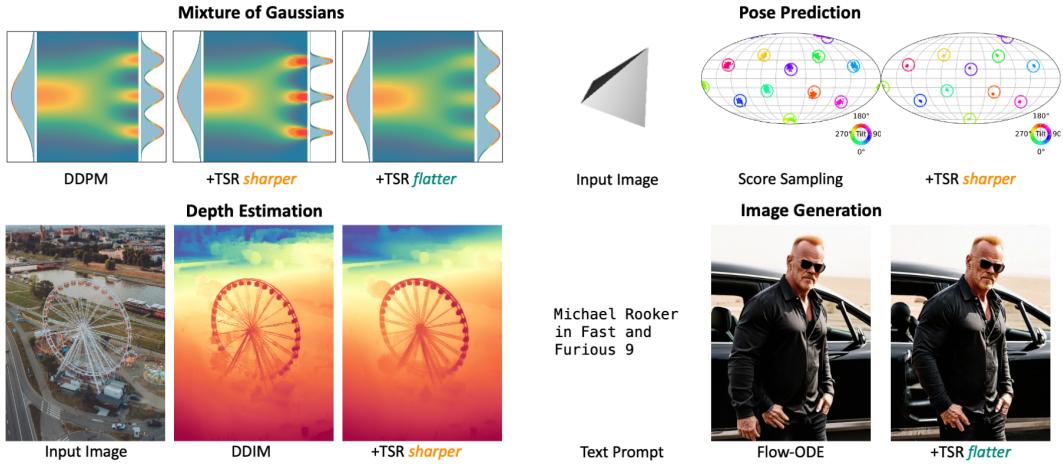


# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 TEMPORAL SCORE RESCALING FOR TEMPERATURE SAMPLING IN DIFFUSION AND FLOW MODELS

005 **Anonymous authors**

006 Paper under double-blind review



025 **Figure 1: Temporal Score Rescaling (TSR)** provides a mechanism to steer the sampling diversity  
026 of diffusion and flow models at inference. *Top-left*: Probability density evolution when sampling a  
027 1D Gaussian mixture with DDPM, and the effects of TSR, which can control the sampling process  
028 to yield sharper or flatter distributions. *Top-right, bottom*: TSR can be applied to any pre-trained  
029 diffusion or flow model, improving performance across diverse domains such as pose prediction,  
030 depth estimation, and image generation.

## 031 ABSTRACT

033 We present a mechanism to steer the sampling diversity of denoising diffusion  
034 and flow matching models, allowing users to sample from a sharper or broader  
035 distribution than the training distribution. We build on the observation that these  
036 models leverage (learned) score functions of noisy data distributions for sam-  
037 pling and show that rescaling these allows one to effectively control a ‘local’  
038 sampling temperature. Notably, this approach does not require any finetun-  
039 ing or alterations to training strategy, and can be applied to any off-the-shelf  
040 model and is compatible with both deterministic and stochastic samplers. We  
041 first validate our framework on toy 2D data, and then demonstrate its applica-  
042 tion for diffusion models trained across five disparate tasks – image generation,  
043 pose estimation, depth prediction, robot manipulation, and protein design. We  
044 find that across these tasks, our approach allows sampling from sharper (or flatter)  
045 distributions, yielding performance gains *e.g.*, depth prediction models bene-  
046 fit from sampling more likely depth estimates, whereas image generation mod-  
047 els perform better when sampling a slightly flatter distribution. Project page:  
048 <https://temporalscorerescaling-anonymous.github.io/>

## 049 1 INTRODUCTION

052 Score-based generative models, such as denoising diffusion (Ho et al., 2020) and flow matching  
053 (Lipman et al., 2023; Liu et al., 2023b), have become ubiquitous across AI applications. Given  
training data  $\{\mathbf{x}^n\}$ , they can model the underlying data distribution  $p(\mathbf{x})$  (or  $p(\mathbf{x}|\mathbf{c})$  for conditional

054 settings) and at inference, they allow drawing samples  $\mathbf{x} \sim p(\mathbf{x})$  e.g., to generate novel images.  
055 However, in certain applications, we may not want to truly sample the original data distribution.  
056 For example, when predicting depth from RGB input, we may want the more likely estimate(s) as  
057 output. In contrast, an artist exploring design choices may want the trained image generative model  
058 to yield more diverse samples, even if they may be somewhat less likely in the data. In this work,  
059 we ask whether we can steer the sampling process of diffusion or flow matching models to output  
060 more likely (or conversely, more diverse) samples than the original training data.

061 This process of trading off sample likelihood and diversity at inference is commonly referred to as  
062 *temperature sampling* (Hinton et al., 2015) – a higher temperature leads to diverse samples, and a  
063 lower temperature leads to more likely ones. While prior methods have investigated temperature  
064 sampling for score-based generative models like denoising diffusion, developing an efficient tem-  
065 perature sampling method for pre-trained diffusion/flow models remains an open challenge. For  
066 example, commonly leveraged techniques like classifier-free guidance (Ho & Salimans, 2022) or  
067 variance-reduced sampling (Yim et al., 2023; Geffner et al., 2025) can trade off sampling diversity  
068 and likelihood, but as we show later, these are not probabilistically interpretable as temperature scal-  
069 ing the data distribution. Conversely, methods such as likelihood-weighted finetuning (Shih et al.,  
070 2023) or Langevin correction (Song et al., 2021b; Du et al., 2023) can indeed allow temperature  
071 sampling, but at the cost of additional training or significantly increased inference-time computa-  
072 tion. In this work, we instead seek to develop a (local) temperature sampling method that is: a)  
073 *training free* i.e., does not require fine-tuning or distilling a pre-trained model, b) compatible with  
074 deterministic samplers e.g., DDIM (Song et al., 2021a), c) efficient i.e., does not increase the number  
075 of score evaluations at inference, and d) provably correct for some simple distributions.

076 Towards developing such an approach, we note that denoising diffusion and flow matching mod-  
077 els define a forward process to induce noisy data distributions  $p(\mathbf{x}_t)$  and train neural networks to  
078 approximate the corresponding score functions  $\nabla \log p(\mathbf{x}_t)$ . We ask whether one can analytically  
079 relate these to the score of the (hypothetical) distributions  $\bar{p}(\mathbf{x}_t)$  that would be induced by the for-  
080 ward process if the original data distribution were temperature scaled. We study the case of mixture  
081 of isotropic Gaussians, and derive a simple (time-dependent) rescaling function. As the reverse  
082 sampling process for sampling flow/diffusion models relies only on the learned score functions, our  
083 derived rescaling thus allows a training-free approach by simply scaling the inferred score at each  
084 inference step. While the analytical derivation is restricted to a simple setting, we show that our  
085 approach can be generally interpreted as a ‘local’ temperature sampling method, where it does not  
086 alter the overall distribution of global modes, but controls the local variance of samples around it.

087 We perform experiments to highlight the broad applicability of TSR. We show that it can efficiently  
088 allow local temperature sampling for denoising diffusion and flow matching models and is compati-  
089 ble with generic stochastic and deterministic samplers. We study diverse applications like image  
090 generation, depth estimation, pose prediction, robot manipulation, and protein generation. Across  
091 these applications, we show that TSR can provide a plug-and-play solution to control the sampling  
092 diversity of pre-trained models and leads to consistent performance gains e.g., allowing more precise  
093 depth and pose inference, or enabling image generation to better match real data distribution.

## 094 2 PRIOR ART

095 **Guided Inference.** A widely adopted mechanism for steering sampling in diffusion and flow mod-  
096 els is to leverage Classifier-Free Guidance (CFG) (Ho & Salimans, 2022). While this allows one  
097 to trade off likelihood and diversity by controlling the effect of the conditioning on the drawn sam-  
098 ples, it is fundamentally different from temperature scaling. Moreover, CFG cannot be applied to  
099 unconditional models and even for conditional ones, requires training with condition dropout. An al-  
100 ternative to CFG by Karras et al. (2024) is to use a ‘bad version’ of the diffusion model for guidance,  
101 but its probabilistic interpretation is unclear and it also requires intermediate checkpoints which are  
102 not widely available even for open-weight models. In comparison, TSR serves as a plug-and-play  
103 technique compatible with any diffusion and flow matching model without any requirement on train-  
104 ing. Moreover, as we empirically demonstrate for image generation, our method is orthogonal to  
105 CFG and can be applied together for further improvement in quality.

106 **Temperature Scaling in Diffusion Models.** We are not the first to consider temperature sampling  
107 in context of diffusion models. In particular, Shih et al. (2023) presented a technique to finetune

108 diffusion (and autoregressive) models for temperature scaled inference. Their approach assigned an  
 109 importance weight to each training sample based on its likelihood approximated by computing its  
 110 ELBO with respect to a pretrained diffusion model on the same data. However, this approach is  
 111 not training-free, making it difficult to leverage for large models and impossible in scenarios where  
 112 training data is unavailable. An alternative training-free approach is to modify the reverse sampling  
 113 by applying a stochastic MCMC corrector at each denoising step (Song et al., 2021b; Du et al.,  
 114 2023). However, this increases the computational cost at inference by an order of magnitude and  
 115 does not support deterministic sampling. In contrast `TSRis` a training-free approach that does not  
 116 increase the inference cost and can be leveraged for stochastic and deterministic sampling.  
 117

118 **Pesudo-temperature Sampling via Noise Scaling.** Perhaps the closest to our approach in terms of  
 119 being efficient and training-free is the technique of ‘Constant Noise Scaling’ (CNS) where one scales  
 120 the stochastic noise at each sampling step by a constant. More formally, following the definition by  
 121 Song et al. (2021b), CNS can be viewed as sampling the following reverse SDE:  
 122

$$d\mathbf{x} = [f(\mathbf{x}, t) - g(t)^2 \nabla \log p_t(\mathbf{x})] dt + \frac{g(t)}{\sqrt{k}} d\bar{\mathbf{w}} \quad (1)$$

123 where  $f(\mathbf{x}, t)$ ,  $g(t)$  denote the drift and diffusion coefficient, and  $d\bar{\mathbf{w}}$  is a standard Wiener process.  
 124 Compared to regular reverse diffusion SDE, the noise term is scaled by a constant  $1/\sqrt{k}$ . While CNS  
 125 is the de facto approach to control sample variance in several domains (Yim et al., 2023; Geffner  
 126 et al., 2025), as Shih et al. (2023) point out, it is only a ‘*pseudo temperature*’ sampling method. Intu-  
 127 itively, the noise-to-score ratio controls the strength of exploration versus converging to distribution  
 128 modes during sampling. By scaling down this ratio by a constant, CNS over-suppresses exploration  
 129 at high noise levels and under-suppresses it at low noise levels, leading to inadequate exploration of  
 130 the data space when the model should recover global structure. We empirically show in Section 4  
 131 that CNS behaves differently from temperature scaling and drop modes even for simple distribu-  
 132 tions. Moreover, CNS only applies to stochastic samplers and struggles with modern flow-matching  
 133 models (see Section 5.1). In contrast, we propose a time-dependent score scaling schedule that pre-  
 134 serves the global structure of the sampled distribution and is compatible with both deterministic and  
 135 stochastic samplers.  
 136

### 3 FORMULATION

#### 3.1 PRELIMINARIES

141 Both diffusion and flow matching models fall under the family of stochastic interpolants (Albergo  
 142 et al., 2023), which convert samples from data distribution  $\mathbf{x}_0 \sim p_0(\mathbf{x})$  to gaussian noise  $\epsilon \sim$   
 143  $\mathcal{N}(0, \mathbf{I})$ . The interpolant process can be defined as:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \epsilon \quad (2)$$

144 Different noise schedules  $\alpha_t, \sigma_t$  correspond to different formulations of stochastic interpolants. For  
 145 example, for flow matching models, it is common to set  $\alpha_t = 1 - t$ ,  $\sigma_t = t$ , while for variance-  
 146 preserving diffusion models, they are defined such that  $\alpha_t^2 + \sigma_t^2 = 1$ .  
 147

148 We can sample from the data distribution by training a model  $\mathbf{s}_\theta(\mathbf{x}, t) = \nabla \log p_t(\mathbf{x})$  that estimates  
 149 the score of the noisy distribution. Starting from  $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ , the sampling process usually solves  
 150 either a reverse-time SDE or a probability flow ODE. In practice, the learned model could pre-  
 151 dict various equivalent parameterization of the score, such as noise  $\epsilon_\theta(\mathbf{x}_t, t)$  (common in denoising  
 152 diffusion) or the probability flow velocity  $\mathbf{v}_\theta(\mathbf{x}_t, t)$  (common in flow matching), which can all be  
 153 expressed as linear combinations of score and  $\mathbf{x}_t$  (See Section 3.3).  
 154

#### 3.2 TEMPORAL SCORE RESCALING

155 Given a pre-training score function  $\mathbf{s}_\theta$ , we are interested in designing a temperature sampling process  
 156 that does not require training or additional computation at inference. In particular, we propose a  
 157 mechanism that achieves *local temperature scaling*, which can steer the variance of the sampled  
 158 distribution while preserving the global distribution structure (e.g., without mode dropping). More  
 159 formally, we define local temperature scaling as the task that takes in a data distribution  $p_0(\mathbf{x})$   
 160 modeled as a mixture of (*an unknown set of*) Gaussians and generates the corresponding ‘sharper’  
 161

162 or ‘flatter’ distributions  $\tilde{p}_0^k(\mathbf{x})$  (parameterized by  $k$ ):  
163

$$164 p_0(\mathbf{x}) \equiv \sum_m w_m \mathcal{N}(\mathbf{x}; \mu_m, \Sigma_m) \Rightarrow \tilde{p}_0^k(\mathbf{x}) \equiv \sum_m w_m \mathcal{N}(\mathbf{x}; \mu_m, \frac{1}{k} \Sigma_m)$$

165 Intuitively,  $\tilde{p}_0^k(\mathbf{x})$  represents a distribution where the variance near each local mode in the data  
166 distribution is scaled by  $\frac{1}{k}$ , while preserving all the means and weights. Such a local scaling effect  
167 is different from the traditional temperature scaling that would change the weights of modes and  
168 alter the distribution structure. We now formulate our problem statement as: How can we alter the  
169 pretrained score function  $\mathbf{s}_\theta$  so that a diffusion or flow sampler yields  $\tilde{p}_0^k(\mathbf{x})$  instead of  $p_0(\mathbf{x})$ ?  
170

