_t(x)) - p_t(x) \frac{d}{dt} \log(q_t(x)) \right) dx$$
  
1356 
$$= \int_{\mathbb{R}^d} \left( \dot{p}_t(x) \log \left( \frac{p_t(x)}{q_t(x)} \right) + \dot{p}_t(x) - p_t(x) \frac{\dot{q}_t(x)}{q_t(x)} \right) dx$$
  
1357  
1358  
1359  
1360

1361 We will consider each term separately as  $T_1, T_2$  and  $T_3$ . For  $T_1$  using the continuity equation, we  
1362 have  
1363

1364 
$$T_1 = \int_{\mathbb{R}^d} \dot{p}_t(x) \log \left( \frac{p_t(x)}{q_t(x)} \right) dx$$
  
1365 
$$= - \int_{\mathbb{R}^d} \nabla \cdot (v_t^X(x) p_t(x)) \log \left( \frac{p_t(x)}{q_t(x)} \right) dx$$
  
1366 
$$= \int_{\mathbb{R}^d} p_t(x) \left\langle v_t^X(x), \nabla \log \left( \frac{p_t(x)}{q_t(x)} \right) \right\rangle dx \quad (\text{Integration by parts})$$
  
1367  
1368  
1369  
1370

1371 Note that  $\int_{\mathbb{R}^d} p_t(x) dx = 1$  for all  $t \in [0, 1]$ . Thus for  $T_2$ , we obtain  
1372

1373 
$$T_2 = \int_{\mathbb{R}^d} \dot{p}_t(x) dx = \frac{d}{dt} \int_{\mathbb{R}^d} p_t(x) dx = \frac{d}{dt} 1 = 0.$$
  
1374

1375 For the final term  $T_3$ , we again use the continuity equation to get  
1376

1377 
$$T_3 = - \int_{\mathbb{R}^d} p_t(x) \frac{\dot{q}_t(x)}{q_t(x)} dx$$
  
1378 
$$= \int_{\mathbb{R}^d} p_t(x) \frac{\nabla \cdot (v_t^Y q_t)}{q_t(x)} dx$$
  
1379 
$$= - \int_{\mathbb{R}^d} q_t(x) \left\langle \nabla \left( \frac{p_t(x)}{q_t(x)} \right), v_t^Y(x) \right\rangle dx \quad (\text{Integration by parts})$$
  
1380  
1381  
1382  
1383  
1384  
1385

1386 Combining all the terms above, we get  
1387

1388 
$$\frac{d\text{KL}(p_t||q_t)}{dt} = \int_{\mathbb{R}^d} p_t(x) \left\langle \nabla \log \left( \frac{p_t(x)}{q_t(x)} \right), v_t^X(x) - v_t^Y(x) \right\rangle dx. \quad (18)$$
  
1389

1390 To obtain an expression for score function in terms of the velocity vector, we use Tweedie's formula  
1391 Efron (2011) which leads us to  
1392

1393 
$$\mathbb{E}[X - Z|X_t = x] = \frac{1}{1-t} \mathbb{E}[X - X_t|X_t = x]$$
  
1394 
$$= \frac{1}{1-t} \mathbb{E}[X|X_t = x] - \frac{x}{1-t}$$
  
1395 
$$= \frac{1}{t(1-t)} (x + (1-t)^2 \nabla \log p_t(x)) - \frac{x}{1-t} \quad (\text{Tweedie's Formula})$$
  
1396  
1397  
1398  
1399  
1400

1401 Similarly, we obtain  
1402

1403 
$$v_t^Y(x) = \mathbb{E}[Y - Z|Y_t = x] = \frac{x}{t} + \frac{1-t}{t} \nabla \log q_t(x). \quad (20)$$

1404  
1405

Plugging in the expressions for the score functions into equation 18, we obtain

1406  
1407

$$\frac{d\text{KL}(p_t||q_t)}{dt} = \int_{\mathbb{R}^d} p_t(x) \left\langle \frac{t}{1-t} (v_t^X(x) - v_t^Y(x)), v_t^X(x) - v_t^Y(x) \right\rangle dx$$

1408  
1409

$$= \frac{t}{1-t} \int_{\mathbb{R}^d} p_t(x) \|v_t^X(x) - v_t^Y(x)\|^2 dx$$

1410  
1411

$$\implies \text{KL}(p_1||q_1) - \text{KL}(p_0||q_0) = \int_0^1 \frac{t}{1-t} \int_{\mathbb{R}^d} p_t(x) \|v_t^X(x) - v_t^Y(x)\|^2 dx dt.$$

1412

Recall that  $p_0 = q_0 = q_1 = \mathcal{N}(0, \mathbf{I})$  and  $p_1 = p^X$ . Thus, we get the desired claim1413  
1414

$$\text{KL}(p^X||\mathcal{N}(0, I)) = \int_0^1 \frac{t}{1-t} \int_{\mathbb{R}^d} p_t(x) \|v_t^X(x) - v_t^Y(x)\|^2 dx dt.$$

1415  
1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458 E TEXT-TO-IMAGE GENERATION  
14591460 E.1 ABLATIONS  
14611462 E.1.1 DIFFERENT CHOICES OF  $K$  AND  $M$   
1463

1464 We report results for various choices of  $K$  and  $M$  for Top –  $K$  of  $M$  sampling for GenAI-Bench  
1465 in Tables 6, 7 and 8. Note that these models are also trained on GenAI-Bench. We also report the  
1466 (mean)score separately for the “Basic” and “Advanced” split in the prompt set. Results for Top – 10  
1467 of 100 sampling for T2I-CompBench++ is given in Table 9. Models for Table 9 were trained on  
1468 the train split of T2I-CompBench++. All results are consistent with the developed theory: both  
1469 GRAFT and P-GRAFT outperform base SDv2 and P-GRAFT, for an appropriate choice of  $N_I$  always  
1470 outperform GRAFT.

