DEBIASING GUIDANCE FOR DISCRETE DIFFUSION WITH SEQUENTIAL MONTE CARLO

Anonymous authors
Paper under double-blind review

ABSTRACT

Discrete diffusion models are a class of generative models that produce samples from an approximated data distribution within a discrete state space. Often, there is a need to target specific regions of the data distribution. Current guidance methods aim to sample from a distribution with mass proportional to $p_0(x_0)p(\zeta|x_0)^\alpha$ but fail to achieve this in practice. We introduce a Sequential Monte Carlo algorithm that generates unbiasedly from this target distribution, utilising the learnt unconditional and guided process. We validate our approach on low-dimensional distributions, controlled images and text generations. Our method provides strong control for text generation while maintaining low perplexity compared to guidance-based approaches.

1 INTRODUCTION

Discrete Diffusion models generate approximate samples from a data distribution $p_0(x_0)$ by gradually evolving samples from a simple base distribution $p_1(x_1)$ through a Continuous-Time Markov Chain (CTMC) (Shi et al., 2024; Campbell et al., 2022; Lou et al., 2024). While these models learn unconditional distributions, practical applications such as graph generation (Vignac et al., 2023), protein co-design (Campbell et al., 2024) or text generation (Lou et al., 2024) require controlled generation. For a conditioning variable ζ and temperature parameter α, we aim to sample from a tempered conditional distribution proportional to $p_0(x_0)p(\zeta|x_0)^\alpha$. When $\alpha = 1$, this recovers the conditional $p_0(x_0|\zeta)$. Higher values ($\alpha > 1$) biases sampling toward the conditioning signal ζ, while lower values ($\alpha < 1$) promotes diversity. To sample from this tempered distribution, guidance methods modify the unconditional diffusion transition rates (Nisonoff et al., 2024). However, guidance fails to sample from the intended distribution (Chidambaram et al., 2024; Bradley & Nakkiran, 2024), producing corrupted samples at high α values. Recent works propose finetuning approaches using reinforcement learning to sample from tempered distributions (Domingo-Enrich et al., 2025; Venkatraman et al., 2024; Fan et al., 2023; Black et al., 2024; Clark et al., 2024; Uehara et al., 2025). In this work, we address the theoretical and practical challenges of sampling from tempered distributions in discrete diffusion models. This work makes the following contributions:

1. We derive the exact transition rates required to sample from the tempered distribution with unnormalised mass $p_0(x_0)p(\zeta|x_0)^\alpha$ in discrete diffusion.
2. We propose an algorithm based on Sequential Monte Carlo, that asymptotically samples from the intended tempered distribution by exploiting previously forgone properties of the guided transition rate matrix, without any additional learning.
3. We validate our approach through low-dimensional experiments and demonstrate its effectiveness on controlled image and text generation, achieving strong conditional control with low perplexity on text dataset.

2 BACKGROUND WORK

2.1 CONTINUOUS TIME MARKOV CHAINS

A Continuous Time Markov Chain (CTMC) $\{X_t\}_{t\in[0,1]}$ is a Markov process on a finite state space $\mathcal{X} = \{1, \ldots, S\}$. The evolution of X_t is governed by a time-dependent rate matrix $R_t : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, which defines infinitesimal transition probabilities $p_{t+\Delta t|t}(x_{t+\Delta t}|x_t)$ given by

$$\delta_{x_{t+\Delta t},x_t} + R_t(x_t, x_{t+\Delta t})\Delta t + o(\Delta t), \tag{1}$$

where $\delta_{a,b}$ is 1 if $a = b$ and 0 otherwise.

The probability mass p_t of X_t evolves according to the Kolmogorov Forward Equation (KFE):

$$\partial_t p_t(x) = \underbrace{\sum_{y\neq x} p_t(y)R_t(y,x)}_{\text{incoming mass}} - \underbrace{\sum_{y\neq x} p_t(x)R_t(x,y)}_{\text{outgoing mass}}, \tag{2}$$

which can be written in vector form (Eq. (3)), where R_t satisfies mass conservation (Eq. (4)). For simplicity we define $R_t(x) := R_t(x,x)$.

$$\partial_t p_t = R_t^{\intercal} p_t, \tag{3}$$

$$R_t(x,x) = -\sum R_t(x,y) \text{ for } y \neq x \tag{4}$$

To obtain approximate samples $\tilde{X}_t$, one may use the Euler sampling algorithm (Campbell et al., 2024), initialising $\tilde{X}_0 \sim p_0$ and updating subsequent samples in intervals of Δt following:

$$\tilde{X}_{t+\Delta t} \sim \tilde{p}^{\text{Euler}}_{t+\Delta t|t}(\cdot|\tilde{x}_t), \qquad \text{with} \qquad \tilde{p}^{\text{Euler}}_{t+\Delta t|t}(\tilde{x}_{t+\Delta t}|\tilde{x}_t) \propto \delta_{\tilde{x}_{t+\Delta t},\tilde{x}_t} + R_t(\tilde{x}_t, \tilde{x}_{t+\Delta t})\Delta t.$$

A more comprehensive overview including the treatment of reverse-time CTMC can be found in Appendix B.

2.2 Discrete Diffusion Model

Discrete diffusion models are generative models that rely on a pair of stochastic processes formulated as Continuous Time Markov Chains on finite state spaces $\mathcal{X}$ with $|\mathcal{X}| = S$: a forward process corrupting the data distribution p_0 into a base distribution p_1, and a learned reverse process reconstructing from p_0 to p_1. Campbell et al. (2022) learn the reverse rate matrix through likelihood objectives, while Lou et al. (2024) instead learn the reverse rate matrix using a score entropy objective. In their work, Lou et al. (2024) propose the masking diffusion model, which extends the state space $\mathcal{X}$ with a masking state $\mathbb{m} = S + 1$ to form an extended state space $\overline{\mathcal{X}} = \{1, \dots, S, \mathbb{m}\}$. At time t, the masking process sends the current state to the masking state $\mathbb{m}$ according to a time-dependent noise schedule $\sigma : [0, 1] \to \mathbb{R}^+$, with transition probability $p^{\text{mask}}_{t|0}$:

$$p^{\text{mask}}_{t|0}(x_t|x_0) = \delta_{x_t,x_0} e^{-\sigma(t)} + \delta_{x_t,\mathbb{m}}(1 - e^{-\sigma(t)}) \tag{5}$$

For sufficiently large $\sigma(1)$, the masking process mixes to a point mass at $\mathbb{m}$ at time $t = 1$. To sample from a masking diffusion model, we start at the masking state $\mathbb{m}$ at $t = 1$ and apply Euler sampling in reverse time with learned rate matrix $\tilde{R}_t$. For high-dimensional data like text, the process extends to $\mathcal{X}^d$ by corrupting each dimension independently (Shi et al., 2024).

3 Guidance

Let $\{X_t\}_{t\in[0,1]}$ be a reverse-time CTMC with probability mass p_t and rate matrix R_t. For a conditioning variable ζ, we write $p_t(\zeta|x_t)$ and $p_t(x_t|\zeta)$ as the conditioned probability desnity of ζ given $X_t = x_t$ and the conditioned probability mass of x_t given ζ respectively. With guidance scale α, we aim to sample from the tempered distribution $p_0(x_0)p_0(\zeta|x_0)^\alpha/\mathcal{Z}_\alpha$, where $\mathcal{Z}_\alpha$ is a normalising constant. When $\alpha = 1$, this recovers $p_0(x_0|\zeta)$. Higher values ($\alpha > 1$) bias sampling toward high likelihood regions of ζ, while lower values promote diversity. Guidance modifies R_t to R_t^α (Definition 3.1) under the premise that it samples correctly from the intended tempered distribution. However, guidance fails to do so when $\alpha \neq 1$ as shown in Figure 1.

Definition 3.1. The guided rate matrix R_t^α is defined as,

$$\forall x \neq y.\ R_t^\alpha(x,y|\zeta) = R_t(x,y)\left[\frac{p_t(\zeta|y)}{p_t(\zeta|x)}\right]^\alpha \tag{6}$$

$$\forall x.\ R_t^\alpha(x,x|\zeta) = -R_t^\alpha(x|\zeta) \tag{7}$$

We defer the learning of guided rate matrix to Section 6.1. One might expect the guided rate matrix R_t^α to represent the time-reversal corrupting process starting at the tempered distribution $p_0(x_0)p_0(\zeta|x_0)^\alpha/\mathcal{Z}_\alpha$. However, as shown in Proposition 3.2, this is not the case.

Proposition 3.2. *Let $M_t[\cdot]$ denote the evolution of probability mass at time t under the corrupting process $p_{t|0}^{corrupt}$ of X_t, the unconditional diffusion model. Define $p_t^{\alpha,true}$ as:*

$$p_t^{\alpha,true} = M_t\big[p_0(\cdot)\, p_0(\zeta|\cdot)^\alpha/\mathcal{Z}_\alpha\big].$$

Then,

$$p_t^{\alpha,true}(x_t) \propto p_t(x_t)\mathbb{E}\big[p_0(\zeta|X_0)^\alpha|X_t = x_t\big].$$

The true tempered rate matrix $R_t^{\alpha,true}$ given by

$$\forall x \neq y,\ R_t^{\alpha,true}(x,y|\zeta) = R_t(x,y)\tfrac{\mathbb{E}[p_0(\zeta|X_0)^\alpha|X_t=y]}{\mathbb{E}[p_0(\zeta|X_0)^\alpha|X_t=x]},$$
$$\forall x,\ R_t^{\alpha,true}(x,x|\zeta) = -\sum_{y\neq x} R_t^{\alpha,true}(x,y|\zeta),$$

satisfies the reverse-time Kolmogorov Forward Equation:

$$\partial_t p_t^{\alpha,true} = -\left[R_t^{\alpha,true}\right]^\intercal p_t^{\alpha,true}.$$

From Proposition 3.2, it follows that $R_t^\alpha = R_t^{\alpha,\text{true}}$ holds only for $\alpha = 1$. Furthermore, for $\alpha = 1$, guidance requires that the base distribution p_1 is independent of the conditioning variable ζ to correctly sample from the conditioned distribution $p_0(X_0|\zeta)$, as shown in Corollary 3.3.

Corollary 3.3. *If $p_1(x_1|\zeta) = p_1(x_1)$, then guidance samples correctly from the conditioned distribution $p_0(x_0|\zeta)$.*

While (Nisonoff et al., 2024) presents similar results to Corollary 3.3, our lemma establishes that the independence condition is necessary.
Learning $R_t^{\alpha,true}$ most often requires expensive simulation-based objectives to estimate the gradients of a reverse path-wise KL divergence (Domingo-Enrich et al., 2025; Denker et al., 2024; Uehara et al., 2025). Therefore, we focus instead on sampling from $p_t^\alpha(x_t) \propto p_t(x)p_t(\zeta|x_t)^\alpha$.

4 Debiasing Guidance with Sequential Monte Carlo

Our objective is to sample from the tempered distribution $\frac{1}{\mathcal{Z}_0^\alpha}p_0(x_0)p_0(\zeta|x_0)^\alpha$. To that end, we first introduce an importance sampling method, which we later leverage through resampling.

4.1 Importance Sampling the Tempered Distribution

We consider the importance sampling method to compute the expectation of a function h under the distribution of probability mass $p_t^\alpha(x_t) = \frac{1}{\mathcal{Z}_t^\alpha}p_t(x_t)p_t(\zeta|x_t)^\alpha$ where $\mathcal{Z}_t^\alpha$ is the normalising constant. We write the expectation as:

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{1}{\mathcal{Z}_t^\alpha}\sum_{x\in\mathcal{X}} h(x)p_t(x)p_t(\zeta|x_t)^\alpha, \tag{8}$$

For importance sampling, we consider a proposal in the form of a reverse time CTMC $\{Y_t\}_{t\in[0,1]}$, with rate matrix Q_t and initial distribution p_1. We construct unnormalised importance weight $\{W_t\}_{t\in[0,1]}$ such that:

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{\mathbb{E}[W_t \cdot h(Y_t)]}{\mathbb{E}[W_t]} \tag{9}$$

where $\mathbb{E}$ is an expectation over the joint distribution $\text{Law}(W_t, Y_t)$.

Proposition 4.1 provides one form of the unnormalised importance weights $\{W_t\}_{t\in[0,1]}$.

Proposition 4.1. *Let proposal $\{Y_t\}_{t\in[0,1]}$ be a reverse-time CTMC with rate matrix Q_t and initial distribution $p_1(\cdot)$. Define $\{W_t\}_{t\in[0,1]}$ by*

$$W_t = \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t}\left(\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}}\right)^{\alpha},$$

where we have path space measures $\mathbb{P}_t^{base} = \mathrm{Law}\{X_\tau\}_{\tau\in[t,1]}$, $\mathbb{P}_t^{\alpha=1} = \mathrm{Law}\{X_\tau \mid \zeta\}_{\tau\in[t,1]}$ and $\mathbb{Q}_t = \mathrm{Law}\{Y_\tau\}_{\tau\in[t,1]}$. Then the Radon-Nikodym derivatives can be written as,

$$\ln\frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t} = \sum_{\substack{\tau\geq t\\ Y_{\tau^+}\neq Y_\tau}} \ln R_t(Y_{\tau^+}, Y_\tau) + \int_1^t Q_\tau(Y_t) - R_\tau(Y_t)\,\mathrm{d}\tau \qquad \ln\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}} = \sum_{\substack{\tau\geq t\\ Y_{\tau^+}\neq Y_\tau}} \ln R_t^{\alpha=1}(Y_{\tau^+}, Y_\tau|\zeta) + \int_1^t R_\tau(Y_\tau) - R_\tau^{\alpha=1}(Y_\tau|\zeta)\mathrm{d}\tau,$$

such that for any p_t^α-integrable function h and time $t\in[0,1]$, we have

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{\mathbb{E}[W_t\cdot h(Y_t)]}{\mathbb{E}[W_t]},$$

where the expectation is taken over Law(W_t, Y_t)

The proof is in Appendix D. The result extends to stochastic differential equations (Appendix D.1).

For practical implementation, we approximate the CTMC by discretising time into T steps. For times $t < s$, we denote the transitions derived from rate matrices Q_t, R_t, and $R_t^{\alpha=1}$ as $q_{t|s}(x_t|x_s)$, $p_{t|s}(x_t|x_s)$, and $p_{t|s}(x_t|x_s,\zeta)$ respectively. To this end, we consider a discretisation of Proposition 4.1 in Appendix E. We also present Algorithm 1 in Section 4.2, a pseudo-code of the discretised version of the proposed sampling method.