171 **Isotropic Gaussian Data.** To instantiate this, we start with a simple scenario where the data are  
172 drawn from a single isotropic Gaussian distribution  $\mathbf{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I})$ . The target is to sample  
173 from the locally scaled distribution  $\tilde{p}_0^k(\mathbf{x}) \equiv \mathcal{N}(\boldsymbol{\mu}, \frac{\sigma^2}{k} \mathbf{I})$ . Under the stochastic interpolant process  
174 (Eq. 2), we define  $p_t(\mathbf{x})$ ,  $\tilde{p}_t^k(\mathbf{x})$  as the noisy distributions at time  $t$  for the original and scaled data  
175 distribution, respectively. Since both the original and scaled data distributions are Gaussian, their  
176 corresponding noisy distribution can also be shown to be Gaussian:  
177

$$179 p_t(\mathbf{x}) = \mathcal{N}(\alpha_t \boldsymbol{\mu}, (\alpha_t^2 \sigma^2 + \sigma_t^2) \mathbf{I}), \quad \tilde{p}_t^k(\mathbf{x}) = \mathcal{N}(\alpha_t \boldsymbol{\mu}, (\alpha_t^2 \frac{\sigma^2}{k} + \sigma_t^2) \mathbf{I}) \quad (3)$$

181 Then, we can derive the corresponding score functions for the above distributions:  
182

$$183 \nabla \log p_t(\mathbf{x}) = -\frac{\mathbf{x} - \alpha_t \boldsymbol{\mu}}{\alpha_t^2 \sigma^2 + \sigma_t^2}, \quad \nabla \log \tilde{p}_t^k(\mathbf{x}) = -\frac{\mathbf{x} - \alpha_t \boldsymbol{\mu}}{\alpha_t^2 \frac{\sigma^2}{k} + \sigma_t^2} \quad (4)$$

185 Comparing the two score functions above, we observe that the score for the scaled distribution and  
186 the score for the original distribution follow a time-dependent linear relationship:  
187

$$188 \nabla \log \tilde{p}_t^k(\mathbf{x}) = \frac{\eta_t \sigma^2 + 1}{\eta_t \frac{\sigma^2}{k} + 1} \nabla \log p_t(\mathbf{x}) \quad (5)$$

191 where  $\eta_t = \alpha_t^2 / \sigma_t^2$  is the signal-to-noise ratio. Note that  $k = 1.0$  recovers the original score. Given  
192 a score estimator  $\mathbf{s}_\theta(\mathbf{x}, t) = \nabla \log p_t(\mathbf{x})$ , we can compute the score of  $\tilde{p}_t^k$  with the above score  
193 rescaling equation and thus sample from  $\tilde{p}_0^k$  from the same sampling process.

194 **Mixture of Gaussians.** We can show that the score ratio relationship (Eq. 5) is also a valid approx-  
195 imation if the data distribution is a mixture of *well-separated* isotropic Gaussians. In Section B, we  
196 prove that the expected error between the score computed by Eq. 5 and the real score are bounded  
197 at all timestep  $t$ . On the high level, we derive an exponential bound for small  $t$ , where the modes  
198 are well-separated and only one Gaussian component dominates. For large  $t$ , we derive a poly-  
199 nomial bound based on the intuition that the distributions are similar to pure noise  $\mathcal{N}(0, \mathbf{I})$ . The error  
200 vanishes at both ends when  $t$  converges to 0 or 1. The maximum error at any intermediate  $t$  also  
201 converges to zero as the modes become more separated. We also empirically verify these results in  
202 Section A.5

### 203 3.3 STEERING INFERENCE IN DIFFUSION AND FLOW MATCHING

205 While the above analytical derivation for a score rescaling function focused on simple distributions,  
206 we empirically find that it can be applied across generic distributions and we operationalize Eq. 5 to  
207 define TSR sampling, a simple algorithm for steering sampling in diffusion and flow models:  
208

#### 209 **Sampling with Temporal Score Rescaling** TSR ( $k, \sigma$ )

210 Given a pre-trained score model  $s_\theta$ , TSR sampling substitutes its score prediction with:

$$212 \tilde{s}_\theta(\mathbf{x}, t) = r_t(k, \sigma) s_\theta(\mathbf{x}, t), \quad r_t(k, \sigma) := \frac{\eta_t \sigma^2 + 1}{\eta_t \frac{\sigma^2}{k} + 1} \quad (6)$$

214 where  $k, \sigma$  are user-defined parameters, and  $\eta_t$  is the signal-to-noise ratio of the forward process.  
215

216 This makes TSR a plug-and-play method compatible with any parameterization of  $s_\theta$  and sampling  
217 algorithm, since conversions between score and model predictions are always linear and invertible.  
218

219 **Denoising Diffusion:** These models are typically instantiated via neural networks  $\epsilon_\theta$  that learn to  
220 predict the noise added. We can infer the predicted score from this noise via a simple linear relation  
221  $s_\theta(\mathbf{x}, t) = -\sigma_t^{-1}\epsilon_\theta(\mathbf{x}, t)$ . We can thus perform TSR sampling in denoising diffusion models by  
222 simply using a rescaled noise prediction  $\tilde{\epsilon}_\theta(\mathbf{x}, t)$  in any diffusion sampler (e.g., DDPM, DDIM):

$$\tilde{\epsilon}_\theta(\mathbf{x}, t) = r_t(k, \sigma)\epsilon_\theta(\mathbf{x}, t) \quad (7)$$

225 **Flow Matching:** For flow matching models predicting the probability flow velocity  $v_\theta(\mathbf{x}, t)$ , the  
226 corresponding score function can be computed by (Ma et al., 2024):  
227

$$\mathbf{s}_\theta(\mathbf{x}, t) = -\frac{\alpha_t v_\theta(\mathbf{x}, t) - \dot{\alpha}_t \mathbf{x}}{\sigma_t(\dot{\alpha}_t \sigma_t - \alpha_t \dot{\sigma}_t)} \quad (8)$$

230 Combining Eq. 6 and Eq. 8, we can derive the corresponding flow velocity  $\tilde{v}_\theta$  for the scaled distribution,  
231 such that  $\tilde{\mathbf{s}}_\theta$  is a proper scaled version of the original score:  
232

$$\tilde{v}_\theta(\mathbf{x}, t) = \alpha_t^{-1}(r_t(k, \sigma)(\alpha_t v_\theta(\mathbf{x}, t) - \dot{\alpha}_t \mathbf{x}) + \dot{\alpha}_t \mathbf{x}) \quad (9)$$

234 Applying this scaled velocity  $\tilde{v}_\theta$  in the flow samplers yields desired samples from the scaled distribution.  
235 Similar conversion can also be derived for other parameterizations of diffusion models like  
236  $x_0$ -prediction and  $v$ -prediction.  
237

## 238 4 ANALYSIS

240 To understand the behavior of TSR, we first empirically validate it on toy data and show it is more  
241 effective in scaling the variance of samples while preserving each local mode compared to existing  
242 approaches. Then, we analyze how the input parameters  $(k, \sigma)$  control TSR and interpret their  
243 meanings in general settings.  
244

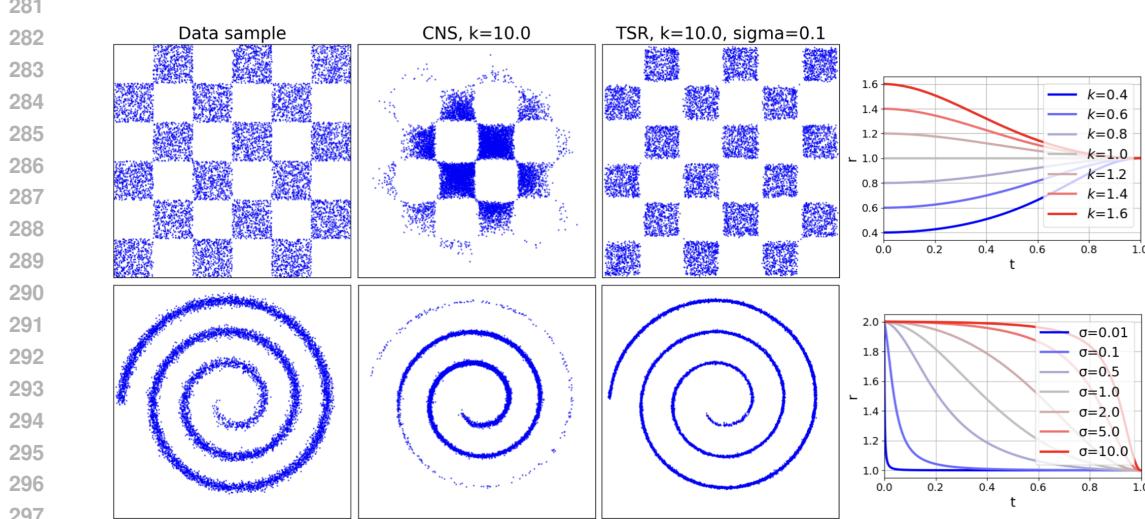
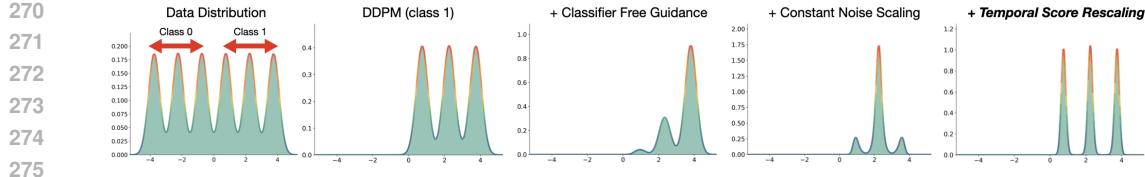
### 245 4.1 VALIDATION ON TOY DISTRIBUTIONS

246 **Mixture of 1D Gaussians.** We begin with a simple conditional generation task using a uniform  
247 mixture of 1D isotropic Gaussians in figure 2, where the left three and right three modes correspond  
248 to two different classes. We apply classifier-free guidance (CFG) with guidance scale 10, constant  
249 noise scaling (CNS) and TSR with  $k = 10$  individually to scale the conditional distribution and evaluate  
250 whether each method preserves all modes under scaling. As shown in figure 2, CFG produces  
251 imbalanced samples, often favoring outer modes, while CNS shifts mass toward central modes at  
252 the expense of others. By contrast, TSR samples evenly across all modes while reducing intra-mode  
253 variance, demonstrating that it preserves the multimodal structure even under conditioning.  
254

255 **General 2D Distributions.** We also apply TSR to unconditional generation on two complex 2D  
256 distributions: checkerboard and swiss roll. We train a small-scale diffusion model for each distribution  
257 and compare the scaled distribution sampled by CNS and TSR in figure 3. We observe that  
258 CNS consistently biases samples toward the central modes, resulting in mode collapse and poor  
259 coverage of peripheral regions. This supports the intuition that reducing noise too aggressively  
260 restricts exploration during the sampling process. In contrast, TSR maintains coverage of the global  
261 distribution while reducing local variance around each mode, producing samples aligning with the  
262 true distribution. These results show that, although derived for isotropic Gaussian data, TSR generalizes  
263 to more complex scenarios and provides consistent improvements in both conditional and  
264 unconditional generation.  
265

### 266 4.2 INTERPRETING RESCALING HYPERPARAMETERS

267 In the derivation of TSR,  $k$  referred to the factor of variance reduction and  $\sigma$  referred to the variance  
268 of the modes in **data distribution**. However, in real-world scenarios with more complex distributions,  
269 the variance of the data distribution is unknown. We provide an intuitive explanation of the role of  
270  $k$  and  $\sigma$  on the rescaling factor  $r_t$  to democratize the practical use of TSR in various scenarios.  
271 Specifically, we show how the rescaling factor  $r_t$  changes over sampling time with different  $k$  and  $\sigma$   
272



305  
306  
307 values in Fig. 3. Intuitively,  $k$  indicates the max/min of the rescaling factor  $r_t$ . As  $t \rightarrow 0$ , signal-to-  
308 noise ratio  $\eta_t \rightarrow \infty$ , and  $r_t \rightarrow k$ . Meanwhile,  $\sigma$  indicates how early we want to steer the sampling  
309 process. The larger  $\sigma$ , the earlier the sampling is steered. A very small  $\sigma$  lets us use the original  
310 diffusion sampling ( $r_t \approx 1.0$ ) and only steer the last few denoising steps.

## 311 5 APPLICATIONS

312  
313  
314 We demonstrate the broad applicability and effectiveness of TSR by applying it to a diverse set  
315 of real-world applications, spanning image generation (Section 5.1), protein design (Section 5.2),  
316 depth estimation (Section 5.3), pose prediction (Section 5.4), and robot manipulation (Section 5.5).  
317 For image generation, we find that a smaller  $k$  enhances details and improves performance, while  
318 for other tasks, a larger  $k$  yields higher accuracy of model predictions.

### 319 5.1 TEXT-TO-IMAGE GENERATION

320  
321 We examine the effect of steering the sampling distribution for diversity versus likelihood with  
322 TSR on Stable Diffusion 3 (Esser et al., 2024), a leading flow matching text-to-image model. As a  
323 creative task, image generation benefits from sampling a flatter distribution, which helps to recover  
more pleasing images with more high frequency details. We evaluate FID (Heusel et al., 2017;



Figure 4: **Qualitative Examples for Varying  $k$ .** TSR allows for tuning the generated outputs to be more diverse and detailed (lower  $k$ ) or more smooth and likely (higher  $k$ ). While neither extreme is desirable, we notice a  $k$  slightly smaller than 1 gives pleasing images with enhanced details.

	SD3		SD2		Flux.1 dev	
	FID $\downarrow$	CLIP $\uparrow$	FID $\downarrow$	CLIP $\uparrow$	FID $\downarrow$	CLIP $\uparrow$
Default Scheduler	24.77	32.82	22.81	33.66	53.99	31.97
+ TSR	<b>22.81</b>	<b>33.05</b>	<b>19.75</b>	<b>33.75</b>	<b>51.79</b>	<b>32.14</b>

Table 1: **Evaluation of Text-to-Image Generation across Models.** TSR consistently improve image quality across Stable Diffusion 3 Esser et al. (2024), Stable Diffusion 2 Rombach et al. (2022), and Flux.1 dev Labs (2024). The optimal  $(k, \sigma)$  found on SD3 generalize effectively to other models. For SD3 and Flux.1 dev, the default scheduler is Euler-ODE. For SD2 the default scheduler is DDPM.