1471 Table 6: VQAScore on GenAI-Bench for  $K = 1$  and  $M = 4$   
1472

Model	Basic	Advanced	Mean
SD v2	74.83	59.19	66.32
GRAFT	77.33	62.76	69.41
P-GRAFT (0.8 $N$ )	76.30	62.18	68.62
P-GRAFT (0.5 $N$ )	78.57	63.38	<b>70.32</b>

1478 Table 7: VQAScore on GenAI-Bench for  $K = 1$  and  $M = 100$   
1479

Model	Basic	Advanced	Mean
SD v2	74.83	59.19	66.32
GRAFT	79.61	64.26	71.2
P-GRAFT (0.75 $N$ )	76.02	62.91	68.89
P-GRAFT (0.5 $N$ )	78.68	64.5	70.97
P-GRAFT (0.25 $N$ )	80.05	64.85	<b>71.79</b>

1489 Table 8: VQAScore on GenAI-Bench for  $K = 25$  and  $M = 100$   
1490

Model	Basic	Advanced	Mean
SD v2	74.83	59.19	66.32
GRAFT	78.01	63.31	70.02
P-GRAFT (0.75 $N$ )	77.36	63.33	69.73
P-GRAFT (0.5 $N$ )	78.18	64.28	70.62
P-GRAFT (0.25 $N$ )	78.77	65.29	<b>71.44</b>

1498 Table 9: VQAScore on T2I-CompBench++ (Val) for  $K = 10$  and  $M = 100$   
1499

Model	Mean
SD v2	69.76
GRAFT	74.66
P-GRAFT (0.25 $N$ )	<b>75.16</b>

1506 E.1.2 CONDITIONAL VARIANCE OF REWARD FOR TEXT-TO-IMAGE GENERATION  
1507

1508 While experimental results in Table 2 already demonstrate the bias-variance tradeoff, we provide  
1509 further evidence of Lemma 4.3 in the context of text-to-image generation. We evaluate conditional  
1510 variance of VQAReward scores of the base SDv2 model in GenAI-Bench. We follow the methodology  
1511 as described in Section 4.1 except that we generate 4 images per prompt for a total of 1600 prompts.  
1512 The results are given in Table 10. It can be seen that even at  $N_I = 0.75N$ , the expected conditional

1512 variance of the reward is significantly smaller than at  $t_N$ . This explains why even  $N_I = 0.75N$  gives  
 1513 a significant gain over the base model as seen in Table 2.  
 1514

1515 Table 10: Expected conditional variance for T2I generation  
 1516

$N_I$	$\mathbb{E} [\text{Var}(r(X_0) X_{t_n})]$
$N$	0.0193
$3N/4$	0.0080
$N/2$	0.0039
$N/4$	0.0019

1523 E.1.3 EFFECT OF LORA RANK  
 1524

1525 We increase the LoRA rank used for fine-tuning and check the impact on the performance. Table 11  
 1526 shows that increasing LoRA rank does not seem to affect performance, indicating that the default  
 1527 LoRA rank is sufficient. Ablations are done on GenAI-Bench with  $M = 100, K = 1$ .  
 1528

1529 Table 11: Effect of LoRa Rank  
 1530

Model	Rank	Mean Reward
P-GRAFT (0.5N)	4	70.97
	6	70.87
	8	70.57
	10	70.84
P-GRAFT (0.25N)	4	71.79
	6	71.84
	8	71.49
	10	71.63

1541 E.1.4 REVERSE STITCHING  
 1542

1543 In P-GRAFT, we always use the fine-tuned model for the first  $(N - N_I)$  steps and then switch to  
 1544 the reference model. We experiment with a reverse stitching strategy, where we use the reference  
 1545 model for the earlier denoising steps and fine-tuned model for the later denoising steps. For switching  
 1546 timestep  $N_I$ , we denote this strategy as RP-GRAFT ( $N_I$ ) - i.e. RP-GRAFT (0.75N) indicates that  
 1547 the base model will be used from  $t_N$  to  $t_{0.75N}$ , after which the fine-tuned model will be used. From  
 1548 Table 12, we observe that this strategy is significantly worse when compared to P-GRAFT - this  
 1549 provides further evidence of the bias-variance tradeoff. Ablations are done with  $M = 100, K = 1$ .  
 1550

1551 Table 12: Ablations on reverse stitching  
 1552

Model	Basic	Advanced	Mean
SDv2	74.83	59.19	66.32
GRAFT	79.61	64.26	71.20
RP-GRAFT (0.75N)	79.23	62.63	70.20
RP-GRAFT (0.5N)	76.60	60.87	68.05
RP-GRAFT (0.25N)	75.74	59.76	67.05

1559 E.2 IMPLEMENTATION DETAILS  
 1560

1561 Since we require samples only from the pre-trained model, sampling and training can be done  
 1562 separately. Therefore, we first perform rejection sampling according to Top – K of M for the chosen  
 1563 values of  $K$  and  $M$ . The selected samples are then used as the dataset for training. If not mentioned  
 1564 explicitly, hyperparameters can be assumed to be the default values for SD 2.0 in the Diffusers library  
 1565 (von Platen et al., 2022).