4.2 Resampling and Sequential Monte Carlo

To approximate samples from the tempered distribution p_0^α, we leverage resampling. Given a set of K weights and samples $(w_t^{(i)}, y_t^{(i)}) \sim \mathrm{Law}(W_t, Y_t)$, resampling yields approximate samples from p_t^α by sampling the categorical distribution,

$$\mathrm{Cat}\left(\left\{\frac{w_t^{(i)}}{\sum_{j=1}^K w_t^{(j)}}\right\}_{i=1}^K, y_t^{(i)}\right), \tag{10}$$

where $y_t^{(i)}$ is sampled with probability $\frac{w_t^{(i)}}{\sum_{j=1}^K w_t^{(j)}}$.

This suggests an algorithm to obtain samples from the tempered distribution $p_0(x_0)p_0(\zeta|x_0)^\alpha$ by resampling at time $t = 0$. However, this approach suffers from a significant challenge: weight degeneracy. As t approaches 0, the variance of weights W_t increases dramatically, requiring an impractical number of samples. Sequential Monte Carlo (SMC) addresses this challenge by performing resampling steps at intermediate times $0 < t < 1$, obtaining approximate samples of p_t^α while resetting the weights to maintain the importance sampling equality. Effectively, this eliminates samples from low-likelihood regions and duplicates those in high-likelihood regions. The resampling step is flexible in both algorithm choice and timing. In Section 6.3, we explore partial resampling where only a subset of samples participate. Resampling is typically triggered when the effective sample size (ESS) falls below a threshold, where ESS is defined as:

$$\mathrm{ESS}_t^K = \left(\sum_{i=1}^K w_t^{(i)}\right)^2 \Big/ \sum_{i=1}^K (w_t^{(i)})^2 \in [1, K] \tag{11}$$

We emphasize that the resampling procedure allows us to leverage importance sampling to generate asymptotically unbiased samples rather than purely computing expectations. The full algorithm to sample from the tempered distribution can be found in Algorithm 1 detailed in Algo. 2 (App. F) for our experiments.

Algorithm 1 Main Algorithm

Require: Number of particles K; Proposal Rate Matrix Proposal Transition Q_t; Unconditional rate matrix R_t; Time grid $\{t_l\}_{l=1}^T$ for potential resampling; ESS threshold ESS_THRESHOLD; Resampling algorithm `resample`.

Initialisation:
Sample $\{y_1^{(i)}\}_{i=1}^K$ i.i.d. from p_1.
Set $\hat{w}_1^{(i)} \leftarrow 1$ for $i = 1, \ldots, K$.
for $l = 1$ to T **do**
  for $i = 1$ to K **do**
    ▷ Step 1: Evolve samples and weights
    Obtain $(y_{l+1}^{(i)}, \hat{w}_{l+1}^{(i)})$ by evolving joint system (Y_t, W_t) from $(y_l^{(i)}, \hat{w}_l^{(i)})$ at time t_l
  end for
  ▷ Step 3: Resample; see Appendix G for other resampling algorithms
  Set ESS $\leftarrow \left(\sum_{i=1}^K \hat{w}_l^{(i)}\right)^2 / \sum_{i=1}^K (\hat{w}_l^{(i)})^2$
  if ESS $\leq$ ESS_THRESHOLD **then**
    Set $y_{l+1}^{(i)}, \hat{w}_{l+1}^{(i)} \leftarrow \texttt{resample}(\{y_l^{(i)}\}, \{\hat{w}_l^{(i)}\})$
  end if
end for
Output:
Particles $\{x_T^{(i)}\}_{i=1}^K$

D	$\lvert\mathcal{V}\rvert$	KL($\pi \| \sigma_{\text{SMC}}$)	KL($\pi \| \sigma_{\text{Guided}}$)
1	50	0.002 ± 0.001	0.004 ± 0.003
2	3	0.035 ± 0.179	0.074 ± 0.082
2	10	0.013 ± 0.012	0.081 ± 0.052
2	100	1.725 ± 0.743	6.690 ± 1.194
2	200	5.948 ± 1.042	12.906 ± 0.542

Table 1: Forward KL divergence between the target tempered distribution π and the empirical distribution of samples from the guided process and SMC algorithm. Method with a lower KL divergence is highlighted.

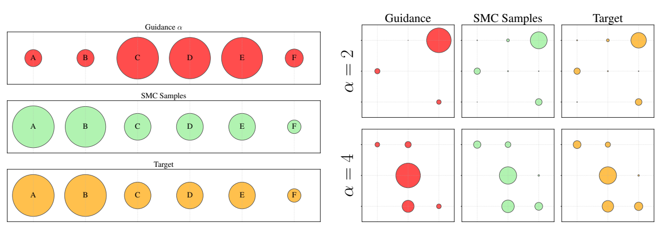

Figure 1: Visualization of 1D (left) and 2D (right) discrete distributions. Orange circles: target distribution; green: SMC approximation; red: standard guidance. Left: six states (A-F) with probability mass shown by disk size. Right: nine states at $\alpha = 2$ (top) and $\alpha = 4$ (bottom).

5 Numerical Verification

We validate our algorithm's ability to sample from the tempered distribution $p_0(x_0)p_0(\zeta|x_0)^\alpha/\mathcal{Z}_\alpha$ using low-dimensional examples. These experiments use explicitly specified probability mass $p_0(x_0)$ and conditional likelihood $p_0(\zeta|x_0)$, enabling direct computation of the guided ratio matrix. This setup verifies algorithmic correctness independent of learning effects and allows evaluation of the KL divergence between sampled and target distributions.

Experimental setup We evaluate on state spaces of dimension one ($\mathcal{V}$) and two ($\mathcal{V}^2$), where $\mathcal{V}$ is a finite vocabulary. Our proposal uses per-dimension Euler sampling (Campbell et al., 2024) with 100 discretization steps and 50,000 samples. For each dimension, we generate 30 different tempered target distributions $\pi \propto p_0(x_0)p_0(\zeta|x_0)^\alpha$ by combining different choices of $p_0(x_0)$, $p_0(\zeta|x_0)$, and temperature α. We sample from each target using both guidance and SMC, then compute KL divergences between the true tempered distribution and the empirical distributions σ^{SMC} and σ_{Guided}. Results in Table 1 show that guided processes significantly deviate from targets while SMC achieves approximate sampling, demonstrating superior sampling across all combinations. Qualitative examples are shown in Figure 1.

6 EXPERIMENTS

Section 5 demonstrated SMC's ability to sample tempered distributions with sufficient particles and discretisation steps. In practice, for image and text generation tasks, computing rate matrices is computationally intensive, limiting feasible particle counts and step sizes. For these practical settings, we use the guided rate matrix as our proposal in SMC, with guidance temperature β distinct from the SMC temperature α, and provide guidance on resampling strategies.

6.1 LEARNING THE GUIDANCE TERM

We consider two approaches to obtain the guided rate matrix R_t^α, **DEFT** and **CFG**.

DEFT For discrete diffusion models, we extend Doob's h-transform Efficient FineTuning (DEFT, Denker et al. (2024)) by showing that the conditional rate matrix $R_t^{\alpha=1}$ decomposes into an unconditional rate matrix R_t and a guidance term $p_t(\zeta|y)/p_t(\zeta|x)$: $R_t^{\alpha=1}(x,y) = R_t(x,y) \times p_t(\zeta|y)/p_t(\zeta|x)$. We model the guidance term with a neural network g and parameterize the conditional rate matrix as $R_t^{\alpha=1}(x,y|\zeta) = R_t(x,y)g(x,y,t)$ while keeping the unconditional rate parameters fixed. We use this guidance term in both pixel-level image generation and controlled text generation.

CFG In text generation, we extend classifier-free guidance (Ho & Salimans, 2022a). We obtain the conditional rate matrix $R_t^{\alpha=1}(x,y|\zeta)$ by finetuning the learned unconditional rate matrix $R_t(x,y|\zeta)$, concatenating conditions to the inputs. For $\alpha > 1$, the guided rate matrix follows $[p_t(\zeta|y)/p_t(\zeta|x)]^\alpha = \left[R_t^{\alpha=1}(x,y|\zeta)/R_t(x,y)\right]^\alpha$

There are other ways of finding a guided rate matrix (Nisonoff et al., 2024) Vignac et al. (2023) considers a first-order Taylor approximation of the guidance term; Kerby & Moon (2024); Li et al. (2024) consider a training-free approach by Monte Carlo estimates of the guidance term.

We present results using the CFG approach for text experiments and the DEFT approach for image experiments. For DEFT results on text, see Appendix J. We begin by discussing partial resampling, a key component of our SMC implementation, followed by a detailed account of our experiments and findings.

6.2 PARTIAL RESAMPLING

SMC methods suffer from mode collapse in high dimensions when using few particles. In this scenario, most particle weights decay rapidly, leaving only a small subset with significant weights. During resampling, particles with low weights are discarded, often resulting in only one or two unique particles. Figure 3 (middle) demonstrates this effect: a batch of 16 particles collapses into identical images. We adopt a partial resampling scheme which effectively mitigates mode collapse by resampling only a subset of particles. We select the K most and least weighted samples for resampling, which preserves sample diversity and maintains unbiased sampling from the target distribution as shown in Figure 3 (right). We detail partial resampling algorithms in Appendix G. In our experiments, we use Algorithm 4, set $K = \lfloor N/4 \rfloor$, resampling half of all generated samples.

6.3 CLASS-CONDITIONAL PIXEL-LEVEL IMAGE GENERATION

We evaluate our method on MNIST class-conditional generation using 28×28 grayscale images with pixel values in $\{0, ..., 255\}$. We first detail the experimental setup; we then study partial resampling for mitigating mode collapse in SMC and analyze its impact on generation control.

Dataset and Training Details We train an unconditional SEDD model (Lou et al., 2024) parameterized by a U-Net (Ronneberger et al., 2015) for 20k steps. Using 20% of MNIST, we finetune a guidance network for 4k steps with DEFT using the SEDD denoising score entropy loss. The guidance network is a U-Net with learned condition embeddings added to each layer, at 50% parameters of the unconditional network.

Does SMC improve conditional control? For class-conditioned generation, we evaluate sample consistency with target conditions using a pre-trained digit classifier. Figure 2 shows accuracy for various SMC temperatures α and number of particle N, using $R_t^{\alpha=1}$ as the proposal. We compute accuracies within each SMC run and confidence intervals across 10 independent runs. Higher SMC temperatures yield better accuracy across all particle counts. Partial resampling achieves higher accuracy with lower variance, which indicates that it effectively reduces mode collapse. We further investigate mode collapse for MNIST in Appendix H

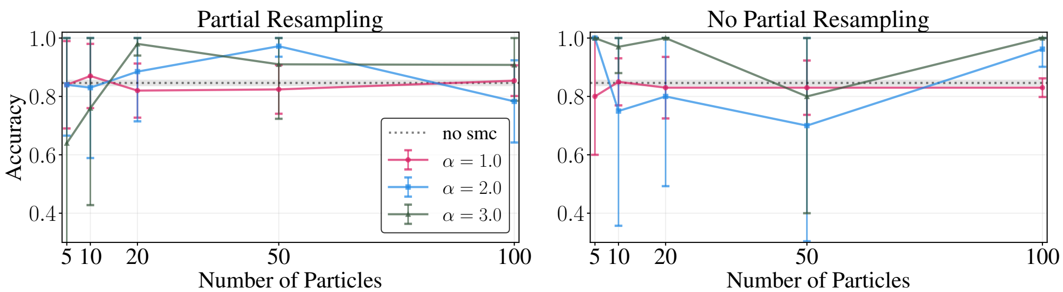

Figure 2: **Accuracy on Class-Conditioned Image generation.** SMC classification accuracy versus particle count for different twist temperatures (α), with and without partial resampling. Partial resampling and higher α achieve better peak accuracy. Dotted lines show baseline accuracy without SMC. Values averaged over 50 runs.

Experiments on MNIST demonstrate that our method successfully addresses two key challenges in conditional generation. Partial resampling effectively mitigates mode collapse while maintaining sample diversity, and higher SMC temperatures improve conditional control across all particle counts. These results validate our theoretical framework for a standard image generation task.

6.4 Controlled Text Generation

We validate the approach on three controlled text generation: Sentiment Controlled Generation, Toxicity Controlled Generation, and Text Infilling. These high-dimensional experiments demonstrate the algorithm's practicality.

6.4.1 Dataset and Training Details

In all experiments, we use the pre-trained 320M non-embedding parameters `sedd-medium`, a Score Entropy Discrete Diffusion (SEDD) model (Lou et al., 2024). When guiding with CFG we finetune `sedd-medium` and when guiding with DEFT we finetune `sedd-small` for 2k steps to serve as the guidance term $p_t(\zeta|y)/p_t(\zeta|x)$ for the large `sedd-medium` model, on task-specific datasets:

Toxicity Controlled Generation Following (Liu et al., 2021), we use human-annotated comments from the Jigsaw Unintended Bias in Toxicity Classification challenge[1], with texts labelled toxic when 50% or more annotators mark them as such. A prompt `The generated text is toxic` is concatenated to the transformer input as the condition.
Sentiment Controlled Generation. We follow the setup in (Amini et al., 2024), using the Stanford Sentiment Treebank, a dataset of movie reviews, as the finetuning dataset. A prompt *"The generated text is of positive sentiment"* or *"The generated text is of negative sentiment"* is concatenated to the transformer input as the condition.
Text Infilling. We finetune on 10% of the OpenWebText training dataset. For each sentence, randomly sampled indices (`START`, `END`) define the token range to mask. The masked sentence ζ is concatenated to the model input throughout generation. The entire partial masked sentence `[PREFIX] [MASK] [SUFFIX]` is passed to the network as the condition

We evaluate samples using fluency and conditional control metrics, averaging across 5 runs and 50 particles for SMC or 250 particles for guidance, with 100 discretization steps and partial resampling. For fluency, we compute generative perplexity (Liu et al., 2021) using `GPT2-XL` as the reference model. For sentiment and toxicity tasks, we use a LoRA-finetuned `GPT2-XL` to account for stylistic differences. For conditional metrics, we use a pre-trained classifier for sentiment accuracy, the Perspective API for toxicity scores, and measure the success rate of generating unmasked tokens.

