
Discriminator Guidance for Autoregressive Diffusion Models

Filip Ekström Kelvinius
Linköping University

Fredrik Lindsten
Linköping University

Abstract

We introduce discriminator guidance in the setting of Autoregressive Diffusion Models. The use of a discriminator to guide a diffusion process has previously been used for continuous diffusion models, and in this work we derive ways of using a discriminator together with a pretrained generative model in the discrete case. First, we show that using an optimal discriminator will correct the pretrained model and enable exact sampling from the underlying data distribution. Second, to account for the realistic scenario of using a sub-optimal discriminator, we derive a sequential Monte Carlo algorithm which iteratively takes the predictions from the discriminator into account during the generation process. We test these approaches on the task of generating molecular graphs and show how the discriminator improves the generative performance over using only the pretrained model.

1 INTRODUCTION

Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020) have in the last couple of years received significant attention, and many improvements and extensions of the original works have been proposed. A recent work by Kim et al. (2023) introduced *discriminator guidance*: a way of improving the score estimation of a pretrained score-based diffusion model by training a discriminator, and during generation combine the pretrained score model with the gradient of the discriminator to obtain an improved estimation of the score. Empirically, this approach improves sample quality.

Score-based models operate on continuous data. One

formulation of a diffusion model for *discrete* data is the Autoregressive Diffusion Model (ARDM) by Hoogetboom et al. (2022a) which builds on the Order-Agnostic Autoregressive Model (OA-ARM) by Uria et al. (2014). This is a significantly different formulation of a diffusion model compared to the score-based model, and it is therefore not straightforward to use the existing formulation of discriminator guidance for this type of model. However, given the improved performance obtained by incorporating discriminator guidance in the continuous setting, it is of interest to develop such techniques also for the ARDM.

To address this challenge, we formulate the following algorithms for incorporating a discriminator into the sampling from an ARDM:

- Autoregressive Discriminator Guidance (ARDG), which is the most similar to the continuous case by Kim et al. (2023): the conditional distribution predicted by the ARDM is corrected by a discriminator, which we show enables exact sampling under the (in practice unrealistic) assumption that we can find a perfect discriminator.
- Sequential Monte Carlo (SMC) algorithms which build on discriminator guidance, but aims at mitigating the accumulation of errors that can occur due to imperfect discriminators for the intermediate sampling steps. This is made possible by the sequential correction of intermediate target distributions through the propagation, weighting, and resampling steps of SMC.

Using these different approaches, we empirically verify that discriminator guidance improves the generative sample quality over regular ARDMs in the setting of generating molecular graphs. Our code is available online¹.

2 BACKGROUND

Consider the task of generating a D -dimensional vector $\mathbf{x} = (x_1, x_2, \dots, x_D)$, where each variable x_i is dis-

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹https://github.com/filipekstrm/graph_ardm




Figure 1: Illustration of the ARDG method when applied to graphs. The variables x_t (nodes and edges in the graph) are one by one assigned values, where nodes and edges with “?” correspond to variables which have not yet been assigned. Assignment of the white node could be done by sampling from the conditional distribution $p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$, which has been learnt by a neural network. In our method, however, a separate discriminator, d_ϕ , has been trained to distinguish between real and fake (generated by the generative model p_θ) samples. With the help of this discriminator, we can correct the distribution $p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$ so that it becomes closer to the true underlying data distribution $p_{\text{data}}(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$.

crete. To this end, we will use an Autoregressive Diffusion Model (ARDM) (Hooeboom et al., 2022a) in combination with a discriminator. This section therefore introduces background on the ARDM, discriminator guidance for *continuous* data, and some additional related work.

We write $p_{\text{data}}(\mathbf{x})$ for the true data distribution, and $p_\theta(\mathbf{x})$ for the generative model, which is effectively an estimate of $p_{\text{data}}(\mathbf{x})$ constructed using neural networks with parameters θ .

2.1 Autoregressive Diffusion Models

To introduce ARDMs, we start from the standard autoregressive model (ARM). For an ARM, the data log-likelihood $\log p_\theta(\mathbf{x})$ is factorized as

$$\log p_\theta(\mathbf{x}) = \sum_{t=1}^D \log p_\theta(x_t|\mathbf{x}_{1:t-1}), \quad (1)$$

where $\mathbf{x}_{1:t-1} := (x_1, \dots, x_{t-1})$. Under the autoregressive model, data can be generated via ancestral sampling, i.e., sample the variables x_t conditioned on the previously sampled values $\mathbf{x}_{1:t-1}$.