1566

**Training on GenAI-Bench:**

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585

1586

1587

1588

The hyperparameters for sampling and training are given in Table 13 and Table 14 respectively. Note that one training epoch is defined as one complete pass over the training dataset. The size of the training dataset depends on the chosen  $K$  and  $M$ . For instance,  $K = 10$  and  $M = 100$  results in 10 images per-prompt, for a total of 16000 images. One training epoch corresponds to a single pass over these 16000 images, which with a batch size of 8 corresponds to 2000 iterations per epoch.

Table 13: Sampling hyperparameters for GenAI-Bench

Sampling Steps	50
Scheduler	EulerDiscreteScheduler
Guidance Scale	7.5

Table 14: Training hyperparameters for GenAI-Bench

Training Epochs	10
Image Resolution	$768 \times 768$
Batch Size	8
Learning Rate	$10^{-4}$
LR Schedule	Constant
LoRA Fine-Tuning	True

**Training on T2I-CompBench++:**

1589

The hyperparameters for sampling and training are given in Table 15 and Table 16 respectively. We use different sampling schedulers for the two datasets to ensure that our results hold irrespective of the choice of the scheduler.

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

Table 15: Sampling hyperparameters for T2I-CompBench++

Sampling Steps	50
Scheduler	DDIMScheduler
$\eta$ (DDIMScheduler specific hyperparameter )	1.0
Guidance Scale	7.5

**P-GRAFT Training:**

1603

1604

1605

Training and sampling using GRAFT is straightforward since standard training and inference scripts can be used out-of-the box: the only additional step need is rejection sampling on the generated samples before training. For P-GRAFT, the following changes are to be made:

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

- While sampling the training data, the intermediate latents should also be saved along with the final denoised iamge/latent. Rejection sampling is to be done on these intermediate latents, but using the rewards corresponding to the final denoised images.
- While training, note that training has to be done by noising the saved intermediate latents. This needs a re-calibration of the noise schedule, since by default, training assumes that we start from completely denoised samples. The easiest way to re-calibrate the noise schedule is by getting a new set of values for the betas parameter, `new_betas` as follows (where  $N_I$  denotes the intermediate step of P-GRAFT):

```
new_betas[0,  $N_I$ ] ← 0
new_betas[ $N_I$ ,  $N$ ] ← betas[ $N_I$ ,  $N$ ]
```

1618

1619

After re-calibrating the noise, we use `new_betas` to get the corresponding `new_alphas` and `new_alphas_cumprod`. It is also necessary to note that while training, the denoiser has been trained to predict  $X_0$  given any noised state  $X_t$  and not the saved intermediate

1620  
1621 Table 16: Training hyperparameters for T2I-CompBench++  
1622  
1623  
1624  
1625  
1626  
1627

Training Epochs	10
Image Resolution	$768 \times 768$
Batch Size	8
Learning Rate	$10^{-4}$
LR Schedule	Constant
LoRA Fine-Tuning	True

1628  
1629  
1630 latent  $X_{t_{N_I}}$ . Let the corresponding saved completely denoised latent be  $X_0$  To ensure that  
1631 the training is consistent, we train using the following strategy:  
1632

1633 Sample  $\epsilon \sim \mathcal{N}(0, \mathbb{I})$   
1634 Get  $X_t \leftarrow \left( \sqrt{\text{new\_alphas\_cumprod}[t]} \right) X_{t_{N_I}} + \left( \sqrt{1 - \text{new\_alphas\_cumprod}[t]} \right) \epsilon$   
1635 Get  $\epsilon' \leftarrow \frac{X_t - \sqrt{\text{alphas\_cumprod}[t]} X_0}{\sqrt{1 - \text{alphas\_cumprod}[t]}}$   
1636 Compute Loss using  $X_t$  and  $\epsilon'$   
1637

1638  
1639 E.3 POLICY GRADIENT ALGORITHMS  
1640

1641 DDPO(Black et al., 2023) is an on-policy policy gradient method for diffusion models that optimizes  
1642 a clipped importance-weighted objective over the denoising trajectory. The original paper reports  
1643 results on experiments using at most 400 prompts. Both prompt sets we consider are significantly  
1644 larger (1600 prompts for GenAI-Bench and 5600 (train) prompts for T2I-CompBench++). This  
1645 difference is crucial, since it has been shown in Deng et al. (2024) that scaling DDPO to large  
1646 prompt sets result in unstable training and subpar performance. We also observe this phenomenon,  
1647 as evidenced by the results in Table 2. As menioned in the main text, we also augment DDPO with  
1648 additional elements in an attempt to improve performance. In particular, we study the following  
1649 variants:  
1650

1. **DDPO:** Clipped importance-weighted policy gradient.
2. **DDPO+KL** DDPO augmented with a stepwise KL regularizer to the (frozen) reference  
model.
3. **DDPO+KL+EMA** DDPO with KL regularization as well as a prompt-wise exponential-  
moving-average baseline for advantage estimation.