[1] `https://bit.ly/3cvG5py`

6.4.2 Ablating Guidance and SMC temperature

We study how guidance temperature β and SMC temperature α impact both conditional control and fluency metrics across toxicity, sentiment, and infilling tasks.

Our findings: SMC-based methods improve considerably on conditional control metrics for low-to-mid guidance temperature β and with high SMC temperature α. For high guidance temperature, the guidance baseline improves conditional metrics but at the cost of higher perplexity. This implies that guidance-generated samples drift from the distribution of the fine-tuning dataset. Hence, we conclude that the optimal approach is to combine SMC with high SMC temperature and mid-to-low guidance temperature (see Appendix I for more results).

Method	Temperature		Sentiment Control		Toxicity Control		Text Infilling			
	α	β	Perplexity	Accuracy ↑	Perplexity	Toxicity ↑	Perplexity	BERT ↑	GLEU ↑	Accuracy ↑
SMC	1.0	1.0	84.410±8.647	0.839±0.047	66.558±9.085	0.451±0.244	65.048±30.951	0.583±0.038	0.701±0.021	0.992±0.015
		1.2	93.225±7.281	0.840±0.065	58.913±1.945	0.384±0.231	-	-	-	-
SMC	2.0	1.0	80.500±3.554	0.882±0.048	62.192±6.998	0.577±0.232	67.062±42.959	0.584±0.032	0.703±0.022	0.996±0.014
		1.2	83.767±7.457	0.877±0.050	88.650±3.325	0.547±0.222	-	-	-	-
SMC	4.0	1.0	77.538±11.209	0.875±0.048	89.125±4.062	0.595±0.231	73.562±26.980	0.587±0.029	0.708±0.019	0.998±0.013
		1.2	85.170±7.253	0.872±0.057	109.200±3.518	0.611±0.208	-	-	-	-
SMC	8.0	1.0	80.283±6.825	0.868±0.059	77.000±3.420	0.561±0.252	80.583±45.478	0.576±0.038	0.699±0.022	0.998±0.015
		1.2	78.388±6.606	0.875±0.060	88.800±1.549	0.558±0.230	-	-	-	-
Guidance	-	1.0	80.283±6.825	0.837±0.056	67.475±8.467	0.455±0.244	63.915±31.375	0.583±0.037	0.701±0.021	0.991±0.018
		1.2	78.388±6.606	0.837±0.006	70.525±10.079	0.505±0.246	-	-	-	-
Reconstruction	-	-	44.673±11.687	0.548±0.105	44.319±13.885	0.130±0.123	53.697±26.543	0.595±0.038	0.706±0.021	1.000±0.000

Table 2: Performance comparison of Sequential Monte Carlo, Guidance, and Reconstruction across three tasks: sentiment control, toxicity control, and text infilling. Results show mean ± standard deviation. α is the SMC temperature and β the guidance temperature. Light blue marks the best overall result per metric, while light purple marks the best result between Guidance and SMC for text infilling. Text infilling results (grey columns) have been run for 1000 steps.

6.4.3 Scaling up Discretisation Step

We further investigate sample qualities by increasing discretization steps to 1000, effectively allocating 10x compute. For text infilling, we employ two additional metrics: BERTScore (Zhang* et al., 2020) with the DeBERTa model (He et al., 2021), and GLEU-4 (Wu et al., 2016). For sentiment and toxicity-controlled generation, we compare against a baseline prompting method, while for infilling, we evaluate against a reconstruction method that overwrites the sampling trajectory with the partially masked sentence. Table 2 summarises our results. SMC with a high SMC temperature significantly outperforms the base guidance method across all tasks. Notably, for toxicity and sentiment-controlled generation our SMC-based method significantly outperforms a simple prompting strategy in generating strong conditions. However, for infilling and despite their theoretical guarantees, guidance and SMC approaches underperform compared to the reconstruction method, even though it is known to be biased (Wu et al., 2024). This indicates an inadequate learning of the guided rate matrix. While SMC can strengthen the base guidance method, it cannot fully compensate for an imperfectly learned guided rate matrix.

7 Conclusion and Future Work

This work addresses a fundamental challenge in diffusion models: accurate sampling from tempered distribution. We make three key contributions. First, we derive the exact transition rate to sample from the tempered distribution; second, we propose an SMC-based algorithm that asymptotically samples from the tempered distribution; third, we demonstrate its practicality in controlled generation. Our empirical results validate the practicality of our theoretically correct algorithm. In low dimensions, it significantly outperforms standard guidance methods. In high dimensions, our method achieves superior control while maintaining sample quality. However, our experiments also reveal that the effectiveness depends on the quality of the learned guided rate matrix.

References

Michael S. Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants, 2023. URL https://arxiv.org/abs/2209.15571.

Michael S Albergo and Eric Vanden-Eijnden. Nets: A non-equilibrium transport sampler. *arXiv preprint arXiv:2410.02711*, 2024.

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL https://arxiv.org/abs/2303.08797.

Afra Amini, Li Du, and Ryan Cotterell. Structured voronoi sampling, 2024. URL https://arxiv.org/abs/2306.03061.

Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning, 2024. URL https://arxiv.org/abs/2305.13301.

Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable video diffusion: Scaling latent video diffusion models to large datasets, 2023. URL https://arxiv.org/abs/2311.15127.

Arwen Bradley and Preetum Nakkiran. Classifier-free guidance is a predictor-corrector, 2024. URL https://arxiv.org/abs/2408.09000.

Andrew Campbell, Joe Benton, Valentin De Bortoli, Tom Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models, 2022. URL https://arxiv.org/abs/2205.14987.

Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi S. Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=kQwSbv0BR4.

Davide Carbone, Mengjian Hua, Simon Coste, and Eric Vanden-Eijnden. Efficient training of energy-based models using jarzynski equality. *Advances in Neural Information Processing Systems*, 36: 52583–52614, 2023.

Muthu Chidambaram, Khashayar Gatmiry, Sitan Chen, Holden Lee, and Jianfeng Lu. What does guidance do? a fine-grained analysis in a simple setting, 2024. URL https://arxiv.org/abs/2409.13074.

Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly fine-tuning diffusion models on differentiable rewards, 2024. URL https://arxiv.org/abs/2309.17400.

Francois R J Cornet, Grigory Bartosh, Mikkel N. Schmidt, and Christian A. Naesseth. Equivariant neural diffusion for molecule generation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=40pE5pFhWl.

Pierre Del Moral and Spiridon Penev. *Stochastic processes: From applications to theory*. CRC Press, Taylor Francis Group, 2017.

Alexander Denker, Francisco Vargas, Shreyas Padhy, Kieran Didi, Simon V. Mathis, Vincent Dutordoir, Riccardo Barbano, Emile Mathieu, Urszula Julia Komorowska, and Pietro Lio. DEFT: efficient finetuning of conditional diffusion models by learning the generalised h-transform. *CoRR*, abs/2406.01781, 2024. doi: 10.48550/ARXIV.2406.01781. URL https://doi.org/10.48550/arXiv.2406.01781.

Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021. URL https://arxiv.org/abs/2105.05233.

Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control, 2025. URL https://arxiv.org/abs/2409.08861.

Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Dpok: Reinforcement learning for fine-tuning text-to-image diffusion models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 79858–79885. Curran Associates, Inc., 2023.

Ishaan Gulrajani and Tatsunori B. Hashimoto. Likelihood-based diffusion language models, 2023. URL https://arxiv.org/abs/2305.18619.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. Deberta: Decoding-enhanced bert with disentangled attention, 2021. URL https://arxiv.org/abs/2006.03654.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022a. URL https://arxiv.org/abs/2207.12598.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022b. URL https://arxiv.org/abs/2207.12598.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. URL https://arxiv.org/abs/2006.11239.

Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d, 2022. URL https://arxiv.org/abs/2203.17003.

Thomas J. Kerby and Kevin R. Moon. Training-free guidance for discrete diffusion models for molecular generation, 2024. URL https://arxiv.org/abs/2409.07359.

Christian Léonard. Some properties of path measures. *Séminaire de Probabilités XLVI*, pp. 207–230, 2014.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding, 2024. URL https://arxiv.org/abs/2408.08252.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023. URL https://arxiv.org/abs/2210.02747.

Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code, 2024. URL https://arxiv.org/abs/2412.06264.

Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. DExperts: Decoding-time controlled text generation with experts and anti-experts. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 6691–6706, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.522. URL https://aclanthology.org/2021.acl-long.522/.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution, 2024. URL https://openreview.net/forum?id=71mqtQdKB9.

L. Martino, V. Elvira, and F. Louzada. Weighting a resampled particle in sequential monte carlo. In *2016 IEEE Statistical Signal Processing Workshop (SSP)*, pp. 1–5, 2016. doi: 10.1109/SSP.2016.7551711.

Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance for discrete state-space diffusion and flow models, 2024. URL https://arxiv.org/abs/2406.01572.

Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis, 2023. URL https://arxiv.org/abs/2307.01952.

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

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. URL https://arxiv.org/abs/1505.04597.

Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized masked diffusion for discrete data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=xcqSOfHt4g.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution, 2020. URL https://arxiv.org/abs/1907.05600.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. URL https://arxiv.org/abs/2011.13456.

Masatoshi Uehara, Yulai Zhao, Chenyu Wang, Xiner Li, Aviv Regev, Sergey Levine, and Tommaso Biancalani. Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review, 2025. URL https://arxiv.org/abs/2501.09685.

Francisco Vargas, Will Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. *arXiv preprint arXiv:2302.13834*, 2023.

Francisco Vargas, Nikolas Nusken, Shreyas Padhy, and Denis Blessing. Transport meets variational inference: Controlled monte carlo diffusions. In *The Twelfth International Conference on Learning Representations: ICLR 2024*, 2024.

Siddarth Venkatraman, Moksh Jain, Luca Scimeca, Minsu Kim, Marcin Sendera, Mohsin Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, Alexandre Adam, Jarrid Rector-Brooks, Yoshua Bengio, Glen Berseth, and Nikolay Malkin. Amortizing intractable inference in diffusion models for vision, language, and control. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=gVTkMsaaGI.

Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=UaAD-Nu86WX.

Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

Luhuan Wu, Brian Trippe, Christian Naesseth, David Blei, and John P Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google's neural machine translation system: Bridging the gap between human and machine translation, 2016. URL https://arxiv.org/abs/1609.08144.

Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SkeHuCVFDr.

A RELATED WORK

Diffusion models have emerged as a powerful class of generative models that can be interpreted through various theoretical frameworks, including score-based modeling (Song & Ermon, 2020; Song et al., 2021), flow matching (Lipman et al., 2023; 2024), variational methods (Ho et al., 2020) and stochastic interpolant (Albergo & Vanden-Eijnden, 2023; Albergo et al., 2023). Their versatility has led to successful applications across diverse domains, from image and video generation (Rombach et al., 2022; Podell et al., 2023; Blattmann et al., 2023) to molecular design (Cornet et al., 2024; Vignac et al., 2023; Hoogeboom et al., 2022), protein design (Watson et al., 2023), and text generation (Shi et al., 2024; Lou et al., 2024; Gulrajani & Hashimoto, 2023), where controlled generation is often crucial.

Two primary approaches have been developed for controlled generation: classifier guidance (Dhariwal & Nichol, 2021) and classifier-free guidance (Ho & Salimans, 2022b). Classifier guidance incorporates gradients from a trained classifier to steer the generation process, while CFG eliminates the need for a separate classifier by jointly training conditional and unconditional models. For scenarios with limited data where training classifiers or conditional models is impractical, DEFT (Denker et al., 2024) demonstrates that fine-tuning a small network can learn classifier scores directly, enabling guidance of unconditional models. However, recent theoretical analyses (Chidambaram et al., 2024; Bradley & Nakkiran, 2024) have revealed fundamental limitations of guidance methods. Specifically, as the guidance temperature parameter increases to strengthen conditioning, these methods fail to sample from their intended tilted target distributions. This limitation is particularly relevant for applications requiring precise control over generated samples.

In the discrete setting, several approaches have been proposed to adapt guidance methods. Nisonoff et al. (2024) adapt Classifier Free Guidance, Vignac et al. (2023) use Taylor expansions to approximate guidance terms with a property-predicting regressor. Li et al. (2024) introduce SVDD, integrating value functions for reward-based sampling without fine-tuning. Kerby & Moon (2024) develop a training-free guidance framework for discrete data generation.

Recent works propose finetuning approaches using reinforcement learning to sample from tempered distributions (Domingo-Enrich et al., 2025; Venkatraman et al., 2024; Fan et al., 2023; Black et al., 2024; Clark et al., 2024; Uehara et al., 2025). Black et al. (2024) introduce Denoising Diffusion Policy Optimization (DDPO), which reformulates the denoising process as a multi-step decision problem. Clark et al. (2024) propose Direct Reward Fine-tuning (DRaFT), demonstrating that backpropagation through the entire sampling procedure can effectively optimize differentiable reward functions. Fan et al. (2023) develop DPOK, which combines policy optimization with KL regularization for fine-tuning text-to-image models, showing improvements in both image-text alignment and image quality compared to supervised fine-tuning approaches.

B A PRIMER ON CONTINUOUS-TIME MARKOV CHAINS

B.1 DEFINITION

A Continuous-Time Markov Chain $\{X_t\}_{t\in[0,1]}$ on finite state space $\mathcal{X} = \{1, \dots, |\mathcal{X}|\}$ is collection of time-indexed random variables taking values on $\mathcal{X}$. A CTMC $\{X_t\}_{t\in[0,1]}$ obeys the Markov property, that is, for all $A \subseteq \mathcal{X}$, $n \in \mathbb{N}$, $t_1, \dots, t_n \in [0,1]$, and $0 \le x_1 < \cdots < x_n \le 1 \in \mathcal{X}$:

$$\mathbb{P}(X_t \in A | X_{t_1} = x_1, \dots X_{t_n} = x_n) = \mathbb{P}(X_t \in A | X_{t_n} = x_n). \tag{12}$$

The initial distribution $p_0(x_0)$ and transition probability $p_{t|s}(x_t|x_s)$ are defined as,

$$p_0(x_0) = \mathbb{P}(X_0 = x_0) \qquad\qquad p_{t|s}(x_t|x_s) = \mathbb{P}(X_t = x_t, X_s = x_s). \tag{13}$$

Notably, a CTMC is **càdlàg** by convention, that is for every $t \ge 0$:

1. $\lim_{s\to t^+} X(s) = X(t)$ (right-continuous at t),
2. $\lim_{s\to t^-} X(s)$ exists (finite left-hand limit at t).