However, the autoregressive model assumes a certain order of the variables x_t . In the *order agnostic* autoregressive model (OA-ARM) (Uria et al., 2014), the order is not viewed as fixed, but a random variable drawn from the set of all permutations S_D of the indices $\{1, \dots, D\}$. Denoting the permutation as $\sigma = (\sigma(1), \dots, \sigma(D))$, the data log-likelihood for OA-

ARM, given the order σ , is written as

$$\log p_\theta(\mathbf{x}|\sigma) = \sum_{t=1}^D \log p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}). \quad (2)$$

Here, and in what follows, we use the shorthand notation $\sigma(<t) := (\sigma(1), \dots, \sigma(t-1))$ for brevity, to denote the first $t-1$ elements of the permutation. The full data log-likelihood can then be written as the expectation over all orders

$$\log p_\theta(\mathbf{x}) = \log \mathbb{E}_\sigma [p_\theta(\mathbf{x}|\sigma)]. \quad (3)$$

To formulate ARDM, Hooeboom et al. (2022a) use the OA-ARM formulation in a diffusion model context: by having a generative process according to Equation (2), i.e., pick an order randomly and then one by one assigning values to the variables (or “unmasking” them) according to this order, the noising process corresponds to randomly “masking out” variables one by one.

An OA-ARM or ARDM can then be trained by maximizing the log-likelihood as

$$\begin{aligned} \log p_\theta(\mathbf{x}) &= \log \mathbb{E}_\sigma [p_\theta(\mathbf{x}|\sigma)] = \mathbb{E}_\sigma [\log p_\theta(\mathbf{x}|\sigma)] \\ &= \mathbb{E}_\sigma \left[\sum_{t=1}^D \log p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}) \right] \\ &= \mathbb{E}_\sigma [D \mathbb{E}_t [\log p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})]] \\ &= D \mathbb{E}_t [\mathbb{E}_{\sigma(<t)} [\mathbb{E}_{\sigma(t)|\sigma(<t)} [\log p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})]]]. \end{aligned} \quad (4)$$

It can be noted that, contrary to Uria et al. (2014) and Hooeboom et al. (2022a), we have not relied on Jensen’s inequality for moving the log inside the

expectation (first equality). The reason for this is that a different σ merely leads to a different factorization of the same likelihood, and hence, $\log p_\theta(\mathbf{x}|\sigma) = \log p_\theta(\mathbf{x})$, $\forall \sigma$.²

As σ is drawn uniformly from the set of all permutations, the inner-most expectation can be rewritten as a sum over all masked elements, and training an ARDM is done by maximizing the log-likelihood

$$\log p_\theta(\mathbf{x}) = D\mathbb{E}_t \left[\mathbb{E}_{\sigma(<t)} \left[\frac{1}{D-t+1} \sum_{k \notin \sigma(<t)} \log p_\theta(x_k | \mathbf{x}_{\sigma(<t)}) \right] \right]. \quad (5)$$

2.2 Discriminator Guidance for Score-based Diffusion Models

In score-based diffusion models (Song and Ermon, 2019; Song et al., 2020), a neural network \mathbf{s}_θ is trained to estimate the score $\nabla \log p_{\text{data}}(\mathbf{x})$, where \mathbf{x} is now a continuous variable. The trained neural network \mathbf{s}_θ (implicitly) defines a distribution p_θ . If the learnt score-function \mathbf{s}_θ deviates from the real score $\nabla \log p_{\text{data}}$, there will be a gap between p_{data} and p_θ . By observing that

$$\nabla \log p_{\text{data}}(\mathbf{x}) = \mathbf{s}_\theta(\mathbf{x}) + \nabla \log \frac{p_{\text{data}}(\mathbf{x})}{p_\theta(\mathbf{x})}, \quad (6)$$

Kim et al. (2023) introduce a correction term $\nabla \log \frac{p_{\text{data}}(\mathbf{x})}{p_\theta(\mathbf{x})}$. This is not tractable, but they explain how it can be estimated by training a discriminator d_ϕ to discriminate between real and generated samples: if $d_\phi(\mathbf{x})$ can accurately predict the probability that a sample \mathbf{x} is drawn from the data distribution, then the ratio $\frac{d_\phi(\mathbf{x})}{1-d_\phi(\mathbf{x})}$ approximates $\frac{p_{\text{data}}(\mathbf{x})}{p_\theta(\mathbf{x})}$ and hence

$$\mathbf{c}_\phi(\mathbf{x}) := \nabla \log \frac{d_\phi(\mathbf{x})}{1-d_\phi(\mathbf{x})} \approx \nabla \log \frac{p_{\text{data}}(\mathbf{x})}{p_\theta(\mathbf{x})}. \quad (7)$$

This correction term is used together with \mathbf{s}_θ so that the estimated score is $\mathbf{s}_\theta(\mathbf{x}) + \mathbf{c}_\phi(\mathbf{x})$, which with a perfectly trained discriminator will correspond to the exact score $\nabla \log p_{\text{data}}(\mathbf{x})$.