1656 **Baseline Implementation:** We use the official PyTorch implementation of DDPO<sup>1</sup> - we further  
1657 adapt the codebase to implement other variants. Fine-tuning is always done on SDv2 using LoRA  
1658 on the UNet only with a LoRA rank of 16. For the results reported in Table 2, we retain the  
1659 hyperparameters used in Black et al. (2023). In particular, we use a PPO clip range of  $10^{-4}$ , gradient  
1660 clipping norm of 1.0, Adam optimizer with  $\beta_1 = 0.9, \beta_2 = 0.999$  and weight decay of  $10^{-4}$ .  
1661 Following the original paper, we train with a relatively high learning rate of  $3 \times 10^{-4}$  since LoRA  
1662 fine-tuning is used. We sample 32 prompts per epoch and train with a batch size of 8, leading to 4  
1663 training iterations per epoch. However, note that each training iteration requires gradients across  
1664 the whole denoising trajectory - this means that within each training iteration, 50 gradient calls are  
1665 needed, corresponding to 50 sampling steps. For GenAI-Bench, training is done for 500 such epochs,  
1666 whereas for T2I-Compbench++, training is done for 800 epochs. With this setup, in Tables 17 and 18,  
1667 we compare the sampling/compute requirements for DDPO and GRAFT/P-GRAFT. In particular,  
1668 note that GRAFT/P-GRAFT already outperforms DDPO with  $K = 1, M = 4$  despite DDPO being  
1669 trained on  $10 \times$  more samples and  $50 \times$  more gradient calls.  
1670

1671 **Additional configurations with base hyperparameters:** With the base hyperparameters described  
1672 above, we also try augmenting DDPO with KL and EMA as described above. The training curves are  
1673 given in Figure 3.

<sup>1</sup><https://github.com/kvab/black/ddpo-pytorch>

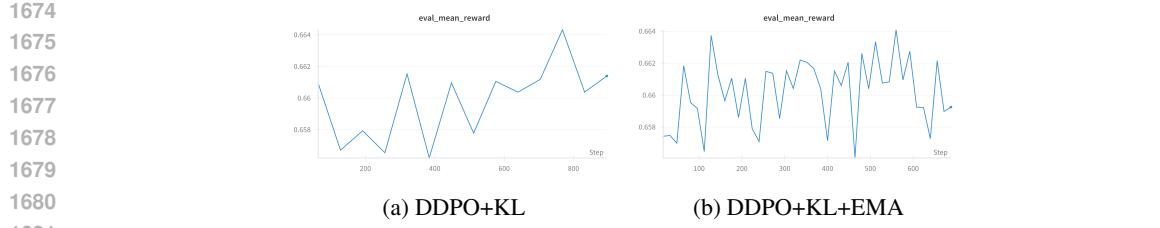


Figure 3: Training curves for the three policy-gradient baselines on GenAI Bench (1,600 prompts) with low value of clipping.

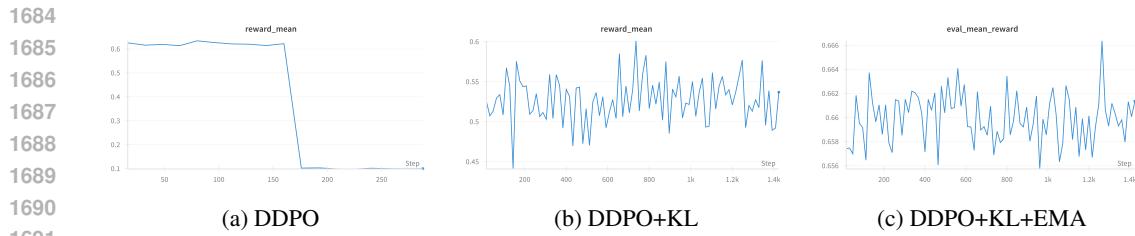


Figure 4: Training curves for the three policy-gradient baselines on GenAI Bench (1,600 prompts) with high value of clipping.

**DDPO:** We also try additional settings for hyperparameters apart from the ones we have reported so far. Sampling uses DDIM with  $T \in [40, 50]$  steps and classifier-free guidance  $g = 5$ . Optimization uses AdamW with learning rates  $\{2 \times 10^{-5}, 10^{-5}\}$ , batch sizes 8/8 (sampling/training), PPO-style clipping  $\epsilon \in \{0.1, 0.2\}$ . Following DDPO, we replay the scheduler to compute per-step log-probabilities on the same trajectories:  $\ell_t = \log p_\theta(x_{t-1} | x_t, c)$  and  $\ell_t^{\text{old}} = \log p_{\theta_0}(x_{t-1} | x_t, c)$ . We use the clipped objective:

$$\mathcal{L}_{\text{DDPO}} = -\mathbb{E}[\min(r_t A, \text{clip}(r_t, 1-\epsilon, 1+\epsilon) A)], \quad r_t = \exp(\ell_t - \ell_t^{\text{old}}), \quad (21)$$

with a centered batchwise advantage  $A$ . Specifically, we experiment with higher clipping range,  $\epsilon \in \{0.1, 0.2\}$  and use a whitened batchwise advantage.  $\ell_t, \ell_t^{\text{old}}$  are obtained by replaying the DDIM scheduler on the same trajectory.