B.2 The Transition Rate Matrix

One way to characterise a CTMC is through the time-dependent **transition rate matrix** R_t (also known as the **generator**), defined as,

$$\forall x, y \in \mathcal{X}.\ R_t(x, y) = \lim_{\Delta t \to 0} \frac{p_{t+\Delta t|t}(x_{t+\Delta t}|x_t) - \delta_{x,y}}{\Delta t} \tag{14}$$

where $\delta_{a,b}$ is 1 if $a = b$ and 0 if otherwise.

By definition, a first-order approximation of the transition probability is:

$$p_{t+\Delta t|t}(x_{t+\Delta t|t}|x_t) = \delta_{x_{t+\Delta t}, x_t} + R_t(x_t, x_{t+\Delta t})\Delta t + o(\Delta t). \tag{15}$$

By noticing that $\sum_{y \in \mathcal{X}} p_{t+\Delta t|t}(y|x) = 1$, we derive the mass conservation property of R_t,

$$R_t(x, x) = -\sum_{y \in \mathcal{X}: x \neq x} R_t(x, y) \tag{16}$$

By convention, we define $R_t(x)$ as,

$$R_t(x) = \sum_{y \in \mathcal{X}: x \neq x} R_t(x, y). \tag{17}$$

B.3 The Kolmogorov Equations

In this section, we outline both Kolmogorov Forward and Backward Equation that govern respectively the change of the probability mass, and expectations of some function over time.

Kolmogorov-Forward Equation. The Kolmogorov Forward Equation (or the continuity equation) governs the change of probability mass over time. It states that the probability mass p_t satisfies the following equation:

$$\forall 0 \leq t \leq 1.\ \partial_t p_t(x_t) = \underbrace{\sum_{y \in \mathcal{X}: x_t \neq y} R_t(y, x_t) p_t(y)}_{\text{incoming probability mass}} - \underbrace{\sum_{y \in \mathcal{X}: x_t \neq y} R_t(x_t, y) p_t(x_t)}_{\text{outgoing probability mass}} \tag{18}$$

In other words, the rate of change of probability mass equals the difference between incoming and outgoing probability flows. In matrix notation, this equation is:

$$\partial_t p_t = R_t^{\intercal} p_t \tag{19}$$

Kolmogorov Backward Equation Given a function $h : \mathcal{X} \to \mathbb{R}$, the Kolmogorov Backward Equation states that the function $u_t(x)$ defined as,

$$u_t(x) = \mathbb{E}[h(X_1)|X_t = x], \tag{20}$$

satisfies the ordinary differential equation,

$$\partial_t u_t(x) = -\sum_y R_t(x, y) u_t(y), \tag{21}$$

writing u_t as a column vector, this can be written in matrix form,

$$\partial_t u_t = -R_t^{\intercal} u_t. \tag{22}$$

B.4 Reverse-time CTMC

Time-reversed CTMC is a core concept of diffusion model. A reverse-time CTMC $\{Y_t\}_{\in[0,1]}$ with time-dependent rate R_t is defined as $Y_t = \hat{Y}_{1-t}$ where $\{\hat{Y}_t\}_{t \in [0,1]}$ is a forward-time CTMC with transition rate R_{1-t}. The reverse-time Kolmogorov Equations are then obtained by applying a change of variable $t \mapsto 1 - t$.

C Proof for Section 3

In this section, we prove Proposition 3.2.

Proposition 3.2. Let $M_t[q]$ denote the evolution of probability mass q up to time t under the corrupting process. Define $p_t^{\alpha,\text{true}}$ as:

$$p_t^{\alpha,\text{true}} = M_t\big[p_0(\cdot)\, p_0(\zeta|\cdot)^\alpha / \mathcal{Z}_\alpha\big].$$

Then,

$$p_t^{\alpha,\text{true}}(x_t) \propto p_t(x_t)\mathbb{E}\big[p_0(\zeta|X_0)^\alpha | X_t = x_t\big].$$

The true tempered rate matrix $R_t^{\alpha,\text{true}}$ given by:

$$\begin{aligned} \forall x \neq y,\ R_t^{\alpha,\text{true}}(x,y|\zeta) &= R_t(x,y)\tfrac{\mathbb{E}[p_0(\zeta|X_0)^\alpha|X_t=y]}{\mathbb{E}[p_0(\zeta|X_0)^\alpha|X_t=x]}, \\ \forall x,\ R_t^{\alpha,\text{true}}(x,x|\zeta) &= -\sum_{y\neq x} R_t^{\alpha,\text{true}}(x,y|\zeta), \end{aligned}$$

satisfies the reverse-time Kolmogorov Forward Equation:

$$\partial_t p_t^{\alpha,\text{true}} = -\left[R_t^{\alpha,\text{true}}\right]^\intercal p_t^{\alpha,\text{true}}.$$

Proof. Recall that $\{X_t\}_{t\in[0,1]}$ is a reversed-time CTMC corresponding to the unconditonal discrete diffusion model. Let $\{p_{t|s}^{\text{corrupt}}(x_t|x_s)\}_{t>s}$ and $\{p_{t|s}(x_t|x_s)\}_{t<s}$ be the transition probability of the forward-time corrupting process and reverse-time generative process. We first show $p_t^\alpha(x_t) \propto p_t(x_t)\mathbb{E}\big[p_0(\zeta|X_0)^\alpha|X_t = x_t\big]$. By definition,

$$\begin{aligned} p_t^{\alpha,\text{true}}(x_t) &= M_t[p_0(\cdot)p(\zeta|\cdot)^\alpha/\mathcal{Z}_\alpha] && (23)\\ &\propto \sum_{x_0} p_{t|0}^{\text{corrupt}}(x_t|x_0)p_0(x_0)p(\zeta|x_0)^\alpha && (24)\\ &= \sum_{x_0} p_{0|t}(x_0|x_t)\frac{p_t(x_t)}{p_0(x_0)}p_0(x_0)p(\zeta|x_0)^\alpha && (25)\\ &= p_t(x_t)\sum_{x_0} p_{0|t}(x_0|x_t)p(\zeta|x_0)^\alpha && (26)\\ &= p_t(x_t)\mathbb{E}[p(\zeta|X_0)^\alpha|X_t = x_t]. && (27) \end{aligned}$$

□

We denote $Z_{t,\alpha}$ the normalising constant of p_t^α such that:

$$p_t^{\alpha,\text{true}}(x_t) = \frac{1}{Z_{t,\alpha}}p_t(x_t)\mathbb{E}[p(\zeta|X_0)^\alpha|X_t = x_t] \tag{28}$$

An other intermediate result that we need is to show that the normalising constant $Z_{t,\alpha}$ is independent of the time t

$$
\begin{aligned}
Z_{t,\alpha} &= \sum_y p_t^\alpha(y) && (29)\\
&= \sum_y p_t(y)\mathbb{E}[p(\zeta|X_0)^\alpha|X_t=y] && (30)\\
&= \sum_y \sum_{x_0} p_t(y)p(\zeta|x_0)^\alpha p_{0|t}(x_0|y) && (31)\\
&= \sum_y \sum_{x_0} p(\zeta|x_0)^\alpha p_{t|0}(y|x_o)p_0(x_0) && (32)\\
&= \sum_{x_0} p(\zeta|x_0)^\alpha p_0(x_0)\left(\sum_y p_{t|0}(y|x_0)\right) && (33)\\
&= \sum_{x_0} p(\zeta|x_0)^\alpha p_0(x_0) && (34)\\
&= \mathbb{E}[p(\zeta|X_0)^\alpha] && (35)
\end{aligned}
$$

We see that $Z_{t,\alpha}$ equals $\mathbb{E}[p(\zeta|X_0)^\alpha]$ which is independent of t and as a consequence $\partial_t Z_{t,\alpha}=0$. We also denote the normalising constant as $Z_\alpha := Z_{t,\alpha}$. We now prove the main result of the proposition with the following derivation:

We first introduce the following notation to make derivation lighter:

$$
w_{t,\alpha}(x_t) = \mathbb{E}[p(\zeta|X_0)^\alpha|X_t=x_t] \tag{36}
$$

Noticing w satisfies a reverse-time Kolmogorov Backward Equation, we have:

$$
\begin{aligned}
\partial_t p_t^\alpha(x_t) &= \partial_t[\frac{1}{Z_\alpha}p_t(x_t)w_{t,\alpha}(x_t)] && (37)\\
&= \frac{1}{Z_\alpha}\left[p_t(x_t)\partial_t w_{t,\alpha}(x_t) + w_{t,\alpha}(x_t)\partial_t p_t(x_t)\right] && (38)\\
&= \frac{1}{Z_\alpha}\left[\sum_y w_{t,\alpha}(y)p_t(x_t)R_t(x_t,y) - \sum_y w_{t,\alpha}(x_t)p_t(y)R_t(y,x_t)\right] && (39)\\
&= \frac{1}{Z_\alpha}\left[\sum_{y\neq x_t} w_{t,\alpha}(y)p_t(x_t)R_t(x_t,y) - \sum_{y\neq x_t} w_{t,\alpha}(x_t)p_t(y)R_t(y,x_t)\right] && (40)\\
&= \frac{1}{Z_\alpha}\left[\sum_{y\neq x_t} p_t(x_t)w_{t,\alpha}(x_t)R_t(x_t,y)\frac{w_{t,\alpha}(y)}{w_{t,\alpha}(x_t)} - \sum_{y\neq x_t} p_t(y)w_{t,\alpha}(y)R_t(y,x_t)\frac{w_{t,\alpha}(x_t)}{w_{t,\alpha}(y)}\right] && (41)\\
&= \sum_{y\neq x_t} p_t^{\alpha,\text{true}}(x_t)R_t^\alpha(x_t,y) - \sum_{y\neq x_t} p_t^{\alpha,\text{true}}(y)R_t^\alpha(y,x_t) && (42)\\
& && (43)
\end{aligned}
$$

and the last equation is the reverse time Forward Kolmogorov Equation which concludes the proof.

D Proof of Proposition 4.1

Proposition 4.1. Let proposal $\{Y_t\}_{t\in[0,1]}$ be a reverse-time CTMC with rate matrix Q_t and initial distribution $p_1(\cdot)$. Define $\{W_t\}_{t\in[0,1]}$ by

$$
W_t = \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t}\left(\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}}\right)^\alpha,
$$

where we have $\mathbb{P}_t^{\text{base}} = \text{Law}\{X_\tau\}_{\tau\in[t,1]}$, $\mathbb{P}^{\alpha=1} = \text{Law}\{X_\tau \mid \zeta\}_{\tau\in[t,1]}$, $\mathbb{Q} = \text{Law}\{Y_\tau\}_{\tau\in[t,1]}$. Then,

$$\ln \frac{\mathrm{d}\mathbb{P}_t{}^{\text{base}}}{\mathrm{d}\mathbb{Q}_t} = \sum_{\substack{\tau\geq t\\ Y_{\tau^+}\neq Y_\tau}} \ln R_t(Y_{\tau^+}, Y_\tau) - \ln Q_t(Y_{\tau^+}, Y s_\tau) + \int_1^t Q_\tau(Y_t) - R_\tau(Y_t)\,\mathrm{d}\tau$$

and

$$\ln \frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}} = \sum_{\substack{\tau\geq t\\ Y_{\tau^+}\neq Y_\tau}} \ln R_t^{\alpha=1}(Y_{\tau^+}, Y_\tau|\zeta) - \ln R_t(Y_{\tau^+}, Y_\tau) + \int_1^t R_\tau(Y_\tau) - R_\tau^{\alpha=1}(Y_\tau|\zeta)\mathrm{d}\tau.$$

For any p_t^α-integrable function h and time $t \in [0,1]$,

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{\mathbb{E}[W_t \cdot h(Y_t)]}{\mathbb{E}[W_t]}.$$

Proof. Via the Radon-Nikodym Theorem, we seek to compute the importance weight (i.e. the Radon-Nikodym derivative [2])

$$\frac{\mathrm{d}\mathbb{P}_t^\alpha}{\mathrm{d}\mathbb{Q}_t} = \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t} \cdot \frac{\mathrm{d}\mathbb{P}_t^\alpha}{\mathrm{d}\mathbb{P}_t^{\text{base}}} = \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t} \cdot p_t(\zeta|Y_t)^\alpha, \tag{44}$$

where $\frac{\mathrm{d}\mathbb{P}_t^\alpha}{\mathrm{d}\mathbb{P}_t^{\text{base}}} = p_t(\zeta|Y_t)^\alpha$ follows by construction of $\mathbb{P}_t^\alpha$ and a direct application of the disintegration theorem (Léonard, 2014) (i.e. the product rule), then considering $\alpha = 1$ it follows that $p_t(\zeta|Y_t)^\alpha = \left(\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}}\right)^\alpha$, and thus the importance weights between our proposal and the tilted path measure reduce to:

$$W_t = \frac{\mathrm{d}\mathbb{P}_t^\alpha}{\mathrm{d}\mathbb{Q}_t} = \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t}\left(\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}^{\text{base}_t}}\right)^\alpha, \tag{45}$$

Finally, the Radon-Nikodym derivative between two arbitrary reverse-time CTMC path measures $\mathbb{P}_t'$ and $\mathbb{Q}_t'$ with transition rate matrix R_τ' and Q_τ' evaluated at the path Y, is known to be the following Appendix C.1 of (Campbell et al., 2024), that is,

$$\frac{\mathrm{d}\mathbb{P}_t'}{\mathrm{d}\mathbb{Q}_t'} = \frac{\exp(-\int_1^t R_\tau'(Y_t)\,\mathrm{d}\tau)}{\exp(-\int_1^t Q_\tau'(Y_t)\,\mathrm{d}\tau)} \prod_{\substack{\tau\geq t\\ Y_{\tau^+}\neq Y_\tau}} \frac{R_t'(Y_{\tau^+}, Y_\tau)}{Q_t'(Y_{\tau^+}, Y_\tau)} \tag{46}$$

We refer the readers to (Campbell et al., 2024) for an accessible sketch of this result, which follows (Del Moral & Penev, 2017). Then, substituting Equation 46 into the RNDs of Equation 45 concludes the proof. ∎

D.1 Itô Process Counterpart

The core result in Proposition Proposition 4.1 is in noticing that $p_t(\zeta|Y_t)^\alpha = \left(\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}}\right)^\alpha$ and decomposing the IS weight via the chain rule and applying Girsanovs Theorem.