2.3 Other Related Work

Discrete Diffusion Models: Some early works on discrete diffusion models are Multinomial Diffusion (Hoogeboom et al., 2021) and D3PM (Austin et al.,

2021). In both these cases, the noise process consists of the variables independently transitioning to different states. In D3PM, an additional noising process is introduced, absorbing state noise: at each time step, variables are independently and with a certain probability "decayed" into an "absorbed" state. ARDM is a continuous-time generalization of this process: in continuous time, at most one variable at a time will decay into the absorbed state, and the noise process hence boils down to finding the order in which the variables decay. If the variables have the same probability of decaying, the order in which they decay will be a draw from the uniform distribution over all permutations of the indices $\{1, \dots, D\}$. The reverse process is to "unmask" the absorbed variables in the reverse order in which they decayed. This reverse order will also be from the uniform distribution, like in OA-ARM.

In our work, we focus on models which operate on discrete state spaces. There are, however, formulations of diffusion models for discrete data which use *continuous* diffusion. Examples of these are EDM (Hoogeboom et al., 2022b), CDCD (Dieleman et al., 2022), and Bit Diffusion (Chen et al., 2023).

Diffusion Models for Graphs: In this work, we will apply our model to the task of generating graphs. A previous work by Niu et al. (2020) use continuous diffusion to generate adjacency matrices, and Jo et al. (2022) use continuous diffusion to generate graphs with node and edge attributes. DiGress (Vignac et al., 2023) is an extension of D3PM to graphs, which generates graphs with node and edge attributes and is the first diffusion model for graphs that use discrete noise.

Generative Adversarial Networks: Generative adversarial networks (GANs) (Goodfellow et al., 2014) are perhaps the class of generative models which is mostly associated with the term "discriminator" as training of GANs is done by training a generative model and discriminator simultaneously in an adversarial manner. In a standard GAN, the discriminator is merely a training artifact and not used for generation. However, follow-up work have found ways of using the discriminator also in the generative process by rejection sampling (Azadi et al., 2018), Metropolis-Hasting sampling (Turner et al., 2019), and Langevin sampling (Che et al., 2020).

SMC and Diffusion Models: Concurrent to our work, Wu et al. (2023); Cardoso et al. (2023) have proposed SMC-based methods for solving inverse problems with diffusion-based priors. The resulting algorithms show some resemblance to our SMC-based samplers, but differ both in that they consider continuous rather than discrete state spaces, and in that they focus on inverse problems and not discriminator guid-

²Put differently, the true likelihood is independent of σ and the introduction of σ as a latent variable only affects the training of the model through the stochastic approximation of the expectation in Equation (4).

ance.

3 DISCRIMINATOR GUIDANCE FOR ARDMS

To formulate a discrete discriminator guidance procedure, we assume we have a pre-trained generative model, p_θ , from which we can sample in an order-agnostic autoregressive manner, i.e.,

$$x_{\sigma(t)} \sim p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}). \quad (8)$$

We denote the real data distribution as p_{data} , and train a discriminator d_ϕ by maximizing

$$\mathcal{L}_\phi = \mathbb{E}_\sigma \left[\mathbb{E}_t \left[\mathbb{E}_{p_{\text{data}}(\mathbf{x}_{\sigma(\leq t)})} [\log d_\phi(\mathbf{x}_{\sigma(\leq t)})] + \mathbb{E}_{p_\theta(\mathbf{x}_{\sigma(\leq t)})} [1 - \log d_\phi(\mathbf{x}_{\sigma(\leq t)})] \right] \right]. \quad (9)$$

Note that the discriminator is assumed to accept any partial sample $\mathbf{x}_{\sigma(\leq t)}$, $t \in \{1, \dots, D\}$ as input, and return the probability that this is a partially masked sample from the data distribution.