**Result.** On 1600 prompts, the learning curve exhibits a short initial rise followed by a sharp collapse after  $\sim 150$  steps (Fig. 4). The setting of 1600 heterogeneous prompts induces high variance and many ratios  $r_t$  saturate at the clipping boundary, producing low-magnitude effective gradients and the observed drop in reward.

**DDPO+KL:** We augment equation 21 with a per-step quadratic penalty to the frozen reference:

$$\mathcal{L}_{\text{DDPO+KL}} = \mathcal{L}_{\text{DDPO}} + \beta \frac{1}{T} \sum_{t=1}^T (\ell_t - \ell_t^{\text{old}})^2, \quad \beta \in \{0.02, 0.005\}. \quad (22)$$

**Result.** The KL term prevents divergence of the policy and eliminates the reward collapse after the first few steps. Even with this, average reward improvements remain limited. Larger  $\beta$  contracts the policy towards the reference, whereas smaller  $\beta$  provides insufficient variance control, yielding small net gains.

**DDPO+KL+EMA (prompt-wise baseline):** To mitigate cross-prompt bias, we maintain for each prompt  $z$ , an EMA of reward and variance,

$$b(z) \leftarrow (1-\alpha)b(z) + \alpha r, \quad v(z) \leftarrow (1-\alpha)v(z) + \alpha(r - b(z))^2,$$

and employ a whitened advantage inside equation 21:  $\widehat{A} = \frac{r-b(z)}{\sqrt{v(z)+\epsilon}} + \eta$ ,  $\eta \sim \mathcal{N}(0, \sigma^2)$ .

**Result.** Training is the most stable among the three variants and exhibits smooth reward trajectories without collapse, yet the absolute improvement in mean reward is modest relative to the base policy.

**PRDP:** We also tried implementing PRDP (Deng et al., 2024) using the PRDP loss function provided in the appendix of the paper since no official code was provided. However, we did not see any significant improvement compared to the baseline despite following the algorithm and hyperparameters closely. One potential reason for this could be that we use LoRA fine-tuning whereas the original paper uses full fine-tuning. Further, we rely on gradient checkpointing for the implementation as well since the backpropogation is through the entire sampling trajectory.

Table 17: Comparison of Sampling Cost and Training Cost for GenAI-Bench

Algorithm	Samples generated	Samples Trained on	Gradient Calls
GRAFT( $K = 10, M = 100$ )	160k	16k	20k
GRAFT ( $K = 1, M = 4$ )	6.4k	1.6k	2k
DDPO	16k	16k	100k

Table 18: Comparison of Sampling Cost and Training Cost for T2I-CompBench++

Algorithm	Samples generated	Samples Trained on	Gradient Calls
GRAFT ( $K = 1, M = 4$ )	22.4k	5.6k	7k
DDPO	25.6k	25.6k	160k

#### E.4 COMPUTE FLOPS ANALYSIS OF P-GRAFT

We compare the compute cost of P-GRAFT and DDPO in terms of total UNet FLOPs. Let  $F_u$  denote the cost of one UNet forward pass at  $64 \times 64$  latent resolution. Following Kaplan et al. (2020), we approximate a backward training step 2 times of a forward step. So if  $F_u$  is the forward step compute, a forward + backward step will incur  $3F_u$ . We assume a batch size of 1 for both algorithms for standardization.

For  $P$  prompts,  $M$  samples per prompt, top- $K$  retained,  $T$  diffusion steps,  $E_{\text{sft}}$  epochs. For the implementation we use the standard stable diffusion training script that only samples a single timestep  $t \in [0, T]$  during training:

$$F_{\text{P-GRAFT}} = \underbrace{P M T}_{\text{sampling}} F_u + \underbrace{E_{\text{sft}} P K}_{\text{training}} 3F_u,$$

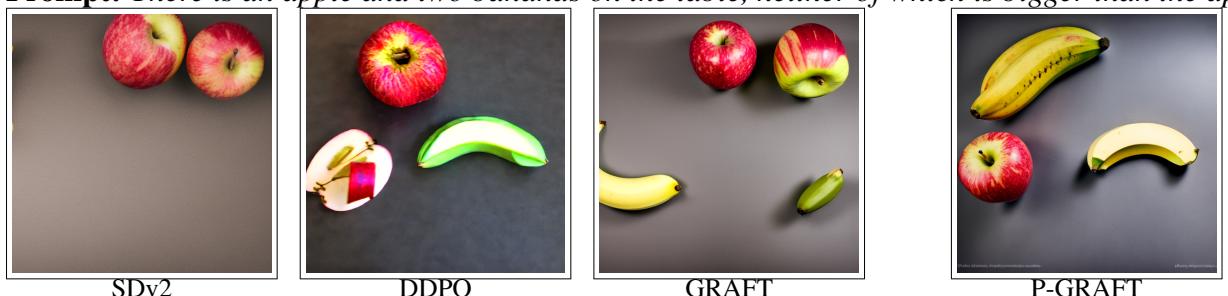
$$F_{\text{DDPO}} = \underbrace{E_{\text{ddpo}} N_{\text{gen}} T}_{\text{trajectories}} \cdot (1 + 3)F_u$$

**GenAI-Bench configuration.** We use  $P=1600$ ,  $M=100$ ,  $K=10$ ,  $T=40$ ,  $E_{\text{sft}}=10$  for P-GRAFT; Trajectories generated per epoch  $N_{\text{gen}}=128$ , and Number of Epochs  $E_{\text{ddpo}}=50$  for DDPO

Table 19: FLOPs in units of forward pass  $F_u$  for GenAI-Bench.