Due to the very modular nature of our result our proposition seamlessly extends to Stochastic Differential Equations.

Proposition D.1. *Let proposal $\{Y_t\}_{t\in[0,1]}$ be a reverse-time SDE with drift $b_t(x)$, diffusion coeficient g_t and initial distribution $p_1(\cdot)$ and X_t also be a reverse-time SDE with drift f_t, diffusion coeficient g_t and initial distribution $p_1(\cdot)$. Define $\{W_t\}_{t\in[0,1]}$ by*

$$W_t = \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t}\left(\frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}}\right)^\alpha, \tag{47}$$

[2] We consider all RND's evaluated at Y_t that is when we write $\frac{\mathrm{d}\mathbb{P}_t'}{\mathrm{d}\mathbb{Q}_t'}$ it is short for $\frac{\mathrm{d}\mathbb{P}_t'}{\mathrm{d}\mathbb{Q}_t'}(Y_t)$.

where we have $\mathbb{P}_t^{base} = \mathrm{Law}\{X_\tau\}_{\tau\in[t,1]}$, $\mathbb{P}^{\alpha=1} = \mathrm{Law}\{X_\tau \mid \zeta\}_{\tau\in[t,1]}$, $\mathbb{Q} = \mathrm{Law}\{Y_\tau^{b,p_1}\}_{\tau\in[t,1]}$. *Then,*

$$\ln \frac{\mathrm{d}\mathbb{P}_t^{\text{base}}}{\mathrm{d}\mathbb{Q}_t} = \int_t^1 \frac{1}{g_\tau^2}(f_\tau(Y_\tau) - b_\tau(Y_\tau))^\top \overleftarrow{\mathrm{d}}\, Y_\tau + \int_t^1 \frac{1}{2g_\tau^2}(||^2 f_\tau(Y_\tau)||^2 - ||b_\tau(Y_\tau)||^2)\mathrm{d}\tau$$

and

$$\ln \frac{\mathrm{d}\mathbb{P}_t^{\alpha=1}}{\mathrm{d}\mathbb{P}_t^{\text{base}}} = \int_t^1 \frac{1}{g_\tau^2}(f_\tau^{\cdot|\zeta}(Y_\tau) - f_\tau(Y_\tau))^\top \overleftarrow{\mathrm{d}}\, Y_\tau + \int_t^1 \frac{1}{2g_\tau^2}(||f_\tau^{\cdot|\zeta}(Y_\tau)||^2 - ||f_\tau(Y_\tau)||^2)\mathrm{d}\tau$$

Where $f_\tau^{\cdot|\zeta}(y_\tau) = f_\tau(y_\tau) - g_\tau^2 \nabla \ln p_\tau(\zeta|y_\tau)$ *as per (Denker et al., 2024, Proposition 2.2).*

For any p_t^α*-integrable function* h *and time* $t \in [0,1]$*,*

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{\mathbb{E}[W_t \cdot h(Y_t)]}{\mathbb{E}[W_t]}.$$

Proof. Equation 47 follows directly from Proposition 4.1. Then, for the RNDs between two reverse time SDEs, we use Equation 64 in Vargas et al. (2024). Note for the VP-SDE case; these expressions can be simplified further as done in (Vargas et al., 2023, Equation 9). Note as with the rest of our results, we require that $p_1(y_1|\zeta) = p_1(y_1)$ which is approximately true in the limit for VP-SDE as $p_1(\cdot|\zeta) \approx \mathcal{N}(0, I)$. □

D.2 Alternative Proof of Proposition 4.1 without Radon-Nikodym Theorem

In this appendix, we present an equivalent statement of Proposition 4.1 and a proof that does not rely on the Radon-Nikodym Theorem.

Proposition 4.1. Let proposal $\{Y_t\}_{t\in[0,1]}$ be a reverse-time CTMC with rate matrix Q_t and initial distribution $p_1(\cdot)$. Define $\{W_t\}_{t\in[0,1]}$ by,

$$W_t = \exp(A_t),$$
$$A_0 = 0,$$
$$A_t = \sum_{\substack{t\le\tau \\ Y_{\tau^+}\neq Y_\tau}} [\ln R_\tau(Y_{\tau^+}, Y_\tau) - \ln Q_\tau(Y_{\tau^+}, Y_\tau)] + \sum_{\substack{t\le\tau \\ Y_{\tau^+}\neq Y_\tau}} \alpha \ln(R_\tau^{\alpha=1}(Y_{\tau^+}, Y_\tau|\zeta)/R_\tau(Y_{\tau^+}, Y_\tau))$$
$$+ \int_1^t \left[Q_\tau(Y_t) - R_\tau(Y_t) + \alpha R_\tau(Y_\tau|\zeta) - \alpha R_\tau^{\alpha=1}(Y_\tau|\zeta)\right] \,\mathrm{d}\tau.$$

For any p_t-integrable function h and time $t \in [0,1]$,

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{\mathbb{E}[W_t \cdot h(Y_t)]}{\mathbb{E}[W_t]}.$$

where the expectation is taken over the law of (Y_t, W_t).

Proof. We closely follow the sketches of (Carbone et al., 2023; Albergo & Vanden-Eijnden, 2024). Assume $p_1(x_1) = p_1(x_1|\zeta)$. We first show that the coupled system (Y_t, A_t) with A_t defined as,

$$W_0 = 0 \tag{48}$$
$$A_t = \sum_{\substack{t\le\tau \\ Y_{\tau^+}\neq Y_\tau}} [\ln R_t(Y_{\tau^+}, Y_\tau) - \ln Q_t(Y_{\tau^+}, Y_\tau)] + \sum_{\substack{t\le\tau \\ Y_{\tau^+}\neq Y_\tau}} \alpha \left[\ln p_t(\zeta|Y_\tau) - \ln p_t(\zeta|Y_{\tau^+})\right] \tag{49}$$
$$+ \int_1^t \left[R_\tau(Y_t) - Q_\tau(Y_t) - \alpha\partial_t \ln p_t(\zeta|Y_t)\right] \,\mathrm{d}\tau, \tag{50}$$

satisfies,

$$\mathbb{E}_{x\sim p_t^\alpha}[h(x)] = \frac{\mathbb{E}[e^{A_t} h(Y_t)]}{\mathbb{E}[e^{A_t}]}, \tag{51}$$

Write $\Delta A_t(x,y) = \ln R_t(x,y) - \ln Q_t(x,y) + \alpha(\ln p_t(\zeta|y) - \ln p_t(\zeta|x))$. Define $f_t(y,a)$ to be the joint density of (Y_t, A_t), the joint density is governed by the (reverse-time) continuity equation,

$$\partial_t f_t(y,a) = -\left[\sum_{z\neq y} Q_t(z,y) f_t(z, a - \Delta A(z,y)) - f_t(y,a) Q_t(y) + \left[Q_t(y) - R_t(y) - \alpha \partial_t \ln p_t(\zeta|y)\right] \partial_a f_t(y,a)\right],$$

$$f_1(y,a) = \delta(a) p_1(x_1) \tag{52}$$

□

Define $g_t(y) = \int_{\mathbb{R}} e^a f_t(y,a) da$. Then, via dominated convergence theorem (i.e. $\partial_t g_t(y) = \int_{\mathbb{R}} e^a \partial_t f_t(y,a) da$) and substituting $\partial_t f_t(y,a)$ from Equation 52, it follows that

$$\partial_t g_t(y) = \int_{\mathbb{R}} e^a \sum_{z\neq y} Q_t(z,y) f_t(z, a - \Delta A(z,y)) da \tag{53}$$

By a change of variable $a' = a - \Delta A(z,y)$, then,

$$\int_{\mathbb{R}} e^a \sum_{z\neq y} Q_t(z,y) f_t(z, a - \Delta A(z,y)) da = \int_{\mathbb{R}} e^{a'} \sum_{z\neq y} e^{\Delta A(z,y)} Q_t(z,y) f_t(z, a') da' \tag{54}$$

$$= \sum_{z\neq y} R_t(z,y) \cdot \frac{p_t(\zeta|y)^\alpha}{p_t(\zeta|z)^\alpha} \cdot g_t(z) \tag{55}$$

And finally, by integration by parts, and noting that $\int_{\mathbb{R}} e^a f_t(y,a) da = Q_t(y) g_t(y)$, we have

$$\int_{\mathbb{R}} e^a \left[Q_t(y) - R_t(y) + \alpha \partial_t \ln p_t(\zeta|y)\right] \partial_a f_t(y,a) da = \left[Q_t(y) - R_t(y) - \alpha \partial_t \ln p_t(\zeta|y)\right] g_t(y). \tag{56}$$

Combining the three cases, we have:

$$g_1(y) = p_1(y) \tag{57}$$

$$\partial_t g_t(y) = -\left[\sum_{z\neq y} R_t(z,y) \cdot \frac{p_t(\zeta|y)^\alpha}{p_t(\zeta|z)^\alpha} \cdot g_t(z) - Q_t(y) g_t(y) + \left[Q_t(y) - R_t(y) - \alpha \partial_t \ln p_t(\zeta|y)\right] g_t(y)\right]. \tag{58}$$

$$\tag{59}$$

Notice that g reduces to a system of ODEs and thus, by the Picard-Lindelöf theorem, has a unique solution. Recall that $\mathcal{Z}_1^\alpha = \sum_x p_1(x) p_1(\zeta|x)^\alpha = p(\zeta)^\alpha$, writing $p(\zeta)$ as the density of ζ. Substituting $g_t(y) = p_t(y) \cdot p_t(\zeta|y)^\alpha / \mathcal{Z}_1^\alpha$, then,

$$\text{L.H.S} = \partial_t g(y) \tag{60}$$

$$= \frac{1}{\mathcal{Z}_1^\alpha} \left[p_t(\zeta|y)^\alpha \partial_t p_t(y) + p_t(y) \partial_t (p_t(\zeta|y)^\alpha)\right] \tag{61}$$

$$\text{R.H.S} = -\frac{1}{\mathcal{Z}_1^\alpha} \left[\sum_{z\neq y} R_t(z,y) \cdot \frac{p_t(\zeta|y)^\alpha}{p_t(\zeta|z)^\alpha} \cdot g_t(z) - Q_t(y) g_t(y) + \left[Q_t(y) - R_t(y) - \alpha \partial_t \ln p_t(\zeta|y)\right] g_t(y)\right] \tag{62}$$

$$= -\frac{1}{\mathcal{Z}_1^\alpha} \left[p_t(\zeta|y)^\alpha \sum_{z\neq y} R_t(z,y) p_t(z) - p_t(\zeta|y)^\alpha \sum_{z\neq y} R_t(y,z) p_t(y) - \alpha \partial_t \ln p_t(\zeta|y) p_t(y) p_t(\zeta|y)^\alpha\right] \tag{63}$$

$$= \frac{1}{\mathcal{Z}_1^\alpha} \left[p_t(\zeta|y)^\alpha \partial_t p_t(y) + p_t(y) \partial_t (p_t(\zeta|y)^\alpha)\right] \tag{64}$$

$$= \text{L.H.S} \tag{65}$$

Since $p_1(y) = p_1(y \mid \zeta)$, it follows that the boundary condition $g_1(y) = p_1(y)$ is fulfilled. Then,

$$\mathbb{E}[e^{A_t}] = \sum_x \int_{\mathbb{R}} e^a f_t(x,a)\mathrm{d}a = \sum_x g_t(x) = \sum_x p_t(x)p_t(\zeta|x)^\alpha / \mathcal{Z}_1^\alpha = \mathcal{Z}_t^\alpha / \mathcal{Z}_1^\alpha \tag{66}$$

Similarly, for an arbitrary function h if we define $g_t^h(x) = \int_{\mathbb{R}} e^a h(x) f_t(x,a)\mathrm{d}a$, we may show that

$$\mathbb{E}[e^{A_t} h(x)] = \sum_x \int_{\mathbb{R}} e^a h(x) f_t(x,a)\mathrm{d}a = \sum_x h(x)p_t(x)p_t(\zeta|x)^\alpha / \mathcal{Z}_1^\alpha. \tag{67}$$

Combining the two equations, we have,

$$\frac{\mathbb{E}[e^{A_t} h(x)]}{\mathbb{E}[e^{A_t}]} = \frac{1}{\mathcal{Z}_t^\alpha} \sum_x h(x)p_t(x)p_t(\zeta|x)^\alpha, \tag{68}$$

as desired.

To finish off the proof, we show that following:

1. $\forall x \neq y.\ \alpha(\ln p_t(\zeta|y) - \ln p_t(\zeta|x)) = \alpha(\ln R_t^{\alpha=1}(x,y) - \ln R_t(x,y))$
2. $\partial_t \ln p_t(\zeta|x_t) = \alpha R_t^{\alpha=1}(x) - \alpha R_t(x)$

The first equation follows from the definition of the guided ratio matrix $R_t^{\alpha=1}(x)$. To show that $\partial_t \ln p_t(\zeta|x_t) = \alpha R_t(x) - \alpha R_t^{\alpha=1}(x)$, recall,

$$p_t(\zeta|x_t) = \mathbb{E}[p_t(\zeta|X_0)|X_t = x_t], \tag{69}$$

Therefore,

$$\partial_t \ln p_t(\zeta|x_t) = \frac{1}{p_t(\zeta|x_t)} \partial_t p_t(\zeta|x_t) \tag{70}$$

$$= \frac{1}{p_t(\zeta|x_t)} \sum_y R_t(x_t,y)p_t(\zeta|y) \quad \text{(Reverse-Time) Kolmogorov Backward Equation} \tag{71}$$

$$= R_t^{\alpha=1}(x_t) - R_t(x_t) \qquad \text{Definition of } R_t^{\alpha=1} \tag{72}$$

This concludes the proof.

E Statement and Derivation of Discretised Weights

In this appendix, we show a discrete time version of Proposition 4.1 and the proof.