Parmar et al., 2022) and CLIP (Radford et al., 2021) scores against a 5k image subset from LAION Aesthetics (Schuhmann et al., 2022) across different CFG guidance scale  $w_{\text{cfg}}$ , TSR parameter  $k$  and  $\sigma$ . We fix the number of sampling steps to 30. In figure 5, we see adjusting  $w_{\text{cfg}}$  makes a trade-off between text-alignment and image fidelity—higher  $w_{\text{cfg}}$  increases CLIP score at the cost of worse FID. Meanwhile, TSR allows for additional improvement beyond the Pareto frontier of CFG. Compared to the regular Euler ODE sampling, **TSR reduces FID score from 24.77 ( $\pm 0.10$ ) to 22.81 ( $\pm 0.13$ ) and increases CLIP score from 32.82 ( $\pm 0.014$ ) to 33.05 ( $\pm 0.018$ )**. These results are averaged over 5 random seeds. TSR achieves the optimal performance with  $k = 0.93, \sigma = 3.0$ . To verify the transferability of these parameters, we apply the same  $(k, \sigma)$  to Flux.1 dev (Labs, 2024) and Stable Diffusion 2 (Rombach et al., 2022) and report the results in Table 1. The optimal  $(k, \sigma)$  found on SD3 consistently improve performance on other models as well, suggesting the robustness of the choice for  $(k, \sigma)$  across models.

Notably, while it is possible to perform stochastic sampling with flow models like SD3, we found that it performs significantly worse than ODE sampling with the same compute budget (see Section A.1), making CNS impractical. We also show in Section A.1 that TSR achieves superior performance with denoising diffusion model (SD2, Rombach et al. (2022)) compared to CNS and other common samplers. Qualitatively, we observe in Fig. 4 that lower  $k$  leads to images with more high-frequency detail (in the extreme case more noise), and higher  $k$  leads to smoother images. We infer that using a smaller  $k$  flattens the modeled distribution and allows better coverage of the desirable image space. Overall, our results highlight that the control over the likelihood-diversity trade-off enabled by TSR is beneficial in image generation.

## 5.2 PROTEIN GENERATION

Generative models have emerged as a powerful paradigm in AI for Science. For example, Protein discovery (Abramson et al., 2024; Jumper et al., 2021; Wu et al., 2024) is an application where

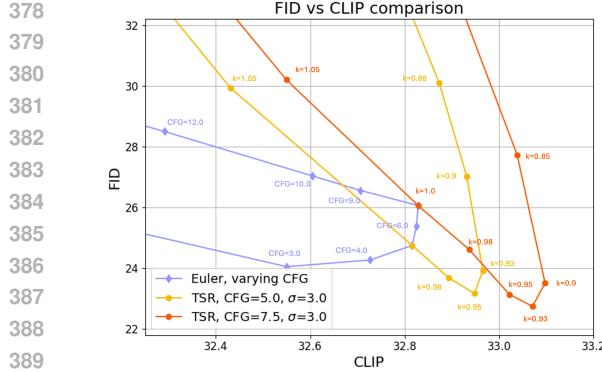


Figure 5: **Image Generation.** TSR achieves better text-alignment (CLIP) and image fidelity (FID), improving upon the Pareto frontier of CFG, which trades off between FID and CLIP.

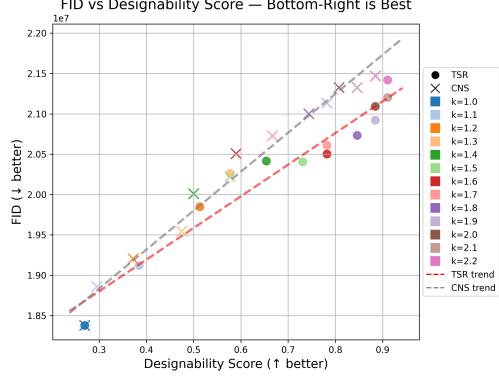


Figure 6: **Protein Generation.** TSR improves the designability score while preserving the diversity (FID) better compared to CNS. The original sampling has a designability score of 0.22.

such models have seen widespread adoption. However, not all generated proteins are valid in the real world. Thus, improving the designability of generated proteins is an important goal. CNS has previously been used to enhance sampling quality (Yim et al., 2023; Geffner et al., 2025).

We conduct experiments with FoldingDiff (Wu et al., 2024), a diffusion-based protein generation method, and compare TSR with CNS. Evaluation uses two complementary metrics: designability score (Wu et al., 2024), measuring structural quality and real-world feasibility, and protein FID (Faltings et al., 2025), capturing distributional similarity and thus diversity. Ideally, a method should achieve a high designability score and a low FID. As shown in fig. 6, samples from TSR lie in the bottom right regions, which shows TSR maintains protein diversity better than CNS, while improving the designability.

### 5.3 DEPTH ESTIMATION

The task of monocular depth estimation is inherently challenging due to its uncertainty—an object may appear large but distant, or small but close. Recent methods (Duan et al., 2024; Saxena et al., 2023; Ke et al., 2024) address this with diffusion models, where different samples correspond to plausible variations or interpolations of the underlying depth structure. We adopt Marigold (Ke et al., 2024), which fine-tunes a pre-trained text-to-image diffusion model for depth estimation and achieves strong results. However, individual samples can be suboptimal due to both the sampling stochasticity and the ambiguity of depth estimation (Ke et al. (2024)). To mitigate this issue, it is desirable to increase the likelihood of each sampled estimate—i.e., to encourage samples to concentrate around the dominant modes of the learned distribution. Doing so reduces sampling variability and suppresses uncertain or noisy depth predictions.

We evaluate on the ETH3D (Schops et al., 2017) and NYUv2 (Nathan Silberman & Fergus, 2012) datasets. As shown in Table 2, TSR outperforms the default DDIM and CNS on prediction accuracy. By sampling from a sharper distribution, TSR yields more probable outputs given the input image. Qualitative comparisons in Fig. 7 further show that TSR produces cleaner depth maps than DDIM, particularly in high-uncertainty regions.

### 5.4 POSE PREDICTION

Previous work Leach et al. (2022); Hsiao et al. (2024); Wang et al. (2023); Zhang et al. (2024) has shown that diffusion models can effectively predict object and camera poses in the  $SO(3)$  space. We demonstrate TSR can improve such models’ accuracy by sampling from a sharper distribution. Our evaluations are based on the  $SO(3)$  diffusion models proposed by Hsiao et al. (2024), where we

	ETH3D		NYUv2	
	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$
DDIM	7.1	90.4	6.0	95.9
+ CNS	6.82	95.6	5.85	<b>96.0</b>
+ TSR	<b>6.68</b>	<b>95.7</b>	<b>5.84</b>	<b>96.0</b>

Table 2: **Quantitative Evaluation of Depth Estimation.** TSR improves depth estimation and outperforms the naive baseline.

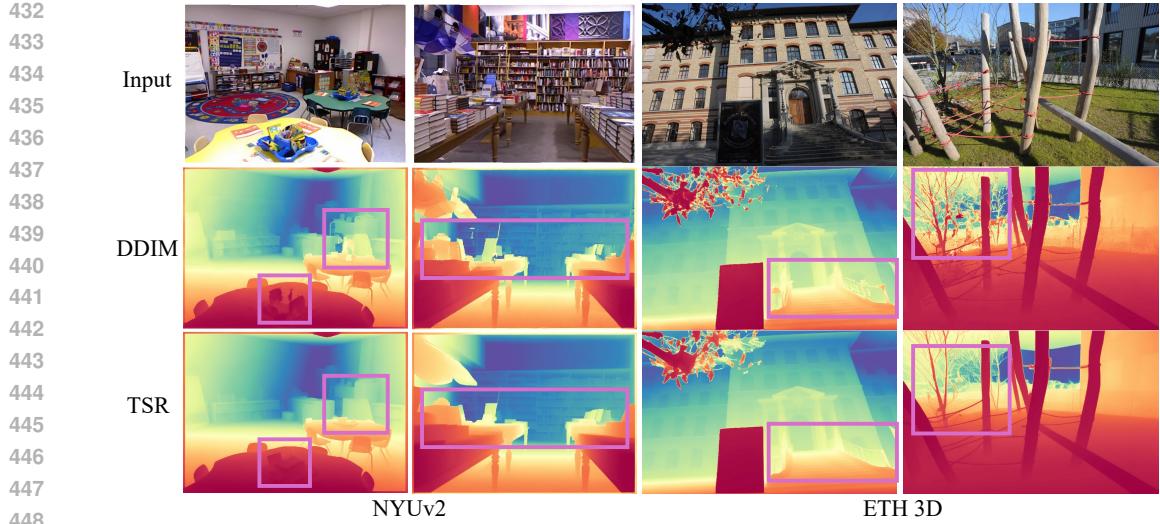


Figure 7: **Qualitative Depth Estimation Comparisons.** Compared to DDIM, TSR with  $k > 1$  predicts cleaner depth in the regions with high uncertainty (highlighted by pink boxes).

452  
453  
454  
455  
456  
457  
458

	Error (deg) $\downarrow$	Acc @ (deg) $\uparrow$		
		0.2	0.5	1.0
Score Sampling	0.444	9.4	68.3	97.9
+ CNS (1600)	<b>0.350</b>	<b>20.0</b>	<b>84.9</b>	<b>99.1</b>
+ TSR (7.0, 0.5)	0.356	18.5	84.0	99.0

459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476

Table 3: **Pose Prediction.** Mean error (deg) and accuracy within thresholds 0.2, 0.5, 1.  $(k, \sigma) = (7.0, 0.5)$  for TSR,  $k = 1600$  for CNS.

apply TSR and evaluate on the SYMSOL dataset Murphy et al. (2021), which contains geometric shapes with a high order of symmetries. We visualize the effect of TSR in 8 where we show the sampled poses on an example image from SYMSOL. TSR samples poses more concentrated around ground truth modes (the circle centers) than the baseline score matching used in Hsiao et al. (2024). In quantitative evaluation (3). TSR predictions have lower average error and higher accuracy under a range of accuracy thresholds compared to score matching sampling, highlighting the benefits of predicting samples close to modes. We find that CNS also reduces pose error, achieving a performance slightly better than TSR on SYMSOL. However, we note that TSR remain robust and applicable over many tasks and sampling methods where constant noise scaling is not possible.

## 5.5 ROBOTIC MANIPULATION

Lastly, we examine the applicability of TSR on predicting robot actions, with a focus on robotic manipulation. One notable difference of this domain compared to others is that the policy only models a distribution of actions for a short horizon, as it is a sequential decision-making problem.

We chose Pi-0 (Black et al., 2025), a generalist robotic flow-matching policy released by Physical Intelligence, finetuned on LIBERO (Liu et al., 2023a), a simulation benchmark for robotic manipulation. Specifically, we evaluate the policy over 10 tasks in the LIBERO-10 benchmark with shared  $(k, \sigma)$  values. The results are in Table 4. Without any further training, TSR improves the performance of 6 tasks and maintains performance for 2 tasks. One notable point is that the two tasks (Task ID 2 and 8) where TSR shows worse performance are precisely those in which the base Pi-0 policy itself exhibits low success rates. This suggests that the suboptimal performance may be due

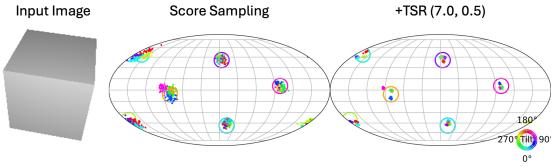


Figure 8: **Predicted poses on SYMSOL.** TSR reduces pose prediction error: each dot marks a sample’s first canonical axis (colored by rotation), while circles denote ground-truth poses.

486 to a common ‘sharpening’ ( $k > 1$ ) hyper-parameter across tasks as this may be suboptimal when  
 487 the policy is not correct, and that tuning  $\text{TSR}$ ’s  $k$  for each task may yield further gains.  
 488

Model / Task ID	0	1	2	3	4	5	6	7	8	9	Average
Pi-0	86.0	<b>97.3</b>	<b>80.7</b>	96.0	<b>84.7</b>	93.3	81.3	94.7	<b>23.3</b>	80.0	81.7
+ TSR(1.25,0.25)	<b>86.7</b>	<b>97.3</b>	77.3	<b>97.3</b>	<b>84.7</b>	<b>96.0</b>	<b>82.7</b>	<b>96.0</b>	21.3	<b>88.7</b>	<b>82.8</b>

492 Table 4: **Results for Robotic Manipulation.** Success rates are computed across 3 seeds, 10 tasks,  
 493 with 50 rollouts per task. The results are computed with the best ( $k, \sigma$ ) for  $\text{TSR}$ .  
 494

## 496 6 DISCUSSION

498 We presented  $\text{TSR}$ , an approach to alter the sampling distribution for pre-trained diffusion and flow  
 499 models. While we demonstrated its efficacy across several (toy and real) tasks, there are fundamen-  
 500 tal limitations worth highlighting. First, unlike temperature scaling,  $\text{TSR}$  can only alter the ‘local’  
 501 sampling and there might exist applications where a global temperature scaling is more desirable  
 502 *e.g.*,  $\text{TSR}$  does not change the weights of the components in a gaussian mixture, only the variance.  
 503 Moreover, while  $\text{TSR}$  does empirically steer the sampling diversity in generic scenarios, the theoreti-  
 504 cal guarantees are limited to simpler settings and one may be able to derive a better algorithm for  
 505 different distributions. Nevertheless, as  $\text{TSR}$  can be readily applied to any off-the-shelf denoising  
 506 diffusion and flow matching model, we believe it would a generally useful technique for the commu-  
 507 nity to explore. In particular, the alternative strategy of ‘constant noise scaling’ is already adopted  
 508 across applications (Yim et al., 2023; Geffner et al., 2025), and our work offers an alternative that is  
 509 empirically better and more widely applicable (*e.g.*, in deterministic sampling).  
 510

## 511 REPRODUCIBILITY STATEMENT

513 We described the proposed algorithm  $\text{TSR}$  for diffusion and flow models in Section 3.3. We include  
 514 the Python code implementation in the supplementary materials. All of the experiments presented  
 515 in this paper are based on open-sourced methods and datasets. The detailed configurations of the  
 516 experiments are described in Section 5 and Section A.  
 517

## 518 REFERENCES

519 Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf  
 520 Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure  
 521 prediction of biomolecular interactions with alphafold 3. In *Nature*, 2024.  
 522

523 Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying  
 524 framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.  
 525

526 Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo  
 527 Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al.  $\pi 0$ : A vision-language-action flow  
 528 model for general robot control. In *RSS*, 2025.  
 529

530 Li Deng. The mnist database of handwritten digit images for machine learning research [best of the  
 531 web]. In *IEEE signal processing magazine*, 2012.  
 532

533 Yilun Du, Conor Durkan, Robin Strudel, Joshua B. Tenenbaum, Sander Dieleman, Rob Fergus,  
 534 Jascha Sohl-Dickstein, Arnaud Doucet, and Will Grathwohl. Reduce, Reuse, Recycle: Composi-  
 535 tional Generation with Energy-Based Diffusion Models and MCMC. In *ICML*, 2023.  
 536

537 Yiquan Duan, Xianda Guo, and Zheng Zhu. Diffusiondepth: Diffusion denoising approach for  
 538 monocular depth estimation. In *ECCV*, 2024.  
 539

540 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam  
 541 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion En-  
 542 glish, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling Rectified Flow  
 543 Transformers for High-Resolution Image Synthesis. In *ICML*, 2024.