Algorithm	Sampling	Training	Total
P-GRAFT ( $K=10, M=100$ )	6.40M	0.48M	6.88M
P-GRAFT ( $K=1, M=4$ )	0.256M	0.048M	0.304M
DDPO ( $E=50, N_{\text{gen}}=128$ )	0.256M	0.768M	1.024M

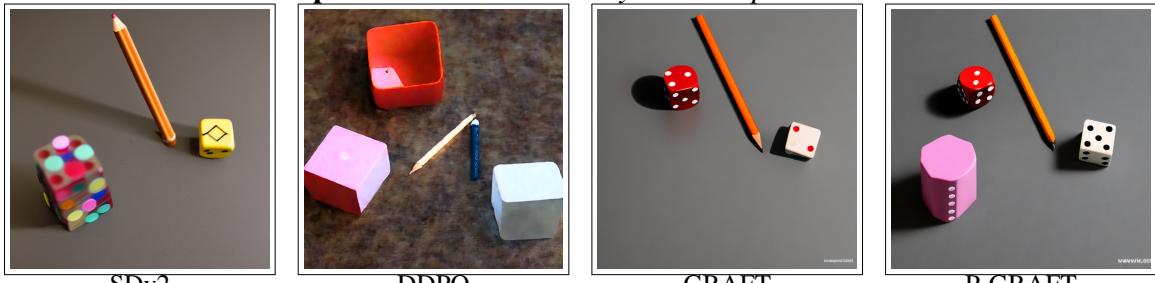
**Discussion.** P-GRAFT’s total compute is dominated by sample generation, while backpropagation is confined to fine-tuning on the selected top- $K$  samples. In contrast, DDPO backpropagates through all  $T$  denoising steps online for every sample, creating a sequential bottleneck. Consequently, despite DDPO’s nominal FLOPs appearing comparable or lower in our regime, its wall-clock time is substantially longer due to stepwise backward passes that are less parallelizable. Moreover, as shown in Table 2, P-GRAFT achieves higher rewards under the reported budgets; and in the compute-matched case ( $K=1$ ; Table 6), P-GRAFT still outperforms DDPO, indicating that gains come from improved optimization and not just additional training compute.

1782 E.5 QUALITATIVE EXAMPLES  
17831784 E.5.1 GENAI-BENCH  
17851786 **Prompt:** Three flowers in the meadow, with only the red rose blooming; the others are not open.  
17871787 **Prompt:** In the yoga room, all the mats are red.  
17881788 **Prompt:** Three policemen working together to direct traffic at a busy intersection.  
17891789 **Prompt:** There is an apple and two bananas on the table, neither of which is bigger than the apple.  
17901829 Figure 5: Qualitative examples on GENAI-BENCH. All results are reported for the same seed across  
1830 different algorithms.  
18311832  
1833  
1834  
1835

1836  
1837

## E.5.2 T2I-COMPBENCH++

1838

**Prompt: a cubic dice and a cylindrical pencil holder**

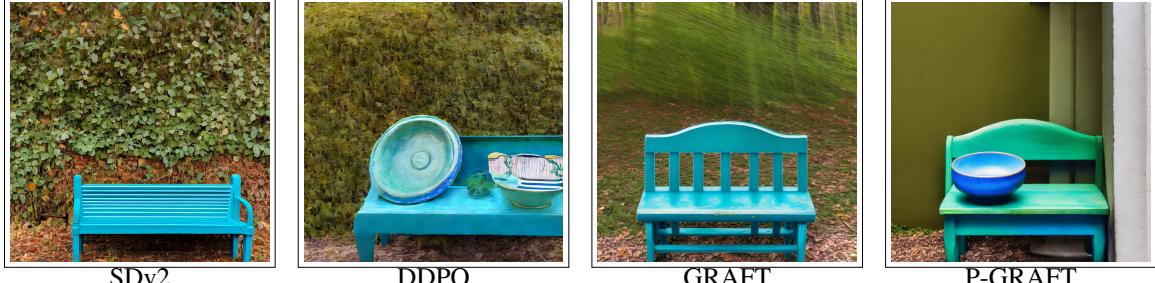
1848

1849

**Prompt: a bee on the bottom of a airplane**

1859

1860

**Prompt: a green bench and a blue bowl**

1870

1871

**Prompt: a green acorn and a brown leaf**

1881

1882

Figure 6: Qualitative examples on T2I-COMPBENCH. All results are reported for the same seed across different algorithms.