Proposition E.1. *Define $\{\hat{W}_{t_k}\}_{k\in\{1,\ldots,T\}}$ as:*

$$\hat{W}_0 = 1,$$
$$\hat{W}_{t_k} = \prod_{1\le j<k} \frac{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})}{q_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})} \cdot \prod_{1\le j<k} \left[\frac{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j},\zeta)}{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})}\right]^\alpha .$$

If $p_1(x_1) = p_1(x_1|\zeta)$, then for any function $h : \mathcal{X} \to \mathbb{R}$ and $k \in \{1,\ldots,T\}$,

$$\frac{1}{\mathcal{Z}_{t_k}^\alpha} \sum_x h(x)p_{t_k}(x)p_{t_k}(\zeta|x)^\alpha = \frac{\mathbb{E}\left[\hat{W}_{t_k} h(Y_{t_k})\right]}{\mathbb{E}\left[\hat{W}_{t_k}\right]},$$

where the expectation is taken over the law of $(Y_t, \hat{W}_t)$.

Proof. First note that,

$$p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j},\zeta) = p^{\text{corrupt}}_{t_j|t_{j+1}}(Y_{t_j}|Y_{t_{j+1}}) \frac{M_t[p_0(x_0|\zeta)](Y_{t_{j+1}})}{M_t[p_0(x_0|\zeta)](Y_{t_j})} \tag{73}$$

$$= p^{\text{corrupt}}_{t_j|t_{j+1}}(Y_{t_j}|Y_{t_{j+1}}) \frac{p_t(Y_{t_{j+1}})p_t(\zeta|Y_{t_{j+1}})}{p_t(Y_{t_j})p_t(\zeta|Y_{t_j})} \tag{74}$$

$$= p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j}) \frac{p_t(\zeta|Y_{t_{j+1}})}{p_t(\zeta|Y_{t_j})}. \tag{75}$$

Then it follows that,

$$\hat{W}_{t_k} = \prod_{1\le j<k} \frac{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})}{q_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})} \prod_{1\le j<k} \left[\frac{p_t(\zeta|Y_{t_{j+1}})}{p_t(\zeta|Y_{t_j})}\right]^\alpha \tag{76}$$

$$= \prod_{1\le j<k} \frac{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})}{q_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})} \left[\frac{p_{t_k}(\zeta|Y_{t_k})}{p_1(\zeta|Y_1)}\right]^\alpha \tag{77}$$

$$= \prod_{1\le j<k} \frac{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})}{q_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})} \left[\frac{p_{t_k}(\zeta|Y_{t_k})}{p(\zeta)}\right]^\alpha \qquad Y_1 = X_1` \perp \zeta \tag{78}$$

$$= \frac{1}{\mathcal{Z}_1^\alpha} \prod_{1\le j<k} \frac{p_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})}{q_{t_{j+1}|t_j}(Y_{t_{j+1}}|Y_{t_j})} \qquad \mathcal{Z}_1^\alpha \left[p_{t_k}(\zeta|Y_{t_k})\right]^\alpha = \sum_x p_1(x)p_1(\zeta|x)^\alpha = p(\zeta)^\alpha. \tag{79}$$

Then,

$$\mathbb{E}\left[\hat{W}_{t_k} h(Y_{t_k})\right] = \sum_{Y_{t_k}} \hat{W}_{t_k} h(Y_{t_k}) q_{t_k}(Y_{t_k})$$

$$= \frac{1}{\mathcal{Z}_1^\alpha} \sum_{\{y_i\}_{i=1}^k : y_k = Y_{t_k}} \hat{W}_{t_k} h(Y_{t_k}) \prod_{j=1}^{k-1} q_1(y_1) q_{t_{j+1}|t_j}(y_{t_{j+1}}|y_{t_j})$$

$$= \frac{1}{\mathcal{Z}_1^\alpha} \sum_{\{x_i\}_{i=1}^k : x_k = Y_{t_k}} h(Y_{t_k}) p_t(\zeta|Y_{t_k})^\alpha q_1(x_1) \prod_{j=1}^{k-1} p_{t_{j+1}|t_j}(x_{t_{j+1}}|x_{t_j})$$

$$= \frac{1}{\mathcal{Z}_1^\alpha} \sum_{Y_{t_k}} h(Y_{t_k}) p_t(\zeta|Y_{t_k})^\alpha p_t(Y_k) \qquad \text{Both } Y_t \text{ and } X_t \text{ starts at } p_1$$

Similarly,

$$\mathbb{E}\left[\hat{W}_{t_k} h(Y_{t_k})\right] = \mathcal{Z}_t^\alpha / \mathcal{Z}_1^\alpha. \tag{80}$$

Combining the two equation lends us to,

$$\frac{1}{\mathcal{Z}_{t_k}^\alpha} \sum_x h(x) p_{t_k}(x) p_{t_k}(\zeta|x)^\alpha = \frac{\mathbb{E}\left[\hat{W}_{t_k} h(Y_{t_k})\right]}{\mathbb{E}\left[\hat{W}_{t_k}\right]} \tag{81}$$

This concludes the proof. □

F SAMPLING WITH APPROXIMATE TRANSITION DISCRETISED WEIGHT

The exact transition probabilities $p_{t_{j+1}|t_j}$ and $q_{t_{j+1}|t_j}$ is typically unknown. In practice, we approximate these transitions using numerical methods such as Euler sampling. As the time step Δt approaches zero, these approximations converge to the true weights W_t. In Algorithm 2, we show the pseudocode of Algorithm 1 using the approximated transition probability $\tilde{q}_{t|s}$ and $\tilde{p}_{t|s}$.

Algorithm 2 Main Algorithm with Approximate Transition

Number of particles K; Approximate Proposal Transition $\tilde{q}_{t|s}$; Approximate transition $\tilde{p}_{t|s}$; Temperature α; Time-grid $\{t_l\}_{l=1}^{T}$; ESS threshold ESS_THRESHOLD; Resampling algorithm `resample`. **Initialisation:** Sample $\{x_1^{(i)}\}_{i=1}^{K}$ i.i.d. from p_1. Set $\hat{w}_1^{(i)} \leftarrow 0$ for $i = 1, \ldots, K$. $l = 1$ to T $i = 1$ to K Step 1: Propagate Sample $x_{l+1}^{(i)} \sim \tilde{q}_{t_{l+1}|t_l}(\cdot|x_l^{(i)})$ Step 2: Update and Normalize Weights Set $\hat{w}_{l+1}^{(i)} \leftarrow \hat{w}_l^{(i)} \cdot \frac{\tilde{p}_{t_{l+1}|t_l}(x_{l+1}^{(i)}|x_l^{(i)})}{\tilde{q}_{t_{l+1}|t_l}\left(x_{l+1}^{(i)}|x_l^{(i)}\right)} \cdot \left[\frac{\tilde{p}_{t_{l+1}|t_l}(x_{l+1}^{(i)}|x_l^{(i)},\zeta)}{\tilde{p}_{t_{l+1}|t_l}\left(x_{l+1}^{(i)}|x_l^{(i)}\right)}\right]^{\alpha}$ Step 3: Resample; see Appendix G for other resampling algorithms Set ESS $\leftarrow \left(\sum_{i=1}^{K} w_l^{(i)}\right)^2 / \sum_{i=1}^{K}(w_l^{(i)})^2$ ESS $\leq$ ESS_THRESHOLD $\tilde{w}_l^{(i)} \leftarrow \hat{w}_l^{(i)} / \sum_i \hat{w}_l^{(i)}$ Set $x_{l+1}^{(i)}, \hat{w}_{l+1}^{(i)} \leftarrow$ `resample`$(\{x_l^{(i)}\}, \{\tilde{w}_l^{(i)}\})$ **Output:** Particles $\{x_T^{(i)}\}_{i=1}^{K}$

G PSEUDOCODE FOR RESAMPLING ALGORITHMS

In this section, we introduce various resampling strategies for Sequential Monte Carlo methods. Ultimately, the goal of resampling algorithm is to replace low weighted particles with high weighted particles to reduce the variance of weights such that the expectation of any function ϕ remains unchanged.

More specifically, consider at iteration l of a SMC algorithm, given K particles $\{x_l^{(i)}\}_{i=1}^K$ and $\{w_l^{(i)}\}_{i=1}^K$ from some joint distribution $\mathbb{Q}$ induced by the SMC algorithm. A resampling algorithm returning $\{\tilde{x}^{(i)}\}_{i=1}^K$ and $\{\tilde{w}^{(i)}\}_{i=1}^K$ generates consistent samples as long as for any bounded and continuous function ϕ, $\mathbb{E}^{\mathbb{Q}}[w^{(i)}\phi(x^{(i)})] = \mathbb{E}^{\mathbb{Q}}[\tilde{w}^{(i)}\phi(\tilde{x}^{(i)})]$ where $\mathbb{E}^{\mathbb{Q}}$ notes an expectation on $\mathbb{Q}$.

G.1 MULTINOMIAL RESAMPLING

The simplest and most common resampling strategy is the multinomial resampling, which involves the particles in the next iteration independently following a categorical distributions distributed according to the normalised weights.

Algorithm 3 Multinomial Resampling

Number of Particles K; Particles $\{x^{(i)}\}_{i=1}^K$; Normalised Weights $\{w^{(i)}\}_{i=1}^K$ $i = 1$ to K Sample $\tilde{x}^{(i)} \sim \sum_{j=1}^K w^{(i)}\delta_{x^{(i)}}(\mathrm{d}x)$ $\tilde{w}^{(i)} = 1/K$ **Output:** Particles $\{\tilde{x}^{(i)}\}_{i=1}^K$ and Weights $\{\tilde{w}^{(i)}\}_{i=1}^K$

G.2 PARTIAL RESAMPLING

We further consider employing a partial resampling strategy adapted from (Martino et al., 2016), which we found to prevent mode collapse of samples effectively.

Algorithm 4 Partial Resampling

Number of Particles K; Particles $\{x^{(i)}\}_{i=1}^K$; Normalised Weights $\{w^{(i)}\}_{i=1}^K$; Resample Size M Set Resample Indices $I = \{i | w^{(i)}$ among the $\lfloor M/2 \rfloor$ highest or $\lceil M/2 \rceil$ lowest$\}$ $i = 1$ to K $i \in I$ Sample $\tilde{x}^{(i)} \sim \sum_{j=1}^K w^{(i)}\delta_{x^{(i)}}(\mathrm{d}x)$ Set $\tilde{w}^{(i)} \leftarrow 1/M \sum_{i \in I} w^{(i)}$ Set $\tilde{x}^{(i)} \leftarrow x^{(i)}$ Set $\tilde{w}^{(i)} \leftarrow w^{(i)}$ **Output:** Particles $\{\tilde{x}^{(i)}\}_{i=1}^K$ and Weights $\{\tilde{w}^{(i)}\}_{i=1}^K$

H Does partial resampling mitigate mode collapse?

The MNIST dataset visually demonstrates mode collapse and the effectiveness of partial resampling (Section 6.2). Using SMC with guidance temperature $\beta = 1$ and SMC temperature $\alpha = 3$, we observe mode collapse in Figure 3 (centre), where 16 particles converge to identical images. Partial resampling (Figure 3, right) maintains both diversity and stronger conditional control compared to naive guidance (Figure 3, left), which occasionally generates incorrect digits.

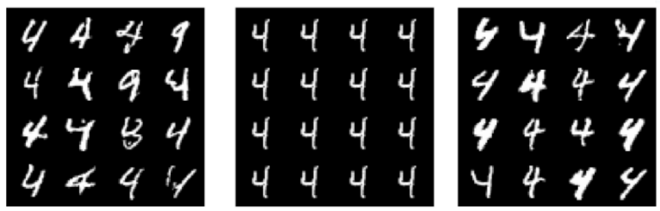

Figure 3: **Class-Conditioned Image Generation** $\zeta = 4$: Left to right: Independent guidance, SMC without partial resampling, SMC with partial resampling. The latter produces high-quality digit 4 samples without mode collapse.

I Figures on text results

Figures 4, 6, and 5 show how α and β affect perplexity and control metrics. With high α, SMC outperforms guidance across tasks, with the gap narrowing at high β but at the cost of deteriorating perplexity, indicating that the proposal is unable to sample from high likelihood region. For toxicity, SMC achieves better control until $\beta > 1.4$ with comparable perplexity. For sentiment, SMC improves accuracy until $\beta > 1.2$ with slightly higher perplexity, while at $\beta > 1.4$ and $\alpha = 1$, it achieves lower perplexity with equal accuracy. For infilling, SMC maintains superior or equal accuracy across all β, with better perplexity at $\beta \in \{0.4, 0.6\}$ or at $\alpha = 1, \beta > 1.4$.

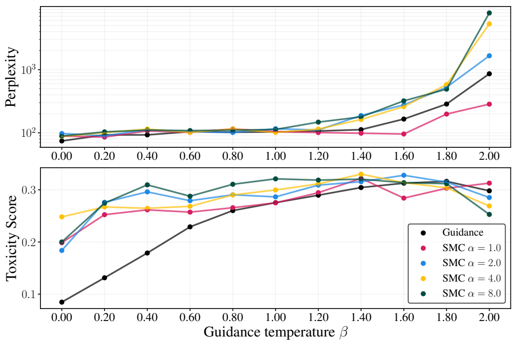

Figure 4: **Toxicity controlled generation:** SMC achieves better toxicity scores across all guidance temperatures β. While baseline methods match SMC's control at high β, perplexity deteriorates.

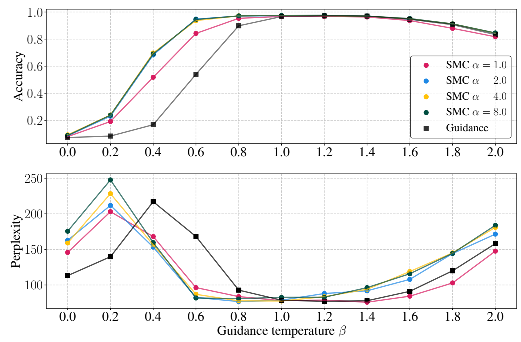

Figure 5: **Text infilling**: SMC achieves higher accuracy across all guidance temperatures, with better perplexity at $\beta \in 0.4, 0.6$ and at $\beta \geq 1.6$ with $\alpha = 1$.