For a fixed generative model p_θ , training the discriminator d_ϕ by minimizing Equation (9) has the optimum $d_\phi^*(\mathbf{x}_{\sigma(\leq t)}) = \frac{p_{\text{data}}(\mathbf{x}_{\sigma(\leq t)})}{p_{\text{data}}(\mathbf{x}_{\sigma(\leq t)}) + p_\theta(\mathbf{x}_{\sigma(\leq t)})}$ (Goodfellow et al., 2014). If we then define

$$W_t(\mathbf{x}_{\sigma(\leq t)}) := \frac{d_\phi(\mathbf{x}_{\sigma(\leq t)})}{1 - d_\phi(\mathbf{x}_{\sigma(\leq t)})}, \quad (10)$$

simple algebra gives, for an optimal discriminator d_ϕ^* ,

$$W_t^*(\mathbf{x}_{\sigma(\leq t)}) = \frac{d_\phi^*(\mathbf{x}_{\sigma(\leq t)})}{1 - d_\phi^*(\mathbf{x}_{\sigma(\leq t)})} = \frac{p_{\text{data}}(\mathbf{x}_{\sigma(\leq t)})}{p_\theta(\mathbf{x}_{\sigma(\leq t)})}. \quad (11)$$

Note that, if we parameterize the discriminator as $d_\phi(\mathbf{x}_{\sigma(\leq t)}) = \text{sigmoid}(f_\phi(\mathbf{x}_{\sigma(\leq t)})) = \frac{\exp(f_\phi(\mathbf{x}_{\sigma(\leq t)}))}{1 + \exp(f_\phi(\mathbf{x}_{\sigma(\leq t)}))}$, it follows that

$$W_t(\mathbf{x}_{\sigma(\leq t)}) = \exp(f_\phi(\mathbf{x}_{\sigma(\leq t)})). \quad (12)$$

In the following sections, we outline our sampling algorithms which combine the pretrained model p_θ and the discriminator d_ϕ .

3.1 Autoregressive Discriminator Guidance

We first derive a discriminator guidance algorithm which is a discrete counterpart of the continuous case (i.e., Kim et al. (2023)). To this end, we start by rewriting

$$p_{\text{data}}(\mathbf{x}_{\sigma(\leq t)}) = p_{\text{data}}(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}) p_{\text{data}}(\mathbf{x}_{\sigma(<t)}). \quad (13)$$

Algorithm 1 Autoregressive Discriminator Guidance

Sample $D \sim p(D)$ {For samples with varying number of elements (see Section 3.3)}

Sample $\sigma \sim p(\sigma)$

Initialize $\mathbf{x}_0 = \emptyset$ {Completely masked}

for t in $1 : D$ **do**

$C_t = 0$ {Normalization constant}

for each value of $x_{\sigma(t)}$ **do**

Compute $d_\phi(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})$

Compute $W_t(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})$ (Equation (15))

$C_t += W_t(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\}) p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})$

end for

Sample

$$x_{\sigma(t)} \sim C_t^{-1} W_t(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\}) p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})$$

end for

This can be done similarly for $p_\theta(\mathbf{x}_{\sigma(\leq t)})$, and with this, we get

$$\begin{aligned} W_t^*(\mathbf{x}_{\sigma(\leq t)}) &= \frac{p_{\text{data}}(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}) p_{\text{data}}(\mathbf{x}_{\sigma(<t)})}{p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}) p_\theta(\mathbf{x}_{\sigma(<t)})} \\ &= \frac{p_{\text{data}}(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})}{p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})} W_{t-1}^*(\mathbf{x}_{\sigma(<t)}) \\ &\iff \\ p_{\text{data}}(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}) &= \frac{W_t^*(\mathbf{x}_{\sigma(\leq t)})}{W_{t-1}^*(\mathbf{x}_{\sigma(<t)})} p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}). \end{aligned} \quad (14)$$

As $x_{\sigma(t)}$ is discrete, the distribution on the right hand side of Equation (14) can be computed as: for each value of $x_{\sigma(t)}$, evaluate the discriminator $d_\phi^*(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})$, compute $W_t^*(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})$, and multiply this with $p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})$. The denominator $W_{t-1}^*(\mathbf{x}_{\sigma(<t)})$ does not depend on $x_{\sigma(t)}$, and can therefore be implicitly computed by normalizing the probabilities $W_t^*(\mathbf{x}_{\sigma(\leq t)}) p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})$. Hence, with a perfectly trained discriminator, we can now at each time step sample from the data distribution, $p_{\text{data}}(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}) \propto W_t^*(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\}) p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)})$. In other words, we use the discriminator to correct the imperfect predictions of the intermediate conditionals made by p_θ . In practice, we will not have access to a perfect discriminator, but we can use the approximation

$$W_t(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\}) = \frac{d_\phi(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})}{1 - d_\phi(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})}, \quad (15)$$

to obtain a guided sampling process to approximately sample from p_{data} . We refer to this procedure as Autoregressive Discriminator Guidance (ARDG). The full procedure can be found in Algorithm 1, and is illustrated in Figure 1. In the algorithm we have assumed