1883

1884

1885

1886

1887

1888

1889

1890  
1891

## F LAYOUT AND MOLECULE GENERATION

1892  
1893

### F.1 INTERLEAVED GIBBS DIFFUSION (IGD)

1894  
1895  
1896  
1897  
1898  
1899

Fine-tuning for both layout generation and molecule generation are done using models pre-trained using the Interleaved Gibbs Diffusion (IGD) (Anil et al., 2025) framework. IGD performs well for discrete-continuous generation tasks with strong constraints between variables - and hence is particularly useful for tasks like layout generation and molecule generation. Further, IGD offers a generalizable framework which can be used for both tasks - while other discrete-continuous diffusion frameworks exist, they are specialized to a particular task, often using domain specific adaptations.

1900  
1901  
1902  
1903  
1904

On a high-level, IGD interleaves diffusion for discrete and continuous elements using Gibbs-style noising and denoising. Essentially, discrete elements are noised using flipping and trained using a binary classification loss. Continuous elements use typical DDPM-style noising and training. While the exact forward and reverse processes are different from DDPM-style processes which we have considered in the main text, the key results follow empirically and theoretically.

1905  
1906

### F.2 LAYOUT GENERATION

1907  
1908  
1909  
1910  
1911  
1912  
1913

**Problem Formulation:** A layout is defined as a set of  $N$  elements  $\{e_i\}_{i=1}^N$ . Each element  $e_i$  is represented by a discrete category  $t_i \in \mathbb{N}$  and a continuous bounding box vector  $\mathbf{p}_i \in \mathbb{R}^4$ . Following (Anil et al., 2025), we use the parameterization  $\mathbf{p}_i = [x_i, y_i, l_i, w_i]^\top$ , where  $(x_i, y_i)$  represents the upper-left corner of the bounding box, and  $(l_i, w_i)$  its length and width, respectively. *Unconditional* generation represents generation with no explicit conditioning for the elements, whereas *Class-Conditional* generation indicates generations conditioned on element categories.

1914  
1915  
1916  
1917  
1918  
1919  
1920

**Implementation Details:** For pre-training, we follow the exact strategy used in (Anil et al., 2025). Fine-tuning is also done with the same hyperparameters used for pre-training. Since the data and model sizes are significantly smaller compared to images, each round of rejection sampling is done on 32768 samples, of which the top 50% samples are selected. For each sampling round, 10000 training iterations are performed with a training batch size of 4096. The results reported in Table 3 are for 20 such sampling rounds. FID computation is done by comparing against the test split of PubLayNet.

1921  
1922

### F.3 MOLECULE GENERATION

1923  
1924  
1925  
1926  
1927

**Problem Formulation:** The task of molecule generation involves synthesizing molecules given a dataset of molecules. A molecule consists of  $n$  atoms denoted by  $\{z_i, \mathbf{p}_i\}_{i=1}^n$ , where  $z_i \in \mathbb{N}$  is the atom's atomic number and  $\mathbf{p}_i \in \mathbb{R}^3$  is the position. A diffusion model is trained to generate such molecules. In this work, we take such a pre-trained model, and try to increase the fraction of stable molecules, as deemed by RDKit.

1928  
1929  
1930  
1931  
1932  
1933  
1934  
1935

**Implementation Details:** For pre-training, we follow the exact strategy used in (Anil et al., 2025). Fine-tuning is also done with the same hyperparameters used for pre-training. Since the data and model sizes are significantly smaller compared to images, each round of rejection sampling is done on 32768 samples. We select all stable molecules, but with the de-duplication strategy described in Section 3 - we find that *this is crucial* to maintain diversity of generated molecules. For each sampling round, 10000 training iterations are performed with a training batch size of 4096. The 1× in Table 4 corresponds to 10 such sampling rounds - 9× therefore corresponds to 90 sampling rounds.

1936  
1937  
1938  
1939  
1940

**Uniqueness of Generated Molecules:** To demonstrate that the fine-tuned models still generate diverse molecules, and do not collapse to generating a few stable molecules, we report the uniqueness metric computed across the generated molecules below. From Table 20, it is clear that the fine-tuned models still generate diverse samples since the uniqueness of the generated molecules remain close to the pre-trained model. Uniqueness is as determined by RDKit.

1941  
1942  
1943

**Effect of de-duplication** We also try out an ablation where we use GRAFT, but without the de-duplication - i.e., we train on all stable molecules irrespective of whether they are unique or not. The results are shown in Figure 7 - without de-duplication, it can be seen that though stability is recovered,

1944

Table 20: Uniqueness of generated molecules

1945

1946

Model	Mol: Stability	Uniqueness
Baseline	$90.50 \pm 0.15$	$95.60 \pm 0.10$
GRAFT	$90.76 \pm 0.20$	$96.04 \pm 0.46$
P-GRAFT( $0.5N$ )	$90.46 \pm 0.27$	$95.70 \pm 0.28$
P-GRAFT( $0.25N$ )	<b><math>92.61 \pm 0.13</math></b>	$95.32 \pm 0.07$

1947

1948

1949

1950

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

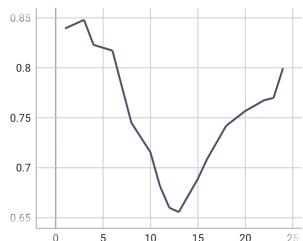
1962

1963

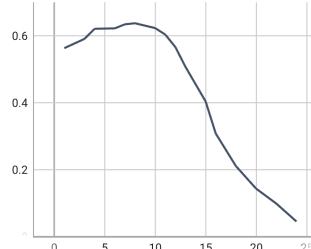
1964

1965

1966



(a) Molecule Stability



(b) Molecule Uniqueness

Figure 7: Molecule Stability and Uniqueness without De-duplication

uniqueness is lost, indicating that the model produces only a small subset of molecules it was initially able to produce.