540 Felix Faltings, Hannes Stark, Tommi Jaakkola, and Regina Barzilay. Protein fid: Improved evalua-  
541 tion of protein structure generative models. *arXiv preprint arXiv:2505.08041*, 2025.  
542

543 Tomas Geffner, Kieran Didi, Zuobai Zhang, Danny Reidenbach, Zhonglin Cao, Jason Yim, Mario  
544 Geiger, Christian Dallago, Emine Kucukbenli, Arash Vahdat, et al. Proteina: Scaling flow-based  
545 protein structure generative models. In *ICLR*, 2025.

546 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.  
547 Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*,  
548 2017.

549 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *stat*,  
550 2015.

552 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS*, 2022.

553

554 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*,  
555 2020.

556 Tsu-Ching Hsiao, Hao-Wei Chen, Hsuan-Kung Yang, and Chun-Yi Lee. Confronting ambiguity in  
557 6d object pose estimation via score-based diffusion on  $se(3)$ . In *CVPR*, 2024.

558

559 John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger,  
560 Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate  
561 protein structure prediction with alphafold. In *Nature*, 2021.

562 Tero Karras, Miika Aittala, Tuomas Kynkänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine.  
563 Guiding a diffusion model with a bad version of itself. *NeurIPS*, 2024.

564

565 Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad  
566 Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In  
567 *CVPR*, 2024.

568

569 Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.

570 Adam Leach, Sebastian M Schmon, Matteo T. Degiacomi, and Chris G. Willcocks. Denoising  
571 diffusion probabilistic models on  $SO(3)$  for rotational alignment. In *ICLR Workshops*, 2022.

572

573 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching  
574 for Generative Modeling. In *ICLR*, 2023.

575

576 Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero:  
577 Benchmarking knowledge transfer for lifelong robot learning. In *NeurIPS*, 2023a.

578

579 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow Straight and Fast: Learning to Generate and  
580 Transfer Data with Rectified Flow. In *ICLR*, 2023b.

581

582 Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Sain-  
583 ing Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant  
584 transformers. In *ECCV*, 2024.

585

586 Kieran A Murphy, Carlos Esteves, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia.  
587 Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold.  
588 In *ICML*, 2021.

589

590 Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support  
591 inference from rgbd images. In *ECCV*, 2012.

592

593 Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in  
594 gan evaluation. In *CVPR*, 2022.

595

596 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
597 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
598 models from natural language supervision. In *ICML*, 2021.

594 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
595 resolution image synthesis with latent diffusion models. In *CVPR*, 2022.  
596

597 Saurabh Saxena, Abhishek Kar, Mohammad Norouzi, and David J Fleet. Monocular depth estima-  
598 tion using diffusion models. *arXiv preprint arXiv:2302.14816*, 2023.

599 Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc  
600 Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and  
601 multi-camera videos. In *CVPR*, 2017.

602

603 Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi  
604 Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An  
605 open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022.

606

607 Andy Shih, Dorsa Sadigh, and Stefano Ermon. Long horizon temperature scaling. In *ICML*, 2023.

608

609 Marta Skreta, Tara Akhoud-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik,  
610 Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-kac correc-  
611 tors in diffusion: Annealing, guidance, and product of experts. In *Forty-second International  
Conference on Machine Learning*, 2025.

612

613 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*,  
2021a.

614

615 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
616 Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *ICLR*,  
2021b.

617

618 Jianyuan Wang, Christian Rupprecht, and David Novotny. PoseDiffusion: Solving pose estimation  
619 via diffusion-aided bundle adjustment. In *ICCV*, 2023.

620

621 Kevin E Wu, Kevin K Yang, Rianne van den Berg, Sarah Alamdari, James Y Zou, Alex X Lu, and  
622 Ava P Amini. Protein structure generation via folding diffusion. In *Nature*, 2024.

623

624 Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay,  
625 and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation. In  
*ICML*, 2023.

626

627 Jason Y Zhang, Amy Lin, Moneish Kumar, Tzu-Hsuan Yang, Deva Ramanan, and Shubham Tul-  
628 siani. Cameras as rays: Pose estimation via ray diffusion. In *ICLR*, 2024.

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

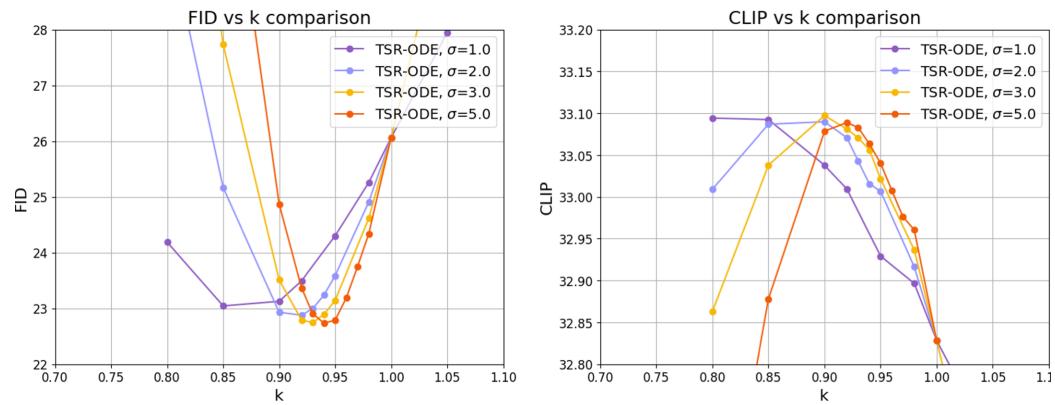
646

647

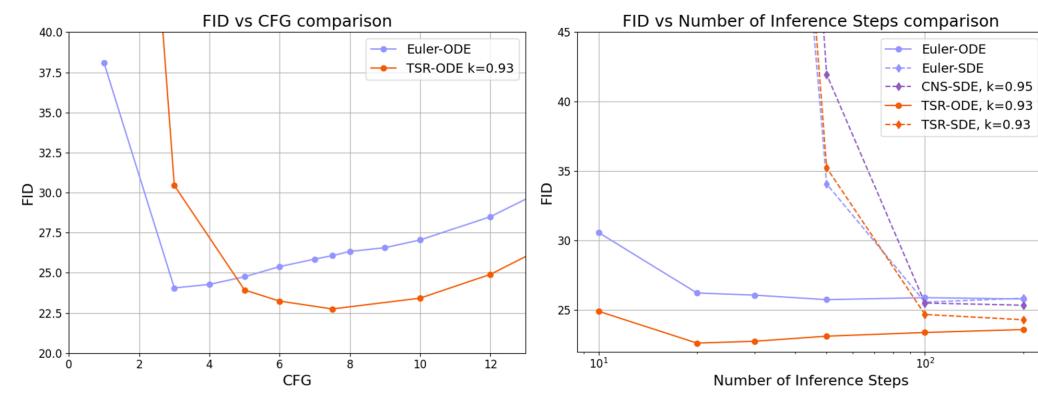
## 648 A ADDITIONAL RESULTS

### 649 A.1 IMAGE GENERATION

650 We show additional qualitative examples generated by Stable Diffusion 3 with TSR in figure 11. We  
 651 show TSR with varying  $k$  and highlight the control over the expression of high-frequency details.  
 652 To better demonstrate the effect of user-input parameters  $(k, \sigma)$  on image quality, we plot FID and  
 653 CLIP scores ablating over different  $(k, \sigma)$  in figure 9. The trends in figure 9 clearly demonstrate that  
 654 TSR with a  $k$  slightly smaller than 1 and various  $\sigma$  values improves both metrics. and the optimal  
 655 performance is achieved at  $k = 0.93, \sigma = 3.0$ . We compare the regular Euler-ODE sampling and  
 656 TSR on different CFG guidance scales in figure 10, highlighting that TSR is orthogonal to CFG and  
 657 improves model performance at various CFG settings. In figure 10, we also present the performance  
 658 of CNS, which has to be applied on a stochastic sampler (Euler-SDE). As Stable Diffusion 3 is a  
 659 flow matching model, stochastic samplers perform significantly worse than ODE samplers, espe-  
 660 cially when the inference steps are less than 100. These results align with the findings in Ma et al.  
 661 (2024). Therefore, CNS does not practically apply to flow matching models like SD3. In Table 5, we  
 662 additionally present quantitative results on Stable Diffusion 2, which is a denoising diffusion-based  
 663 model. While CNS can improve over DDPM sampling, it is outperformed by TSR.  
 664



665 Figure 9: **Ablations over the TSR parameters  $(k, \sigma)$  on Stable Diffusion 3**



666 Figure 10: **Ablations over CFG scale and samplers** *Left:* Comparing regular sampling with  
 667 TSR with various CFG guidance scale on Stable Diffusion 3. *Right:* Comparing deterministic and  
 668 stochastic samplers with TSR or CNS. Stochastic sampling is much worse on flow models, making  
 669 CNS impractical.