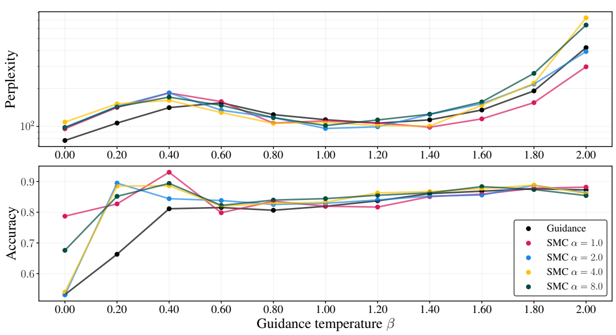

Figure 6: **Sentiment control generation:** For guidance temperatures $\beta \in \{0.2, 0.4\}$, SMC approaches achieve notably better accuracy compared to guidance, though at the cost of slightly higher perplexity. For guidance temperature greater than 1.4 and for $\alpha = 1$, SMC method achieves lower perplexity compared to guidance.

J ADDITIONAL TEXT GENERATION RESULT

We display results on text generation control when using the DEFT guidance term. We find that:

For toxicity-controlled generation, SMC-based methods outperform guidance on control metrics across the guidance temperature range $[0.0, 2.0]$, which is even more apparent for higher SMC temperature α. This is at the expense of a slightly degraded perplexity compared to guidance. At high guidance temperature, the guidance method is on par with SMC-based methods on toxicity score, at the expense of a significantly deteriorated perplexity. This indicates that at high guidance temperatures, the samples deviate from the fine-tuned data distribution. In that scenario, a good balance between perplexity and control generation is in the mid-to-low range of guidance temperature.

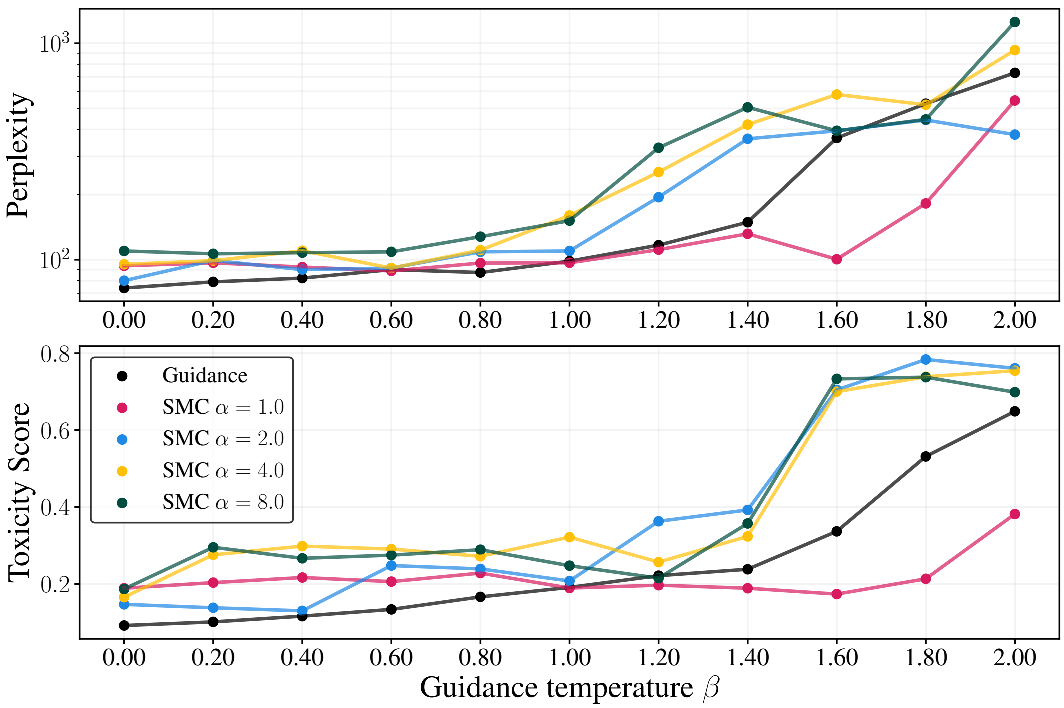

Figure 7: **Toxicity controlled generation:** Guidance is performed with DEFT. SMC methods show improved toxicity scores across all proposal temperatures β, particularly for high β, demonstrating enhanced generation control.

For text infilling, SMC-based methods achieve superior accuracy for guidance temperatures $\beta \in [0.2, 0.8]$. However, at high guidance temperatures, both SMC and guidance methods attain comparable optimal accuracy and perplexity. In this specific scenario, despite lacking theoretical guarantees, the guidance method proves more practical due to its lower computational cost while maintaining equivalent performance.

For sentiment-controlled generation, SMC methods achieve higher accuracy for guidance temperatures $\beta < 0.6$, which comes at the cost of a slight increase in perplexity. At higher guidance temperatures ($\beta > 1.2$), both SMC and standard guidance converge to similar accuracy levels around 0.85, with comparable perplexity until $\beta = 1.6$. Beyond this point, all methods show significant perplexity deterioration, with SMC methods exhibiting slightly higher perplexity than standard guidance.

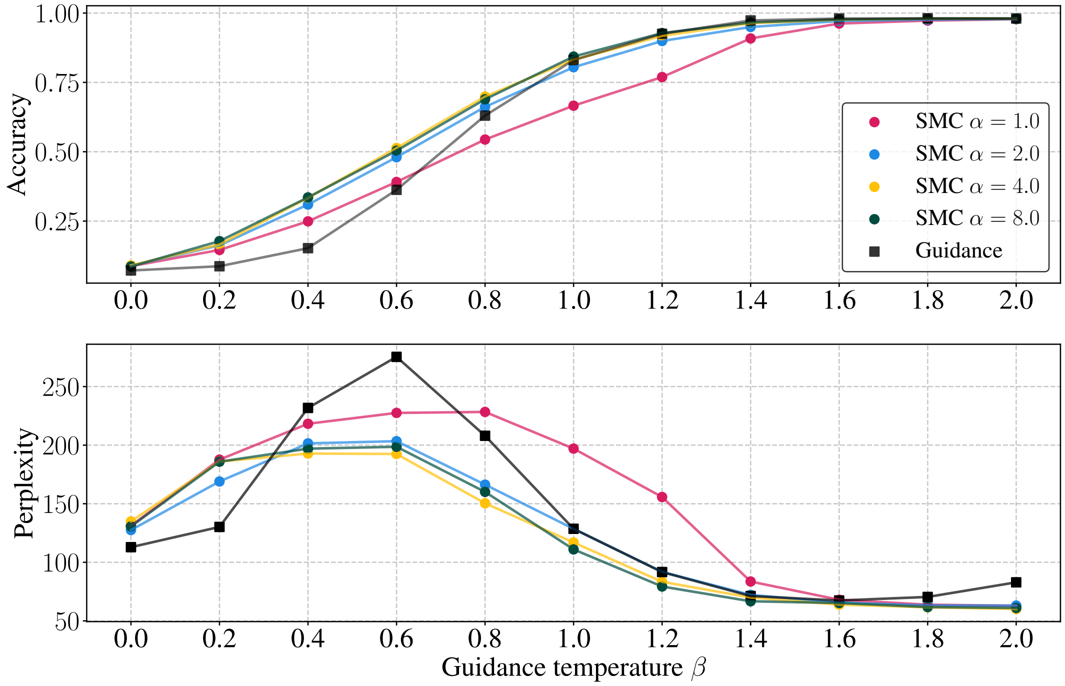

Figure 8: **Text infilling:** Guidance is performed with DEFT. SMC based methods show improved accuracy for $\beta \in \{0.2, 0.4, 0.6, 0.8\}$ while keeping a lower perplexity for $\beta \in \{0.4, 0.6\}$

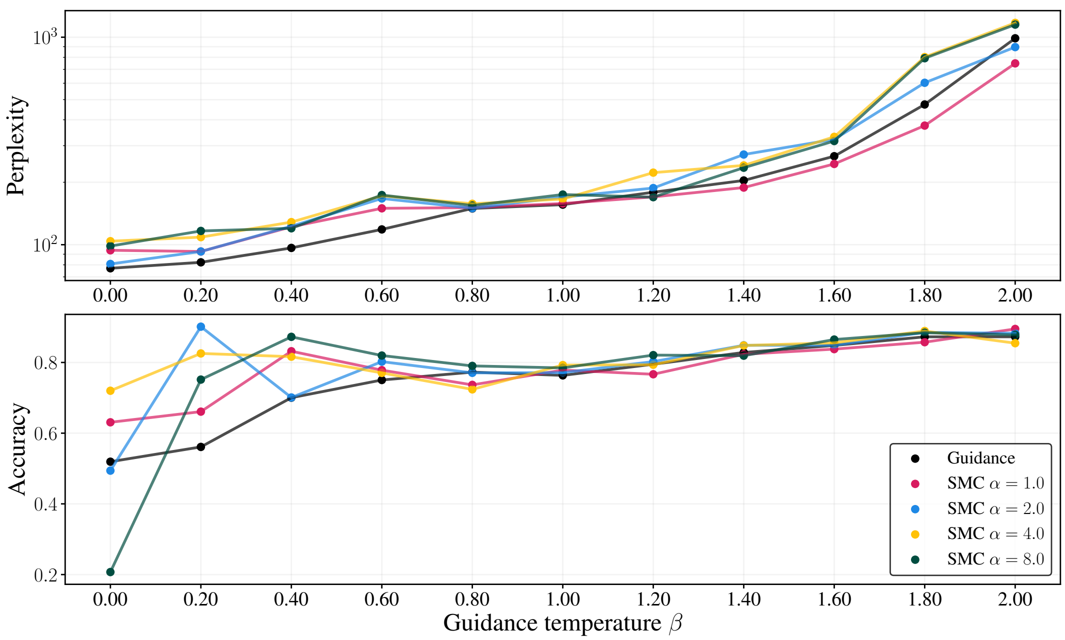

Figure 9: **Sentiment controlled generation**: Guidance is performed with DEFT. For guidance temperature lower than 0.6 SMC methods show improved accuracy while having higher perplexity. For larger guidance temperature accuracy and perplexity are relatively similar.

K SAMPLES

ant EOT a very funny and rewarding small comedy . EOT a rather else altogether EOT genuinely sweet and EOT evokes its low-budget constraints fairly well EOT , evocative EOT 's spooky . EOT an interesting concept EOT here 's the treat . EOT a solidly entertaining film in its own right EOT supremely goofy pleasure EOT demonstrated extraordinary faith EOT an unflinching meditation on how business can change those who enter EOT love the film . EOT – deep down you 've a family EOT a fascinating glimpse EOT of place with its oscar-worthy predecessors EOT pounding, offering fine acting moments , EOT it 's a surest bet EOT an engaging fantasy EOT , the material actually springs to life EOT director credit gyllen brings together caine and grant EOT star charisma EOT consistent EOT what 's best about dustin gondry 's sophisticated read my lips EOT EOT ... scotches . EOT 's surprisingly faithful EOT EOT dazzling EOT make up for the execution 's weak comic buttons . EOT curiously funny, intriguing , EOT just that satisfying exploration of love and companionship EOT EOT emotional sympathy EOT make it entertaining , EOT ungainly work EOT strongest films EOT funny EOT to sharp writing EOT make for a great director 's to make, everlasting EOT innocence EOT though not technically proficient or as EOT remark to ' believer ' in some time EOT sometimes insightful writer EOT mr. honorably EOT and unforgettable characters EOT cool, cocky, and – dare i say twice – hollywood-esque EOT powerful story EOT surprisingly sweet EOT engaging, compelling tale EOT beautiful EOT vibrant air EOT but ultimately winning film EOT its simple expressiveness EOT enriched than limited by its exceptional lead performances EOT of impressively challenging social drama EOT the rich formalism EOT EOT keep you EOT of laughs EOT quieter EOT the praises EOT some monster movie EOT as his penchant for documentaries EOT is latin star to young international hip-hop audiences EOT often funny romantic comedy EOT smoother EOT only enhance his good looks EOT immensely enjoyable EOT one of the most idiosyncratic comedies EOT oddly watchable EOT ambitious EOT directed personalities EOT interesting, i must say , EOT faithful , EOT well, EOT well lived EOT compelling and gripping EOT everyone 's comic relief EOT truly funny EOT enjoy EOT elegant delivery EOT idea EOT familiarity EOT to hold you well past its 90 minutes EOT find love EOT fresh sense EOT has spliceed together bits and pieces of hagney 's reign EOT for good, scary movies EOT gaining this movie EOT flaccid satire , EOT EOT respectable, sweetly adventurous EOT a film EOT smart and subtle EOT vibrant and EOT their fathers EOT her maternal fury of faith EOT EOT protagonists EOT is surprisingly insightful and fun and entertaining EOT that acts as if it never had any idea what EOT , originality and technical skill EOT argue EOT above all, blue crush is nothing to overlook . EOT sulky and beautiful rendering of an intense mystery about the world 's greatest teacher EOT daring teacher EOT cinematic stamp EOT warm , EOT its inviting EOT sexy EOT is realistically terrifying EOT a surprising entree EOT invigoratingly irresistible EOT action-power appeal EOT slyly debated film EOT mention, engaging premise EOT , there are enough laughs to keep on clicking through EOT puts us off-center in the unfolding of bielinsky 's clever scenario , EOT the local flavor and french musicmaking traditions EOT elegance EOT it has a real " story " EOT at least the funniest idea EOT works well on different levels, until it is actually funny . EOT , cleverly crafted and EOT scenes brimmed EOT allows you to forgive its basic humanity EOT that tickles the chills of EOT subtle EOT the film a stunning lesson in human-scale acting , EOT 's a film school exercise with dignity and wit that " should have been right to grow up EOT a riot like john ritter 's brothers ... EOT thumbs up EOT seemingly EOT they 're the target audience because EOT sensitive performances and EOT magnificent landscape EOT caric ballot EOT bring more goodies . EOT as cohesive , EOT some first-time director EOT of the finest kind EOT a spirited EOT 're out tonight EOT style and color EOT a cleverly written and thoroughly winning portrayal of titular kid a woodman who 's become a self image with kids EOT emotional evolution EOT she 's a tempting bouquet – if you live a minute ... come and see the charm of the movie EOT is credible EOT a classical actress EOT delicious camera work EOT a treat , EOT a more absorbing piece EOT all builds up to the easy, endearing knockout . EOT still charming – likely his next film wo n't be the greatest date movie EOT , sweet EOT . EOT most good-hearted EOT a few good ideas EOT a feel-good thriller EOT , illuminating study EOT large metaphor EOT successful EOT fresh is " fresh " again EOT tell this tale EOT felt and affecting EOT