**Fine-Tuning without Predictor-Corrector:** IGD makes use of a version of predictor-corrector method (Lezama et al., 2022; Zhao et al., 2024; Campbell et al., 2022; Gat et al., 2024) termed ReDeNoise at inference-time to further improve generations. The results reported so far make use of this predictor-corrector. While ReDeNoise improves performance significantly, it comes at the cost of higher inference-time compute. We report results of the baseline and fine-tuned version without ReDeNoise in Table 21. Both GRAFT and P-GRAFT still show improvement over the baseline, even without ReDeNoise.

Table 21: Results for Molecule Generation without ReDeNoise

Model	Mol: Stability	Sampling Steps
Baseline	84.00	-
GRAFT	87.13	$9 \times$
P-GRAFT ( $0.5N$ )	84.57	$1 \times$
P-GRAFT ( $0.25N$ )	<b>88.36</b>	$1 \times$

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997

1998  
1999

## G INVERSE NOISE CORRECTION

2000  
2001

2002 The pre-trained models, corresponding to TRAIN\_FLOW function in Algorithm 3, are trained using  
 2003 the NSCNpp architecture and hyperparameters from the official codebase of Song et al. (2020b) with  
 2004 minor changes which we describe below. The noise corrector model is also trained with the same  
 2005 architecture except that the number of channels are reduced from the original 128 to 64 channels  
 2006 - this leads to a reduction in parameter count by  $\approx 4\times$ . For the pre-trained model, we train with  
 2007 num\_scales = 2000, positional embeddings and a batch size of 128. For the noise corrector model,  
 2008 we use the same hyperparameters except for num\_scales = 1000. FID with 50000 samples is  
 2009 used to measure the performance, as is standard in the literature. Note that a separate noise corrector  
 2010 model is trained for each choice of  $\eta$  in Algorithm 3, i.e., for the results reported in Table 5, separate  
 2011 noise corrector models are trained for pre-trained steps of 100 and 200.

2012  
2013  
2014

**CelebA-HQ:** For the baseline pre-trained flow model, we use the checkpoint after 330k iterations,  
 since this gave the lowest FID. For noise corrector model training, we use this checkpoint to generate  
 the inverted noise dataset and train on it for 150k iterations.

2015  
2016  
2017  
2018  
2019  
2020

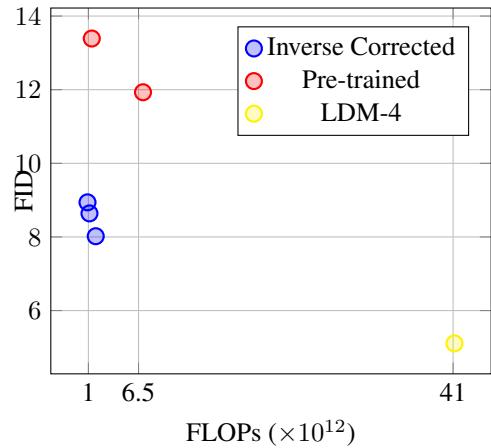
**LSUN-Church:** For the baseline pre-trained flow model, we use the checkpoint after 350k iterations,  
 since this gave the lowest FID. For noise corrector model training, we use this checkpoint to generate  
 the inverted noise dataset and train on it for 55k iterations. Note that Backward Euler (Algorithm 6)  
 suffered from numerical instability, which we hypothesize is due to plain backgrounds, when done on  
 LSUN-Church. To alleviate this issue, we perturb the images with a small Gaussian noise  $\mathcal{N}(0, \sigma^2 \mathbf{I})$ ,  
 with  $\sigma = 10^{-3}$ .

2021  
2022

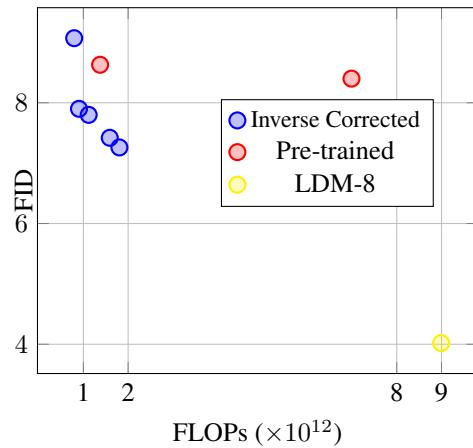
## G.2 FLOPs COMPARISON

2023  
2024

We present a comparison of the exact FLOPs used for inference:



(a) Celeb HQ



(b) LSUN Church

Figure 8: **FLOPs vs FID:** The inverse corrected model achieves better FID despite incurring lower FLOPs. Corresponding LDM models have been added for both datasets for reference.

2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051