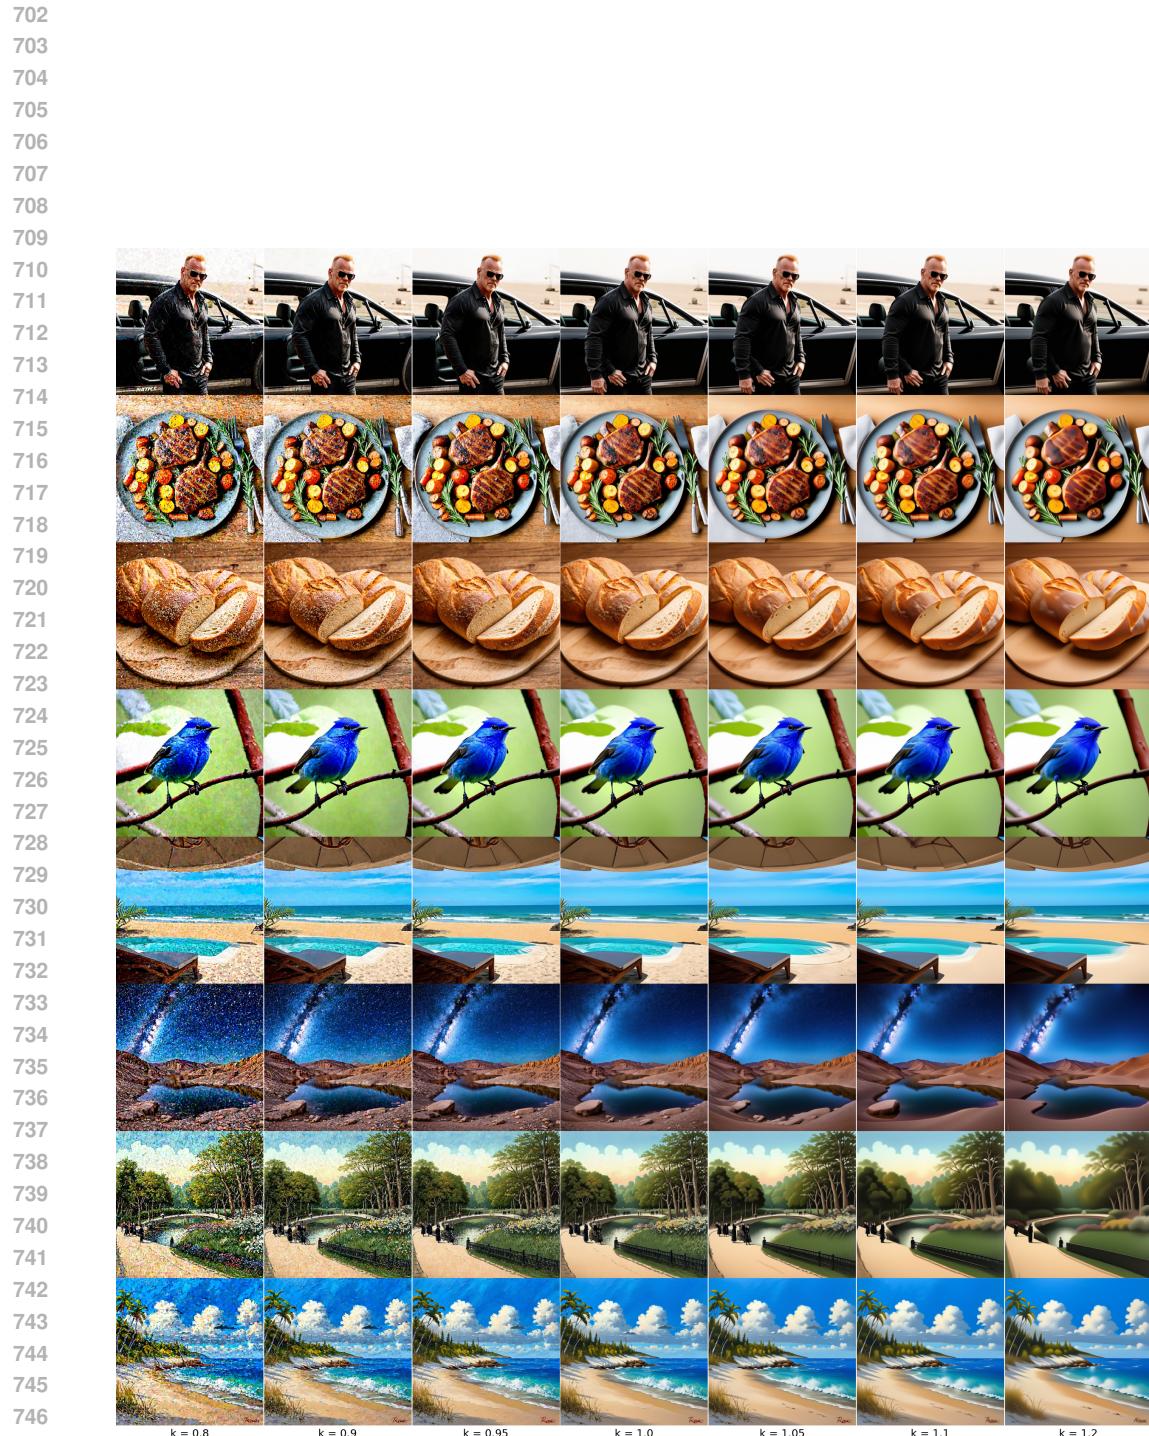


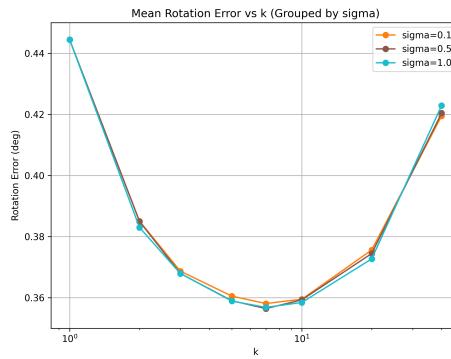
Figure 11: Additional qualitative examples of TSR on Stable Diffusion 3

		FID $\downarrow$	CLIP $\uparrow$
756	DDIM	21.28	33.54
757	+ TSR (0.95, 1.0)	<b>20.05</b>	<b>33.61</b>
758	DDPM	22.81	33.66
759	+ Constant Noise Scaling (0.96)	19.87	33.68
760	+ TSR (0.9, 1.0)	<b>19.57</b>	<b>33.77</b>
761	EulerDiscrete	22.11	33.54
762	+ TSR (0.95, 1.0)	<b>19.95</b>	<b>33.61</b>
763			
764			

765 **Table 5: Results on Stable Diffusion 2.** TSR improve FID and CLIP on various samplers, outper-  
 766 forming CNS.  
 767

## 768 A.2 POSE PREDICTION

769 We show more pose prediction results in figure 13. TSR predicts tighter samples around the ground  
 770 truth mode, which can be observed by the low spread of sampled poses compared to score sam-  
 771 pling. We also include ablations over parameters  $(k, \sigma)$  in figure 12, showing that TSR consistently  
 772 improves accuracy by increasing  $k$ , across different  $\sigma$ .  
 773



788 **Figure 12: Ablations over  $(k, \sigma)$  on pose prediction.** TSR with various  $(k, \sigma)$  configurations  
 789 effectively outperforms the baseline sampling method  $k = 1$ . While TSR is not sensitive to  $\sigma$  in  
 790 pose estimation, TSR reaches optimal performance with  $k \approx 7$ .  
 791

## 792 A.3 DEPTH ESTIMATION

793 We also show the effect of  $\sigma$  and  $k$  on the AbsRel metric in fig. 14. Compared with the DDIM sample  
 794 ( $k = 1$ ), TSR demonstrates consistent performance gain in various  $(k, \sigma)$  configurations. We also  
 795 include more depth samples and comparisons fig. 15. A consistent improvement of TSR result can  
 796 be observed, compared to the DDIM samples.  
 797

798  
 799  
 800  
 801  
 802  
 803  
 804  
 805  
 806  
 807  
 808  
 809

810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
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

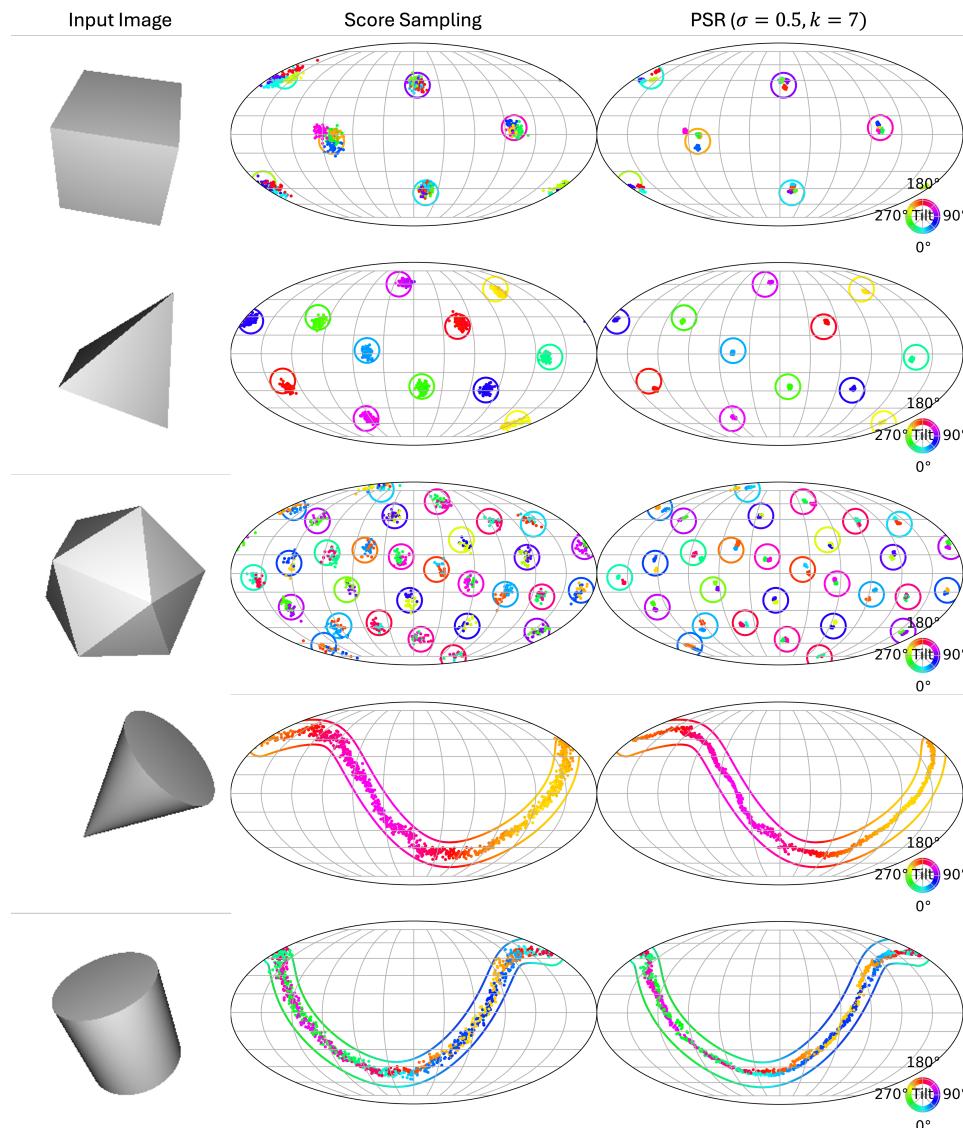


Figure 13: **More predicted poses on SYMSOL.** We show all 5 classes of shapes in SYMSOL. We use  $\sigma = 1, k = 7$  for these visualizations. TSR consistently reduces prediction error across all classes compared to score sampling. We modify the location of samples to exaggerate error by a factor of 15 to show the visual difference given plotting constraints.

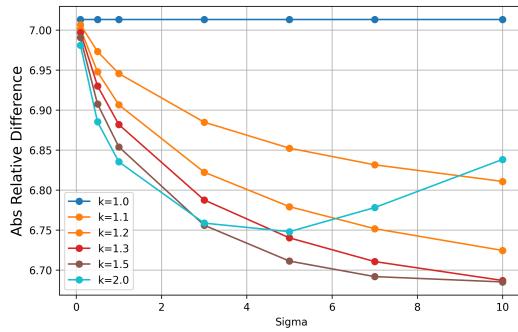


Figure 14: **Effects of  $(k, \sigma)$  on depth estimation.** Comparing with DDIM sample ( $k = 1$ ), TSR demonstrates consistent performance gains in various  $(k, \sigma)$  configurations.

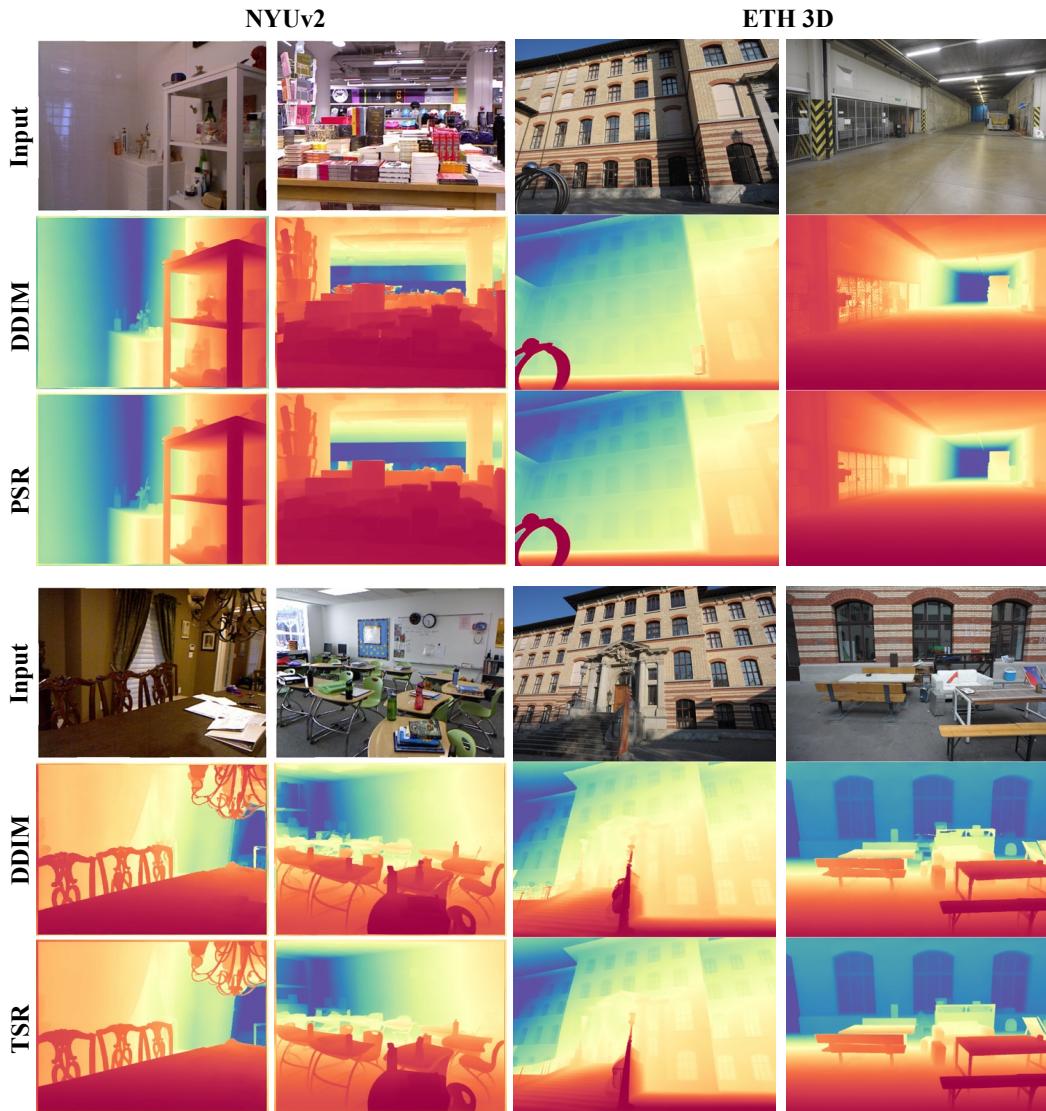


Figure 15: **More Depth Prediction Comparison.** We include more samples from NYUv2 and ETH3D. TSR demonstrates consistent improvement compared to the DDIM samples.

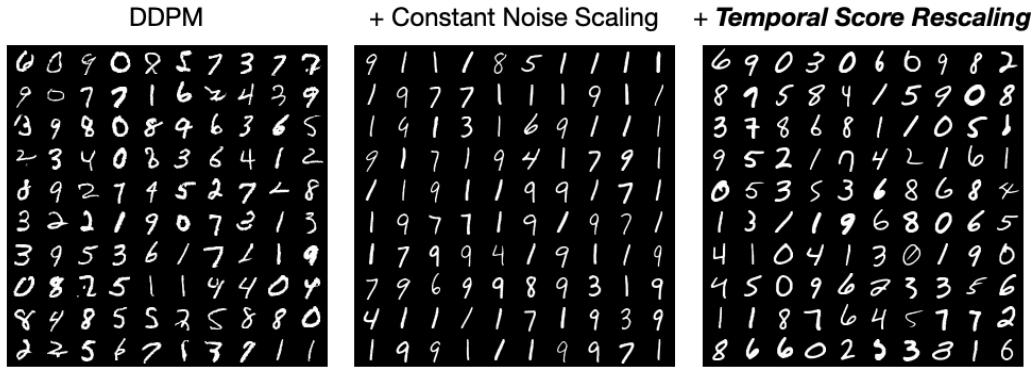
918  
919

## A.4 QUANTIFYING MODE COLLAPSE

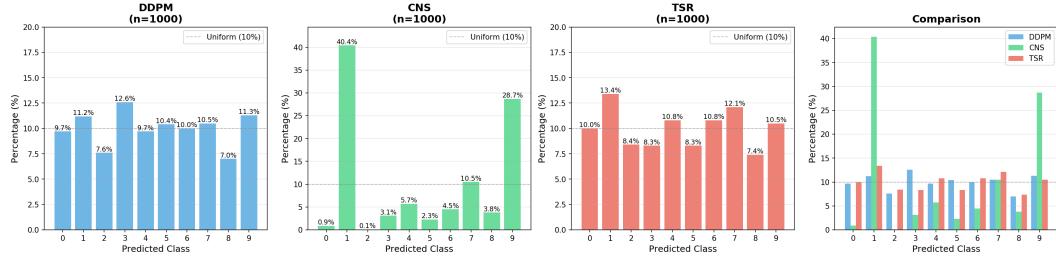
920 To systematically evaluate the mode-collapse behavior of temperature-scaling approaches (Constant  
 921 Noise Scaling and TSR), we train an unconditional DDPM on the MNIST dataset (Deng (2012))  
 922 and apply each sampling method. We additionally train a classifier to label generated samples and  
 923 assess whether Constant Noise Scaling or TSR exhibits mode drop, i.e., produces an imbalanced  
 924 distribution of digits.