Figure 10: Sentiment control - positive samples - SMC $\beta = 1.0$, $\alpha = 2.0$. EOT indicates End Of Text

having a difficult time telling what time in the past year . `EOT` a pale successor `EOT` an odd bit of an erratic `EOT` killing time `EOT` wobbly `EOT` mean-spirited `EOT` 's no energy `EOT` carries little insight on either end of the other `EOT` dangerous comedy `EOT` a slightly flat , `EOT` the little the seams `EOT` `EOT` it drags `EOT` perverse `EOT` this bland `EOT` the cheapness `EOT` poorly `EOT` any teen movie starring slackers or `EOT` succeeded `EOT` brain strain `EOT` ill writing `EOT` may also not have been the grossest american bike movie of 2002 . `EOT` absurdities `EOT` it 's too simpleminded – `EOT` has the idea been sealed in a jar `EOT` this obvious rip-off `EOT` racist `EOT` is wit, or even good acting , `EOT` more like satire than illuminating examination `EOT` despite snow dogs, leaves us cold . `EOT` the action is predictable with many subplots and `EOT` very insecure `EOT` 's still flibbled together its clichés . `EOT` bad combination `EOT` particularly suspenseful at times `EOT` deep technical blues vs blues `EOT` are so smeary, blurry and jagged `EOT` form 51 is n't really scary enough for " midnight flick ", and `EOT` thought you or i were paying full price for this, and rent those instead `EOT` is a minor miracle `EOT` handbags in saving private ryan `EOT` like ' to suck ' `EOT` its clichéd scenes of war `EOT` poorly dubbed dialogue, and `EOT` 's hard to imagine a product more repellent than adam sandler 's 2002 `EOT` `EOT` weighted a sci-fi story that 's bottom-heavy and tired `EOT` wrong things `EOT` stiflingly `EOT` `EOT` endless exposition `EOT` simply overcooked `EOT` evasive `EOT` call an " ambitious failure `EOT` `EOT` even less ambitious failure `EOT` ridiculous and sloppy `EOT` no air conditioning `EOT` the ragged pacing `EOT` a loud, crassly predictable rehash `EOT` how lame it all is `EOT` ends world traveler 's lackluster title `EOT` have done with all the subtlety from earlier had taken in favor of more user-friendly computer tutorial `EOT` `EOT` mind settles into a irksome parade of human frailty and death . `EOT` smug and exploitative `EOT` the whining `EOT` this latest comedy tangents of pratfalls, injuries and pranks `EOT` to anyone who suffered through halfway through david rifkin 's sweet home ' alabama `EOT` old pickup `EOT` should have been ordered off the screen `EOT` like the plague `EOT` neges toward fuzziness `EOT` its skewed acting , `EOT` an odd, painful and obnoxious attempt at a documentary `EOT` , repetitious and contrived `EOT` lacks a moment of considerable poignancy . `EOT` chaos and `EOT` `EOT` cold, emotionally opaque `EOT` stands nowhere `EOT` frustrating `EOT` dull `EOT` suffered `EOT` characteristic quirks `EOT` half-step `EOT` , it 's hopeless . `EOT` there 's nothing new to be taken from ' fatal attraction ' 4ever `EOT` is painfully bad . `EOT` unfaithful hollywood fluff `EOT` have little interest in jez begley 's book `EOT` no foundation for disney `EOT` `EOT` , herzog is the director 's logistically and `EOT` is it not hagiographic, though `EOT` need more than plot `EOT` rocks and holes `EOT` with lots of solips `EOT` ving `EOT` at times maddeningly repetitive . `EOT` rendered all round square edges, blurry images and murky and `EOT` 's cloying, manipulative `EOT` trial `EOT` on sub-zero animation `EOT` an annoying , `EOT` went to the restroom but want my money back . `EOT` water torture `EOT` pompous `EOT` rigid idea `EOT` insufferable `EOT` as boring and pedestrian `EOT` , heavy-handed `EOT` marginal insights `EOT` from lost ring to snatch on stage `EOT` hampered by a shabby script `EOT` gold is a real hollywood dog story `EOT` limpid and pretentious endeavor `EOT` is subtler than it might have been . `EOT` the film with relentlessly nasty situations `EOT` 's not handled well `EOT` wastes dialogue `EOT` sensible violence `EOT` of all the male junkie post-about obsessive relationships `EOT` is a testament to the author 's work `EOT` have had room for more creative action . `EOT` empty farce `EOT` liar zone `EOT` could call it tacky `EOT` painfully awful `EOT` forget about the fact that kennedy has nothing to offer up his act `EOT` the faintest hopefulness , `EOT` `EOT` flat `EOT` , spiteful idiots `EOT` is superficial and `EOT` describe characters ' frustrations `EOT` end zone `EOT` familiar and thoroughly recycled plot . `EOT` tired `EOT` aimless direction, pathetic acting, heavy dialogue, tissue-thin characters `EOT` `EOT` of inept filmmaking `EOT` last 15 years `EOT` the picture just `EOT` nothing except they ca n't – anything, really, seems to matter very much `EOT` a noticeable amount of its own `EOT` neat `EOT` best elsewhere `EOT` waste ; `EOT` human behavior , `EOT` too fancy `EOT` derivancy `EOT` distinct comic direction `EOT` this m-out-l lesbian comedy `EOT` does n't give a complete picture . `EOT` shallow plot `EOT` could have any interest in silly fluff `EOT` , tedious , `EOT` 's not clear what this turkey is

Figure 11: Sentiment control - negative samples - SMC $\beta = 1.0$, $\alpha = 2.0$. `EOT` indicates End Of Text

or it really costs that far to finish rail. `EOT` I always thought the part time NDP would be on the hook for this, but then they had to master the ndp's fiscal policies. `EOT` Fess up........ `EOT` He managed it brilliantly. `EOT` You are wrong in that regard. Harper suggested enough to leave the potential harm to children (and adults) to see. `EOT` Agreed. Further increasing the income of Quebec will likely lead to increased income tax charges for Ontario governments attempts to do same. `EOT` rumps.. the globe loses again.. its only readership is NK `EOT` once again putting alt lunatic whining over the real picture bud `EOT` Alaska hasn't cut the PF yet Walker over spent our PFD for the same 40 years. Walker has 3 more session to do some cuts cutting back PFD dividends not PFD raiding the new gas line. Enough Permanent fund cut already Walker and Governor Walker and Governor Walker. Time for legislature and real leadership in Alaska needing new taxes. `EOT` Yes Bravo Orange Pail! `EOT` Sorry son but don't look forward to Ducks game, it's always good to find out who has the better quarterback. `EOT` Trump is fleeing the Trump administration right now. `EOT` Sorry, you've undermined the whole point of your comment ... `EOT` How about everyone familiar with Dale Wright? `EOT` One less Hawaii Democrat. `EOT` Where's Cory Gardner. `EOT` Silly Senate. `EOT` Happy New Air Canada - what a touch of sweet irony. `EOT` Great comments! `EOT` Seems like your actions out in front of the world for all people to see. `EOT` there is no illegal world building `EOT` Trudeau ? `EOT` Not to fast.....tax Vancouver got it fair share. `EOT` LOL....... `EOT` Maybe our Senators Cory and Murkowski will vote for a repeal. `EOT` This is about criminal charges underway and emails continue to be the target. Mueller will investigate the names Hillary, let Mueller have his look and hopefully does find Hillary exists. `EOT` I'm sure snakes are the least human kind. `EOT` Great pleasant surprise. `EOT` I agree with you! `EOT` The problem is that guys like Meredith and Crapwell will just fade away. `EOT` Alceste three times is running as a republican. `EOT` Wrong case. Can't find a place right there. `EOT` http://JustPeace.org That's Fairbanks, Fairbanks `EOT` Yep. His opening rhinos ring louder every posting. `EOT` Not China nor NK. Neither have threatened NK. `EOT` Maybe they break the law along with other citizens who drive motorized vehicles. `EOT` What were you wearing? Wear the ladies locker? `EOT` Get a clue who play each string on the violin. `EOT` I have been discouraged about owning a home when rental prices have gone down significantly. It seems like there is only reason for gray market rentals any more. If I could find an area that had a small amount of rental units and never close it with no decision but rent, then I could buy in. `EOT` More Liberal tabloidism fleeced by paid leftist disinformation efforts. Good riddance! `EOT` Didn't happen to read your post either. `EOT` Wrong Old Spice.... `EOT` IRT Gary Crum, we don't need any new information when it comes to responding to the DP. `EOT` why? `EOT` Then why tax the PFD earnings when Anchorage borough and developers pay local property taxes with PFD earnings? `EOT` Agreed! Pay you to Netflix any network programming than you get in actual services you pay for. `EOT` You wanted Hillary is a lot different than wishing Hillary lost...... `EOT` Does every car require a plate handle or do democrats have ? `EOT` They have not given up anything re producing soft fishing. The dairy industry will not be on any side of the deal when NAFTA gets a bad looks up. First side to lose the deal is at risk. `EOT` Agreed. `EOT` I assume it isn't you Gary `EOT` Yeahhhh trump juice! `EOT` I remember long ago the Cons and the NDP were slamming each other behind the scenes. `EOT` Your tolerance must be getting really old. `EOT` Wouldn't say Amy hasn't heard here. `EOT` So Berkie and Caldwell belong in jail? `EOT` the big question is "which will they write it" `EOT` How was this deficit announced 6 months ago? `EOT` Very good article thanks for posting Vic. `EOT` Completely untrue. South America is producing the most efficient portion of the world's oil soon. BTW, if your planning on expanding pipelines, move ahead and avoid Alberta. `EOT` Wow. Can't even conceive of imagining how naïve trumpsters are. `EOT` Nah. It was a big one. `EOT` Prove it. Tax my dividend too. `EOT` Did't the letter writer say Lame County? `EOT` Yep! That's who you "democratic" voted into office. `EOT` `EOT` Agreed! Harper! Has to go. `EOT` Sorry but not quite. The data you are referencing is one game show. `EOT` So Alaska needs our own modern rail system? We don't need a modern system of satellites like the majority of lights on the nation. `EOT` I find it hilarious that idiot trump does nothing to mitigate

Figure 12: Toxicity control - non toxic samples - SMC $\beta = 1.2$, $\alpha = 4.0$. `EOT` indicates End Of Text

when commuters were most busy......... EOT Eh. Used that logic. EOT Better hope the coming recession never happens EOT Enough with your infomeration Bombing homes works better. EOT Just keep Trump in. It could be our destiny to get rid of all Obama's puppet masters EOT Here's my only thought regarding your comment: Keep up! EOT How did I need any further contribution to that? EOT Someone please refute my fact. EOT The phony Liberals who "thought" a love of country above party will have no spine. As always, the press will focus on facts and Republicans will relentlessly call out real Liberalism. EOT That's a fact. Comey's firing does have been the direct result of Comey's conduct acting inappropriately for US intelligence. Comey is not the person fired out of political motivation, but Trump was rejected as for the US Criminal Code but rather the federal judge has rejected him. EOT ": An extension for the appeal period or even an extension for each waiver signed may be all Kobachak's congressional breaks." As I pointed out in a Civil Comments form on behalf of Kobachak in 2016 there was nothing further to see in Kobachak's case. He picked a fight and all is well with his formerly a non-friend and ally. EOT So? The author is anonymous? at least he doesn't have a secret. BTW, you seem to assume that Kobachak doesn't share your own psychology. Darn it. EOT I am a retired member of the Sierra Club, although I am also a member of the California chapter at https://fight.org/ and the Northeast chapter at https://www.sanfrust.org/ The health issue in Puerto Rico is significant and I think about it from a personal perspective. Trump is really sparking passions in his comments so will face criticism. Aside from that aways, he needs to understand what actually took place in Texas City and events that resulted in this disaster there. EOT Our government is only better than you. EOT Now that Trudeau has been given free rein in unilateral withdrawal from the U.S., then why would we subject ourselves to further a significant trade relationship that the Chinese and Russians absolutely want? It can only work here if we remove big historical barriers such as the Suzuki's and tackle these for our benefit. The world is entering a new era of mutual respect. EOT This is more of exactly what's expected. EOT That's wording that selectively is used. Only a few extreme views dominate the discussion. These are only three of 10. EOT More that some people do disagree with the process as they may have alternative theories. I also suggest you get educated and understand the current science and the reality of climate change. EOT Aloha, go year long in different neighborhoods like season after year. I see day in and day out the exact same homeless; no different penalty for visiting Eugene; same as anywhere else. EOT Thanks for the linking C of B page to Quebec. Donna can't cut it. EOT The crime is not unusual, per say. They are allowed their privacy, but selectively. No tunnels to break, but it is a crime. EOT The IPCC failing to release public information regarding the case runs contrary to the article's recommendation. Quite clearly his fate is not to be ended in death as was indicated by the way he was raised when confronted with simply not being born. EOT Again, are you not proud of your experience and reach? EOT I agree. This analysis does not take into account that the content of most of our English-language newspapers is particularly bias. To what appears to be happening in Canada a few days, but every once in a while a Canadian newspaper is going to highlight certain facts because of the importance a story that few would know. It has also been said that many newspapers insist on showing all news every week and not just the stories that are the most important. EOT If i can looked at things differently and see more from the top - it emerges I see less as 'completely out of line' as others feel I have looked more towards the 'top'. EOT We are all doing the best we do, sometimes it takes the next thing. EOT Having a minimum wage for the smaller employees in Alaska must be expensive because the businesses have to do their work at home thru government, education and the overall economy. My monthly cost of organic food is \$10.71 per hour plus. So how can I go through the trouble and earn that many per hour rather than having a loss on these goods? I paid lucky so I will have to stock up while in retirement, but that is still very expensive. Remember you not only need the help with government but with income, earned through education, and the rest of the economy. EOT Sure, a lot of kids are taking hitches and missing their focus. The selection of character, athleticism and tougher focus is what determines the chance of success. EOT It's not everything. He did hate the west in his first election. EOT It's the First Amendment folks to blame ... I say to "vide wolf, flock"

Figure 13: Toxicity control - non toxic samples - Guidance $\beta = 1.2$. EOT indicates End Of Text