925 Figures 16 and 17 summarize the results, using  $k = 5.0$  for Constant Noise Scaling (CNS) and  
 926  $(k, \sigma) = (5.0, 1.0)$  for TSR. CNS disproportionately generates digits ‘1’ (40.4%) and ‘9’ (28.7%),  
 927 likely because their straight or curved components appear frequently across other digits, making  
 928 them easier to synthesize under decreased noise. In contrast, TSR produces a distribution of digits  
 929 that closely matches that of DDPM, indicating that it preserves all modes. Furthermore, TSR gener-  
 930 ates noticeably clearer samples than DDPM, demonstrating the benefit of tempered sampling.

931 In summary, TSR maintains mode coverage on MNIST while improving sample quality.  
 932



946 Figure 16: Samples generated on MNIST using DDPM, Constant Noise Scaling (CNS), and TSR.  
 947 CNS tends to favor generating 1 and 9 while making TSR produces clearer digits while preserving  
 948 diversity across modes.



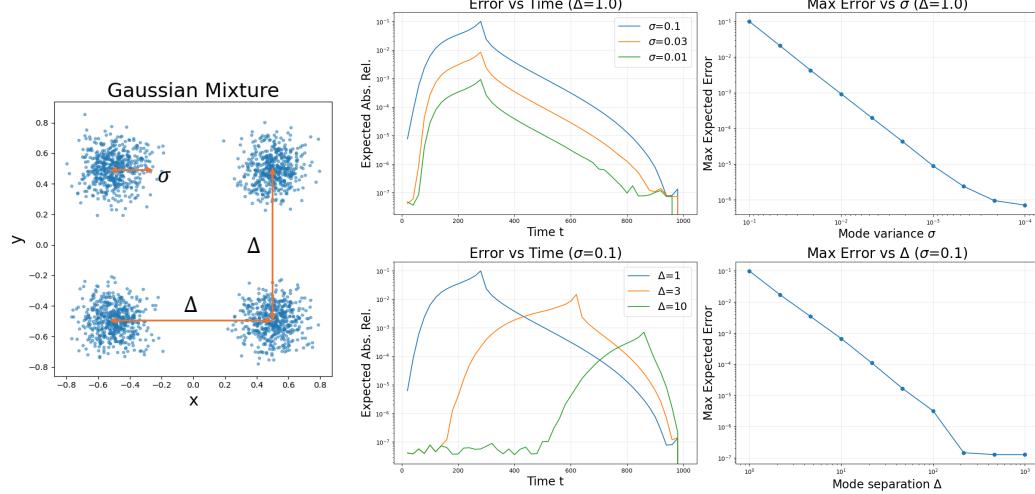
959 Figure 17: Class distribution of generated MNIST samples under DDPM, CNS, and TSR (CNS:  
 960  $k = 5.0$ ; TSR:  $(k, \sigma) = (5.0, 1.0)$ ). CNS exhibits mode imbalance, whereas TSR maintains a  
 961 balanced distribution consistent with the dataset.

962  
963  
964  
965  
966  
967  
968  
969  
970  
971

972    **A.5 EMPIRICAL ANALYSIS ON SCORE APPROXIMATION**  
 973

974    To empirically analyze the our proposed approximation, we conduct an experiment using a 2D  
 975    mixture of four Gaussian distributions, which is visualized in figure 18. We denote the distance  
 976    between neighboring modes as  $\Delta$  and the variance of each mode as  $\sigma$ . Setting the scaling parameter  
 977     $k = 2$ , we systematically vary  $\Delta$  and  $\sigma$  to study the behavior of the approximation error. We quantify  
 978    the deviation by computing the expected absolute relative difference (Abs. Rel.) between the score  
 979    estimated by TSR and the ground-truth score.

980    As illustrated in figure 18, the error vanishes at both ends in the range of timestep  $t$  peaks at interme-  
 981    diate  $t$ . Furthermore, we analyze the maximum error occurring across all timesteps with respect to  $\sigma$   
 982    and  $\Delta$ . The results demonstrate that the maximum error vanishes as the mode variance  $\sigma$  decreases  
 983    or the mode separation  $\Delta$  increases, verifying that the approximation becomes exact as the modes  
 984    are more separated. These results empirically confirm the theoretical bound we proved in Section B.



1002    **Figure 18: Empirical Score Approximation Error:** For the mixture of gaussians depicted in the  
 1003    left, with mode distance  $\Delta$  and mode variance  $\sigma$ , we compute the expected error of TSR approxi-  
 1004    mation at  $k = 2$ . The maximum error is bounded and decreases as  $\sigma$  decreases or  $\sigma$  increases (right  
 1005    column).

1008    **A.6 ADDITIONAL COMPARISON WITH CONSTANT SCORE SCALING**

1009    A less common method that can have similar effect as CNS in temperature sampling is constant  
 1010    score scaling (CSS). Instead of scaling down the noise term like CNS in Eq. 10, CSS constantly  
 1011    scale the score prediciton at each diffusion step, which is equivalent to solving the following reverse  
 1012    SDE:

$$d\mathbf{x} = [f(\mathbf{x}, t) - kg(t)^2 \nabla \log p_t(\mathbf{x})]dt + g(t)d\bar{\mathbf{w}} \quad (10)$$

1016    This method is adopted by Skreta et al. (2025). In figure 19, we additionally evaluate and compare  
 1017    this method on the checkerboard distribution, observing a similar mode-collapse behavior as CNS.  
 1018    We use the same setup as in figure 3.

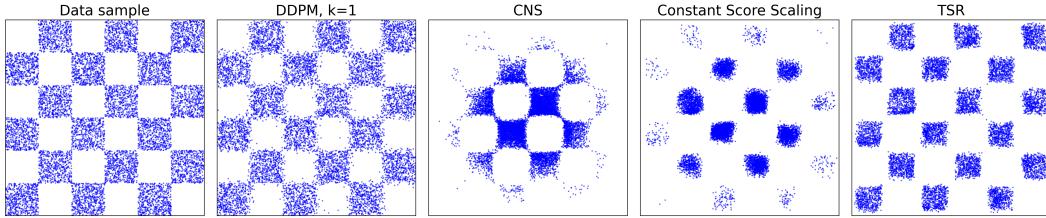


Figure 19: **Evaluating Constant Score Scaling (CSS):** On the 2D checkerboard distribution, both CNS and CSS demonstrates mode-dropping behavior, while only TSR preserves all modes.

## B PROOF FOR TSR FOR MIXTURE OF WELL-SEPARATED GAUSSIANS

We show that for a mixture of well-separated Gaussians, the score approximation in TSR is valid, with the approximation error vanishing asymptotically.

We begin by introducing the notation and defining the estimation error in Section B.1. Our main result is stated in Section B.2. The proof of this result is given in Section B.3, supported by several lemmas whose proofs are provided in Section B.4.

Notations:

- $\alpha_t, \sigma_t, k$ : Diffusion/flow schedule coefficients and sharpening factor.
- $p_t^k(\mathbf{x})$ : Induced distribution at time  $t$  given the data distribution sharpened by  $k$ .
- $\Delta \gg \delta$ : Distance between two mixture means at  $t = 0$ . Define  $\Delta_t = \alpha_t \Delta$ .
- $\sigma$ : Variance of each Gaussian in the mixture at  $t = 0$ .
- $\sigma_{t,k}^2 \equiv \frac{\alpha_t^2 \sigma^2}{k} + \sigma_t^2$ : Variance of each Gaussian at time  $t$  with sharpening factor  $k$ .
- $\delta_{t,n}(\mathbf{x}) \equiv \mathbf{x} - \alpha_t \boldsymbol{\mu}_n$ : Offset vector from  $\mathbf{x}$  to the center of the  $n$ -th Gaussian at diffusion time  $t$ .
- $p_{t,n}^k(\mathbf{x}) \propto \exp\left(-\frac{\|\delta_{t,n}(\mathbf{x})\|^2}{\sigma_{t,k}^2}\right)$ : Unnormalized density of  $\mathbf{x}$  under the  $n$ -th Gaussian.
- $w_{t,n}^k(\mathbf{x}) \equiv \frac{p_{t,n}^k(\mathbf{x})}{\sum_m p_{t,m}^k(\mathbf{x})}$ : Responsibility of the  $n$ -th Gaussian for  $\mathbf{x}$ .
- $N$ : Number of Gaussians in the mixture. Dependent on the dataset only.
- $d$ : Dimensionality of the data. i.e.  $d = 2$  for 2D Gaussian Mixture.
- $\Delta_{\max} = \max_{i,j} |\boldsymbol{\mu}_i - \boldsymbol{\mu}_j|$ : Maximum pairwise distance between Gaussian means in the mixture. For a general dataset, this term is bounded by  $(N - 1)\Delta$ .

### B.1 ERROR IN TSR SCORE APPROXIMATION.

**Score.** The score of the original data is given by:

$$\nabla \log p_t(\mathbf{x}) = -\frac{1}{(\alpha_t^2 \sigma^2 + \sigma_t^2)} \sum_n w_{t,n}^1(\mathbf{x}) \delta_{t,n}(\mathbf{x})$$

For the target distribution  $p^k(\mathbf{x}_0) = \sum_i \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_i, \frac{\sigma^2}{k} \mathcal{I})$  the corresponding noisy distribution  $p^k(\mathbf{x}_t) = \sum_i \mathcal{N}(\mathbf{x}; \alpha_t \boldsymbol{\mu}_i, (\frac{\alpha_t^2 \sigma^2}{k} + \sigma_t^2) \mathcal{I})$ , we have:

$$\nabla \log p_t^k(\mathbf{x}) = -\frac{1}{(\alpha_t^2 \sigma^2/k + \sigma_t^2)} \sum_n w_{t,n}^k(\mathbf{x}) \delta_{t,n}(\mathbf{x})$$

1080 In TSR, we approximate the score of  $p^k(\mathbf{x}_t)$  by:  
1081

$$\begin{aligned} 1082 \nabla \log \tilde{p}_t^k(\mathbf{x}) &\approx \frac{\alpha_t^2 \sigma^2 + \sigma_t^2}{\alpha_t^2 \sigma^2 / k + \sigma_t^2} \nabla \log p_t(\mathbf{x}) = \frac{\sigma_{t,1}^2}{\sigma_{t,k}^2} \left( -\frac{1}{\sigma_{t,1}^2} \sum_n w_{t,n}^1 \delta_{t,n}(\mathbf{x}) \right) \\ 1083 \\ 1084 &= -\frac{1}{\sigma_{t,k}^2} \sum_n w_{t,n}^1 \delta_{t,n}(\mathbf{x}) \\ 1085 \\ 1086 \end{aligned}$$

1087 **Definition B.1** (Error in TSR Score Approximation). Define the amount of error in the score ap-  
1088 proximation as the expected difference between the scores:  
1089

$$1090 \text{Error}(t) = \mathbb{E}_{\mathbf{x} \sim p_t^k} \frac{1}{\sigma_{t,k}^2} \left\| \sum_n (w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})) \delta_{t,n}(\mathbf{x}) \right\| \\ 1091 \\ 1092$$

## 1093 B.2 UPPER BOUND OF THE ERROR

1095 The objective of this proof is to establish a bound on the error term  $\text{Error}(t)$ . Our main results are  
1096 as follows:

1097 **Theorem B.2** (Upper Bound of the Error). *For  $\text{Error}(t)$ , there exists two upper bounds:*

$$\begin{aligned} 1099 \text{Error}(t) &\leq B_{\text{exp}} = 6 \cdot \frac{\alpha_t \Delta_{\max}}{\sigma_{t,k}^2} \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right) \\ 1100 \\ 1101 \\ 1102 \text{Error}(t) &\leq B_{\text{poly}} = \frac{\alpha_t \Delta_{\max}}{4\sigma_{t,k}^2} \left( \frac{1}{\sigma_{t,k}^2} - \frac{1}{\sigma_{t,1}^2} \right) N(d\sigma_{t,k}^2 + \alpha_t^2 \Delta_{\max}^2) \\ 1103 \\ 1104 \end{aligned}$$

1105 **Theorem B.3** (Vanishing Behavior of Error). *Assuming  $\sigma = \epsilon \Delta$ , when  $1 - \alpha_t^2 > \sqrt{\epsilon}$ , we have:*

$$1106 B_{\text{poly}} \sim O(\sqrt{\epsilon}) \\ 1107$$

1108 ; When  $1 - \alpha_t^2 \leq \sqrt{\epsilon}$ , (i.e.  $\alpha_t \approx 1$ ) we have:

$$1109 B_{\text{exp}} \sim O\left(\frac{1}{\sqrt{\epsilon} \Delta} \exp\left(-\frac{1}{\sqrt{\epsilon}}\right)\right) \\ 1110 \\ 1111$$

1112 **Conclusion.** Combining theorem B.2 and theorem B.3, when  $\epsilon \rightarrow 0$ , we have  $\text{Error}(t) \rightarrow 0$ .  
1113 Therefore, when the Gaussians are well-separated ( $\epsilon \rightarrow 0$ ), the approximation error vanishes to 0.  
1114

## 1115 B.3 PROOF OF THEOREM

1117 Before proving the theorems, we first state several lemmas that are useful to the proof, whose proof  
1118 will be given in the next section.  
1119

1120 **Lemma B.4.** *The TSR approximation error  $\text{Error}(t)$  is bounded as follows:*

$$1121 \text{Error}(t) \leq \frac{\alpha_t \Delta_{\max}}{\sigma_{t,k}^2} \mathbb{E}_{\mathbf{x} \sim p_t^k} \| \text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) \| \quad (11) \\ 1122 \\ 1123$$

1124 , where  $\text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) = \sum_n \|w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})\|$ .  
1125

1126 **Lemma B.5.** *There exists a polynomial bound for  $\mathbb{E}_{\mathbf{x} \sim p_t^k} \| \text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) \|$ :*

$$1127 \mathbb{E}_{\mathbf{x} \sim p_t^k} \| \text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) \| \leq 6 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right) \\ 1128 \\ 1129$$

1131 **Lemma B.6.** *There exists an exponential bound for  $\mathbb{E}_{\mathbf{x} \sim p_t^k} \| \text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) \|$ :*  
1132

$$1133 \mathbb{E}_{\mathbf{x} \sim p_t^k} \| \text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) \| \leq \frac{1}{4} \left( \frac{1}{\sigma_{t,k}^2} - \frac{1}{\sigma_{t,1}^2} \right) N(d\sigma_{t,k}^2 + \alpha_t^2 \Delta_{\max}^2)$$

1134 *Proof of Theorem B.2.* Combining Lemma B.4 and Lemma B.5, we obtain the polynomial bound  
 1135 for  $Error(t)$ .

1136 Similarly, Lemma B.4 and Lemma B.6 will give us the exponential bound for  $Error(t)$ .

1138  $\square$

1139  
 1140 *Proof of Theorem B.3.* For simplicity, we assume diffusion scheduling, that is,  $\sigma_t^2 = 1 - \alpha_t^2$  in this  
 1141 part. We also assume  $\sigma = \epsilon\Delta$ . As the dataset is fixed, we can rewrite  $\Delta_{\max} = c\Delta$ , where  $c$  is a  
 1142 constant that only depends on the dataset.

### 1143 Vanishing of Polynomial Bound

1144 Following the polynomial bound from B.2, we have:

$$1145 \begin{aligned} B_{\text{poly}} &= N \frac{\alpha_t \Delta_{\max}}{\sigma_{t,k}^2} \left( \frac{(1 - 1/k)\sigma^2 \alpha_t^2}{\sigma_{t,1}^2 \sigma_{t,k}^2} \right) (d\sigma_{t,1}^2 + \alpha_t^2 \Delta_{\max}^2) \\ 1146 &= N(1 - 1/k) \frac{\alpha_t^3 \Delta_{\max} \sigma^2}{\sigma_{t,k}^4} \left( d + \frac{\alpha_t^2 \Delta_{\max}^2}{\sigma_{t,1}^2} \right) \end{aligned}$$

1151 Consider  $1 - \alpha_t^2 > \sqrt{\epsilon}\Delta^2$ , we have:  $\sigma_{t,k}^2 = \alpha_t^2 \sigma^2 / k + (1 - \alpha_t^2) > (1 - \alpha_t^2) > \sqrt{\epsilon}\Delta^2$ . Therefore,  
 1152 we have:

$$1153 \frac{\alpha_t^3 \Delta_{\max} \sigma^2}{\sigma_{t,k}^4} d \leq \frac{c \alpha_t^3 \epsilon^2 \Delta^3}{\epsilon \Delta^4} d = \alpha_t^3 c d \frac{\epsilon}{\Delta}$$

1154 Since  $\alpha_t \leq 1$  and  $c$  and  $d$  are constant given a dataset, we can absorb them into a constant. Therefore,  
 1155  $\frac{\alpha_t^3 c \Delta_{\max} \sigma^2}{\sigma_{t,k}^4} d \leq C_1 \frac{\epsilon}{\Delta}$ , for some  $C_1 = O(cd)$ .

1156 Similarly to previously proved, for the second term,  $\frac{\alpha_t^3 \Delta_{\max} \sigma^2}{\sigma_{t,k}^4} \cdot \frac{\alpha_t^2 \Delta_{\max}^2}{\sigma_{t,1}^2}$ , we have:

$$1157 \frac{\alpha_t^5 \sigma^2 \Delta_{\max}^3}{\sigma_{t,k}^4 \sigma_{t,1}^2} \leq \frac{\alpha_t^5 c^3 \Delta^3 (\epsilon^2 \Delta^2)}{\epsilon \Delta^4 \cdot \sqrt{\epsilon} \Delta^2} \leq C_2 \frac{\sqrt{\epsilon}}{\Delta}$$

1158 , where  $C_2$  is a constant term based on the dataset (and  $\alpha_t$ ).

1159 Therefore, we have the following.

$$1160 B_{\text{poly}} \leq C_1 \frac{\epsilon}{\Delta} + C_2 \frac{\sqrt{\epsilon}}{\Delta} \leq C \frac{\sqrt{\epsilon}}{\Delta}$$

1161 We can see that the polynomial bound is  $O(\sqrt{\epsilon})$  for such  $\alpha_t$ , which goes to 0 as  $\epsilon \rightarrow 0$

### 1162 Vanishing of Exponential Bound

1163 Assuming the diffusion schedule, and consider  $\alpha_t$  such that  $1 - \alpha_t^2 < \sqrt{\epsilon}\Delta^2$ , we have:

$$1164 B_{\text{exp}} = 6 \frac{\alpha_t \Delta_{\max}}{\alpha_t^2 \sigma^2 / k + 1 - \alpha_t^2} \exp\left(-\frac{\alpha_t^2 \Delta^2}{8(\alpha_t^2 \sigma^2 + 1 - \alpha_t^2)}\right)$$

1165 With our assumption of  $\sigma = \epsilon\Delta$ , for a small  $\epsilon$ :

$$1166 \begin{aligned} \alpha_{t,1}^2 &= \alpha_t^2 \sigma^2 + 1 - \alpha_t^2 = \alpha_t^2 \epsilon^2 \Delta^2 + (1 - \alpha_t^2) \\ 1167 &\leq 2(1 - \alpha_t^2) \leq 2\sqrt{\epsilon} \Delta^2 \\ 1168 &-\frac{\alpha_t^2 \Delta^2}{8\alpha_{t,1}^2} \leq -\frac{\alpha_t^2 \Delta^2}{8 \cdot 2\sqrt{\epsilon} \Delta^2} = -\frac{\alpha_t^2}{16\sqrt{\epsilon}} \end{aligned}$$

1169 Therefore,

$$1170 \exp\left(-\frac{\alpha_t^2 \Delta^2}{8\alpha_{t,1}^2}\right) \leq \exp\left(-\frac{\alpha_t^2}{16\sqrt{\epsilon}}\right)$$

1171 As  $\alpha_t^2 \sigma^2 / k + 1 - \alpha_t^2$  is dominant by  $1 - \alpha_t^2$ , we have  $\alpha_t^2 \sigma^2 / k + 1 - \alpha_t^2 \approx 1 - \alpha_t^2$ .

1188 Therefore, we have:

$$1190 \quad B_{\exp} \leq 6 \frac{\alpha_t \Delta_{\max}}{\alpha_t^2 \sigma^2 / k + 1 - \alpha_t^2} \exp\left(-\frac{\alpha_t^2}{16\sqrt{\epsilon}}\right) \approx \frac{6c\alpha_t}{\sqrt{\epsilon}\Delta} \exp\left(-\frac{\alpha_t^2}{16\sqrt{\epsilon}}\right)$$

1192 As we consider  $\alpha_t$  such that  $1 - \alpha_t^2 < \sqrt{\epsilon}\Delta^2$ , then we can write the exponential bound as  
1193  $O\left(\frac{1}{\sqrt{\epsilon}\Delta} \exp\left(-\frac{1}{\sqrt{\epsilon}}\right)\right)$ , which also vanishes as  $\epsilon \rightarrow 0$ .

1194 **Conclusion**

1195 In both cases, at least one bound is vanishingly small as  $\epsilon \rightarrow 0$ .

1196  $\square$

1200 **B.4 PROOF OF LEMMA**

1202 *Proof of Lemma B.4. Upper bound of the error*

1203 Using the triangle inequality and the fact that  $\sum_n w_{t,n}^1(\mathbf{x}) = 1$  and  $\sum_n w_{t,n}^k(\mathbf{x}) = 1$ , we have the  
1204 following result:

1205

$$\begin{aligned} 1207 \quad \text{Error}(t) &= \mathbb{E}_{\mathbf{x} \sim p_t^k} \frac{1}{\sigma_{t,k}^2} \left\| \sum_n (w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})) \boldsymbol{\delta}_{t,n}(\mathbf{x}) \right\| \\ 1208 &\leq \frac{1}{\sigma_{t,k}^2} \mathbb{E}_{\mathbf{x} \sim p_t^k} \left\| \sum_n (w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})) \|\boldsymbol{\delta}_{t,n}(\mathbf{x})\| \right\| \\ 1209 &\leq \frac{1}{\sigma_{t,k}^2} \mathbb{E}_{\mathbf{x} \sim p_t^k} \left\| \sum_n ((w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})) \alpha_t \boldsymbol{\delta}_{\max}) \right\| \\ 1210 &\leq \frac{\alpha_t \boldsymbol{\delta}_{\max}}{\sigma_{t,k}^2} \mathbb{E}_{\mathbf{x} \sim p_t^k} \sum_n \|w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})\| \end{aligned}$$

1217 Therefore, the approximation error is bounded as follows:

$$1219 \quad \text{Error}(t) \leq \frac{\alpha_t \Delta_{\max}}{\sigma_{t,k}^2} \mathbb{E}_{\mathbf{x} \sim p_t^k} \|\text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x}))\| \quad (12)$$

1221 , where  $\text{dist}(w_t^1(\mathbf{x}), w_t^k(\mathbf{x})) = \sum_n \|w_{t,n}^1(\mathbf{x}) - w_{t,n}^k(\mathbf{x})\|$ .

1223  $\square$

1224 *Proof of Lemma B.5. Exponential Bound*

1226 Following our problem setting, we have:

1227

$$1228 \quad p_t(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \mathcal{N}(x; \alpha_t \boldsymbol{\mu}_i, (\alpha_t^2 \sigma^2 + \sigma_t^2) I).$$

1231 and

$$1232 \quad q_t(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \mathcal{N}(x; \alpha_t \boldsymbol{\mu}_i, (\frac{\alpha_t^2 \sigma^2}{k} + \sigma_t^2) I).$$

1234 , where  $p_t(\mathbf{x})$  is the original distribution, and  $q_t(\mathbf{x})$  is the desired distribution with altered variance.

1235 For each  $\mathbf{x}$ , the responsibility vector under a mixture is defined as:

1236

$$1238 \quad r^{(p)}(\mathbf{x}) = (r_1^{(p)}(\mathbf{x}), \dots, r_N^{(p)}(\mathbf{x}))$$

1239

1240 , where  $r_i^{(p)}(\mathbf{x}) = \frac{\mathcal{N}(x; \alpha_t \boldsymbol{\mu}_i, \alpha_t^2 \sigma^2 + \sigma_t^2)}{\sum_{j=1}^N \mathcal{N}(x; \alpha_t \boldsymbol{\mu}_j, \alpha_t^2 \sigma^2 + \sigma_t^2)}$ .  $r^{(q)}(\mathbf{x})$  is defined analogously as  $r_i^{(q)}(\mathbf{x}) =$   
1241  $\frac{\mathcal{N}(x; \alpha_t \boldsymbol{\mu}_i, \alpha_t^2 \sigma^2 / k + \sigma_t^2)}{\sum_{j=1}^N \mathcal{N}(x; \alpha_t \boldsymbol{\mu}_j, \alpha_t^2 \sigma^2 / k + \sigma_t^2)}$ .

1242 Now we have  $\mathbb{E}_{\mathbf{x} \sim p_t^k} \|dist(w_t^1(\mathbf{x}), w_t^k(\mathbf{x}))\| = \mathbb{E}_{\mathbf{x} \sim p_t^k} [D(\mathbf{x})]$ , where  $D(\mathbf{x}) := \|r^{(p)}(\mathbf{x}) - r^{(q)}(\mathbf{x})\|_1$ .

1245 Define  $i(\mathbf{x}) = \max_i r_i$ , and  $e_i$  as the one-hot vector where the  $i$ th entry is one. Using the triangle  
1246 inequality, we have:

1247  $D(\mathbf{x}) = \|r^{(p)} - r^{(q)}\|_1 \leq \|r^{(p)} - e_{i_p(\mathbf{x})}\|_1 + \|e_{i_p(\mathbf{x})} - e_{i_q(\mathbf{x})}\|_1 + \|e_{i_q(\mathbf{x})} - r^{(q)}\|_1.$

1249 , and that

1250  $\|r^{(p)} - e_{i_p(\mathbf{x})}\|_1 = 2(1 - r_{i_p(\mathbf{x})}^p(\mathbf{x}))$

1252  $\|e_{i_p} - e_{i_q}\|_1 = 2 * \mathbf{1}\{i_p \neq i_q\}$

1253  $\|r^{(q)} - e_{i_q(\mathbf{x})}\|_1 = 2(1 - r_{i_q(\mathbf{x})}^q(\mathbf{x}))$

1255 **Concentration of responsibilities for the true component** Let:

1257  $\epsilon := \max_{i \neq j} \mathbb{P}_{x \sim \mathcal{N}(\mu_i, \sigma^2)} [\|x - \mu_j\| < \|x - \mu_i\|]$

1259 That is, the probability that a sample from component  $i$  is closer to another component  $j$ . Then:

1261  $\mathbb{E}_{x \sim p} \left[ 1 - \max_j r_j^{(p)}(\mathbf{x}) \right] \leq \epsilon \Rightarrow \mathbb{E}_{x \sim p} [D(\mathbf{x})] \approx 2\epsilon$

1264 Recall  $\Delta := \min_{i \neq j} \|\mu_i - \mu_j\|$  to be the minimum pairwise distance between the means. Using  
1265 Gaussian tail bounds, we can approximate:

1266  $\epsilon \approx \exp \left( -\frac{\Delta^2}{8\sigma^2} \right)$

1269 Hence, we have:

1270  $E_{x \sim p_t^k} (2(1 - r_{i_p(\mathbf{x})}^p(\mathbf{x}))) \leq 2 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right)$

1274  $E_{x \sim p_t^k} (2(1 - r_{i_q(\mathbf{x})}^q(\mathbf{x}))) \leq 2 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,k}^2} \right)$

1277 **Bounding  $\Pr(i_p \neq i_q)$**

1278 As  $p_t(\mathbf{x})$  and  $q_t(\mathbf{x})$  share the same modes, we have  $\Pr(i_p \neq i_q) \leq \sum_i \Pr(i_p \neq i_q \mid x \sim$   
1279 component  $i$ )  $\Pr(x \text{ from } i)$ , which can also be bounded using Gaussian tail bounds as above.

1281 Therefore, we have:

1282  $E_{x \sim p_t^k} (D(\mathbf{x})) \leq E_x (\|r^{(p)} - e_{i_p(\mathbf{x})}\|_1) + E_x (\|e_{i_p(\mathbf{x})} - e_{i_q(\mathbf{x})}\|_1) + E_x (\|e_{i_q(\mathbf{x})} - r^{(q)}\|_1)$   
1283  $= E_x (2(1 - r_{i_p(\mathbf{x})}^p(\mathbf{x}))) + E_x (2 * \mathbf{1}\{i_p \neq i_q\}) + E_x (2(1 - r_{i_q(\mathbf{x})}^q(\mathbf{x})))$   
1284  $\leq 2 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right) + (\exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right) + \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,k}^2} \right)) + 2 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,k}^2} \right)$   
1285  $\leq 6 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right)$

1291 Finally:

1293  $E_{x \sim p_t^k} (D(\mathbf{x})) \leq 6 \cdot \exp \left( -\frac{\alpha_t^2 \Delta^2}{8\sigma_{t,1}^2} \right)$

1294  $\square$

1296 **Proof of Lemma B.6. Polynomial Bound**

1297 We consider the softmax representation of the responsibilities:

1299

$$1300 w_t^k(\mathbf{x}) = \text{softmax}(z_t^k(\mathbf{x})), \quad \text{where} \quad z_{t,n}^k(\mathbf{x}) := -\frac{\|\mathbf{x} - \alpha_t \boldsymbol{\mu}_n\|^2}{2\sigma_{t,k}^2}.$$

1301

1302 . Using the Softmax Lipschitz bound that  $\|\text{softmax}(z) - \text{softmax}(z')\|_1 \leq 1/2\|z - z'\|_1$ , we  
1303 have:

1304

$$1306 \quad \|w_t^k(\mathbf{x}) - w_t^1(\mathbf{x})\|_1 \leq \frac{1}{2} \|z_t^k(\mathbf{x}) - z_t^1(\mathbf{x})\|_1.$$

1308

1309 Compute the logits difference coordinatewise:

1310

$$1311 \quad z_{t,n}^k(\mathbf{x}) - z_{t,n}^1(\mathbf{x}) = -\frac{\|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2}{2\sigma_{t,k}^2} + \frac{\|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2}{2\sigma_{t,1}^2}$$

1312

$$1313 \quad = \frac{1}{2} \left( \frac{1}{\sigma_{t,1}^2} - \frac{1}{\sigma_{t,k}^2} \right) \|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2.$$

1314

1315 Adding absolute values,

1316

$$1319 \quad \|z_t^k(\mathbf{x}) - z_t^1(\mathbf{x})\|_1 = \frac{1}{4} \left( \frac{1}{\sigma_{t,k}^2} - \frac{1}{\sigma_{t,1}^2} \right) \sum_{n=1}^N \|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2$$

1320

1321 **Bounding**  $\mathbb{E}_x \left[ \sum_{n=1}^N \|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2 \right]$

1324 Let  $x \sim p_t^k$  be drawn from the mixture with means  $\{\alpha_t \boldsymbol{\mu}_i\}$  and variance  $\sigma_{t,k}^2$ . Write expectation as  
1325 mixture-average:

1326

$$1327 \quad \mathbb{E}_x \left[ \sum_{n=1}^N \|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2 \right] = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{x \sim \mathcal{N}(\alpha_t \boldsymbol{\mu}_i, \sigma_{t,k}^2 I)} \left[ \sum_{n=1}^N \|x - \alpha_t \boldsymbol{\mu}_n\|^2 \right].$$

1329

1330 When the sample was generated from component  $i$ , for any other  $n$ , we have

1331

$$1332 \quad \mathbb{E}\|x - \alpha_t \boldsymbol{\mu}_n\|^2 = \mathbb{E}\|x - \alpha_t \boldsymbol{\mu}_i + \alpha_t \boldsymbol{\mu}_i - \alpha_t \boldsymbol{\mu}_n\|^2 = \mathbb{E}\|x - \alpha_t \boldsymbol{\mu}_i\|^2 + \|\alpha_t \boldsymbol{\mu}_i - \alpha_t \boldsymbol{\mu}_n\|^2$$

1333 , because the cross-term has zero mean.

1334 Since the first term equals the trace of the covariance  $= d\sigma_{t,1}^2$ , we have:

1335

$$1337 \quad \mathbb{E}\|x - \alpha_t \boldsymbol{\mu}_n\|^2 = d\sigma_{t,1}^2 + \|\alpha_t(\boldsymbol{\mu}_i - \boldsymbol{\mu}_n)\|^2$$

1338

1339 Summing over all  $N$  (including  $n=i$ , for which the pairwise term is zero) gives  
1340  $\mathbb{E}_{x \sim \mathcal{N}(\alpha_t \boldsymbol{\mu}_i, \sigma_{t,1}^2 I)} \left[ \sum_{n=1}^N \boldsymbol{\delta}_{t,n}(\mathbf{x}) \right] = N d\sigma_{t,1}^2 + \sum_{n=1}^N \|\alpha_t(\boldsymbol{\mu}_i - \boldsymbol{\mu}_n)\|^2.$

1341

1342 Now, bound the pairwise squared distances by the diameter squared:  $\|\alpha_t(\boldsymbol{\mu}_{t,i} - \boldsymbol{\mu}_{t,n})\|^2 \leq \alpha_t^2 \Delta_{\max}^2$ .

1343 Therefore, we have:  $\mathbb{E}_x \left[ \sum_{n=1}^N \|\boldsymbol{\delta}_{t,n}(\mathbf{x})\|^2 \right] \leq N(d\sigma_{t,1}^2 + \alpha_t^2 \Delta_{\max}^2)$ .

1344 We then have the polynomial bound as:

1345

$$1347 \quad \mathbb{E}_{x \sim p} [D(\mathbf{x})] \leq \frac{1}{4} \left( \frac{1}{\sigma_{t,k}^2} - \frac{1}{\sigma_{t,1}^2} \right) N(d\sigma_{t,k}^2 + \alpha_t^2 \Delta_{\max}^2)$$

1348

1349

□

1350    **C CONSTANT NOISE SCALING**

1352    In this section, we provide a more detailed analysis of Constant Noise Scaling. As discussed in Sec-  
 1353    tion 2, CNS has been adopted as a practical technique to control trade-off sample variance and diver-  
 1354    sity. We intuitively explain and empirically verify that CNS does not correspond to true temperature  
 1355    scaling. We now provide a more rigorous proof that CNS cannot produce the temperature-scaled  
 1356    distribution. Following Song et al. (2021b), a regular score-based model  $\mathbf{s}_\theta(\mathbf{x}, t) = \nabla \log p_t(\mathbf{x})$   
 1357    trained on data distribution  $p_0(\mathbf{x})$  can sample by solving the reverse time diffusion SDE:

1358    
$$d\mathbf{x} = [f(t)\mathbf{x} - g(t)^2 \mathbf{s}_\theta(\mathbf{x}, t)]dt + g(t)d\bar{\mathbf{w}} \quad (13)$$

1360    where  $f(t), g(t)$  are the time-dependent drift and diffusion coefficients,  $d\bar{\mathbf{w}}$  is the standard Wiener  
 1361    process. CNS solves the following SDE instead:

1362    
$$d\mathbf{x} = [f(t)\mathbf{x} - (\frac{g(t)}{\sqrt{k}})^2 (k\mathbf{s}_\theta(\mathbf{x}, t))]dt + \frac{g(t)}{\sqrt{k}} d\bar{\mathbf{w}} \quad (14)$$

1365    Practically, CNS scales the stochastic noise added at each sampling step by  $1/\sqrt{k}$ . When  $k > 1$ ,  
 1366    less noise is added and the process generates samples with reduced variance, and vice versa. To  
 1367    analyze the relationship between CNS and temperature scaling, we denote the temperature-scaled  
 1368    data distribution  $q_0(\mathbf{x})$ , such that  $q_0(\mathbf{x}) \propto p_0(\mathbf{x})^k$ .

1369    **Theorem C.1.** *For general data distribution  $p_0(\mathbf{x})$ , there is no prior distribution  $q'_T(\mathbf{x})$ , such that  
 1370    Eq. 14 starts from  $q'_T(\mathbf{x})$  and generate the temperature scaled distribution  $q_0(\mathbf{x}) \propto p_0(\mathbf{x})^k$ .*

1372    *Proof.* We start by considering the following forward SDE:

1374    
$$d\mathbf{x} = f(t)\mathbf{x}dt + \frac{g(t)}{\sqrt{k}} d\mathbf{w} \quad (15)$$

1377    Let the initial distribution at  $t = 0$  be  $q_0(\mathbf{x})$ , we define the time-dependent distribution generated by  
 1378    this forward SDE as  $q_t(\mathbf{x})$ . Then, one corresponding reverse SDE that can sample  $q_0(\mathbf{x})$  takes the  
 1379    form of

1380    
$$d\mathbf{x} = [f(t)\mathbf{x} - (\frac{g(t)}{\sqrt{k}})^2 (\nabla \log q_t(\mathbf{x}))]dt + \frac{g(t)}{\sqrt{k}} d\bar{\mathbf{w}} \quad (16)$$

1382    Comparing Eq. 14 and Eq. 16, we can infer the following Lemma:

1384    **Lemma C.2.** *The CNS reverse-time SDE Eq. 14 and the SDE Eq. 16 are equivalent if and only if  
 1385     $\nabla \log q_t(\mathbf{x}) = k\mathbf{s}_\theta(\mathbf{x}, t)$  for all time  $t$ .*

1386    By construction, Eq. 16 evolves from  $q_T(\mathbf{x})$  to  $q_0(\mathbf{x})$ . Now we assume CNS (Eq. 14) starts from  
 1387    the same prior distribution  $q_T(\mathbf{x}) = \mathcal{N}(0, \frac{1}{k}\mathbf{I})$ , by Lemma C.2, CNS correctly perform temperature  
 1388    scaling and sample from  $q_0(\mathbf{x})$  if and only if  $\nabla \log q_t(\mathbf{x}) = k\mathbf{s}_\theta(\mathbf{x}, t)$ . Now we show that this  
 1389    condition is not true in general.

1391    **Left Side:** To compute  $q_t(\mathbf{x})$ , we need to solve the SDE in Eq. 15. For an initial condition  $x = X_0$ ,  
 1392    the solution  $X(t)$  is given by the following stochastic interpolant:

1393    
$$X(t) = \alpha_q(t)X_0 + \sigma_q(t)\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (17)$$

1395    
$$\alpha_q(t) = \int_0^t f(s)ds = \alpha_t$$

1396    
$$\sigma_q(t) = \int_0^t \frac{g(s)^2}{k} \exp(-2 \int_0^s f(u)du)ds = \frac{\sigma_t}{k}$$

1401    Therefore, we can compute the  $q_t(\mathbf{x})$  by

1403    
$$q_t(\mathbf{x}) = \int q_0(\mathbf{y})\mathcal{N}(\mathbf{x}; \alpha_t\mathbf{y}, \frac{\sigma_t^2}{k}\mathbf{I})d\mathbf{y} \quad (18)$$

1404      **Right Side.** For the original diffusion process without scaling, we can compute the noisy distribution  
 1405       $p_t(\mathbf{x})$  at time  $t$  as

$$1406 \quad p_t(\mathbf{x}) = \int p_0(\mathbf{y}) \mathcal{N}(\mathbf{x}; \alpha_t \mathbf{y}, \sigma_t^2 \mathbf{I}) d\mathbf{y} \quad (19)$$

1409      Comparing Eq. 18 and Eq. 19, we can infer that  $\nabla \log q_t(\mathbf{x}) \neq k \mathbf{s}_\theta(\mathbf{x}, t)$  for general distribution.  
 1410      One simple counterexample is where  $p_0(\mathbf{x})$  is a mixture of Gaussians. By previous reasoning, CNS  
 1411      cannot generate  $q_0(\mathbf{x})$  if the prior distribution is  $q_T(\mathbf{x})$ .

1412      What if we allow initial samples drawn from distributions other than  $q_T(\mathbf{x})$ ? We consider the special  
 1413      case where  $p_0(\mathbf{x}) = \mathcal{N}(0, \mathbf{I})$ , then  $p_t(\mathbf{x}) = p_0(\mathbf{x})$ ,  $q_t(\mathbf{x}) = q_0(\mathbf{x})$ . The condition  $\nabla \log q_t(\mathbf{x}) =$   
 1414       $k \mathbf{s}_\theta(\mathbf{x}, t)$  trivially holds true. By Lemma C.2, CNS(Eq. 14) and Eq. 16 are equivalent. Therefore,  
 1415      CNS can generate  $q_0(\mathbf{x})$  if and only if the prior distribution at time  $T$  is the same as  $q_T(\mathbf{x})$ . For any  
 1416      other prior distribution, CNS would not be able to generate  $q_0(\mathbf{x})$ .

1417      In conclusion, there does not exist an prior distribution  $q'_T(\mathbf{x})$ , from which CNS can always generate  
 1418      the temperature scaled distribution  $q_0(\mathbf{x})$

1419       $\square$

## 1420      D THE USE OF LARGE LANGUAGE MODELS (LLMs)

1421      We utilize LLMs to aid and refine some of the writing in the paper, such as correcting potential  
 1422      grammatical errors and suggesting more suitable expressions based on our original writing in some  
 1423      paragraphs.

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