Algorithm 2 Bootstrap particle filter discriminator guidance (BSDG). All operations for $i = \{1, \dots, N\}$

Sample $D \sim p(D)$ {For samples with varying number of elements (see Section 3.3)}
 Sample $\sigma \sim p(\sigma)$
 Set $\mathbf{x}_0^i = \emptyset$ and $w_0^i = 1/N$ {Completely masked}
for t in $1 : D$ **do**
 if ESS too low **then**
 Resample $\{\mathbf{x}_{\sigma(<t)}^i\}_{i=1}^N$ and set $w_{t-1}^i \equiv 1/N$
 end if
 Sample $x_{\sigma(t)}^i \sim p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}^i)$ and set $\mathbf{x}_{\sigma(\leq t)}^i = (\mathbf{x}_{\sigma(<t)}^i, x_{\sigma(t)}^i)$
 Compute $\tilde{w}_t^i = w_{t-1}^i W_t(\mathbf{x}_{\sigma(\leq t)}^i) / W_{t-1}(\mathbf{x}_{\sigma(<t)}^i)$
 Normalize $w_t^i = \tilde{w}_t^i / \sum_{j=1}^N \tilde{w}_t^j$
end for
 Sample $k \sim \text{Categorical}(N, \{w_D^i\}_{i=1}^N)$
return $\mathbf{x} = \mathbf{x}_\sigma^k$

that the generation order (permutation) can be sampled before the main loop: $\sigma \sim p(\sigma)$ (which is typically the case for ARDMs), but it is also possible to sample this one variable at a time $\sigma(t) \sim p(\sigma(t) | \sigma(<t))$ as part of the loop.

3.2 Sequential Monte Carlo

A possible drawback with ARDG is that (as any "plain" autoregressive model) it is unable to correct, at some later iteration, for mistakes made at earlier iterations during the generation. However, intuitively we might expect that the discriminator becomes more accurate for large t , since it then has access to more unmasked elements and thus more information. We therefore propose an extension to ARDG based on Sequential Monte Carlo (SMC; see, e.g., Naesseth et al. (2019)).

The motivation is that SMC can potentially be more resilient to an accumulation of errors. The reason is that SMC has a built-in correction mechanism, in the sense that "errors" in the intermediate target distributions of the algorithm (details below) are corrected for when transitioning from one iteration to the next. Therefore, it is sufficient that the final target distribution is error-free to obtain an algorithm that is consistent, i.e., asymptotically unbiased as the number of SMC samples ("particles") increases (Naesseth et al., 2019). In our setting, this means that as long as the discriminator $d_\phi(\mathbf{x})$ which has access to a complete, unmasked sample is (a close approximation to) the optimal classifier, the guided sampling process will asymptotically generate samples from (a close approximation of) the data distribution.

In practice we view the number of SMC particles N

Algorithm 3 Fully adapted SMC discriminator guidance (FADG). All operations for $i = \{1, \dots, N\}$

Sample $D \sim p(D)$ {For samples with varying number of elements (see Section 3.3)}
 Sample $\sigma \sim p(\sigma)$
 Set $\mathbf{x}_0^i = \emptyset$ and $w_0^i = 1/N$ {Completely masked}
for t in $1 : D$ **do**
 Compute $C_{t-1}^i = \sum_{x_{\sigma(t)}} \frac{W_t(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(<t)}^i)}$
 Compute $\tilde{w}_t^i = w_{t-1}^i C_{t-1}^i$
 Normalize $w_t^i = \tilde{w}_t^i / \sum_{j=1}^N \tilde{w}_t^j$
 if ESS too low **then**
 Resample $\{\mathbf{x}_{\sigma(<t)}^i, C_{t-1}^i\}_{i=1}^N$ and set $w_t^i \equiv 1/N$
 end if
 Sample

$$x_{\sigma(t)}^i \sim \frac{1}{C_{t-1}^i} \frac{W_t(\{\mathbf{x}_{\sigma(<t)}, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(<t)}^i)} p_\theta(x_{\sigma(t)} | \mathbf{x}_{\sigma(<t)}^i)$$

 and set $\mathbf{x}_{\sigma(\leq t)}^i = (\mathbf{x}_{\sigma(<t)}^i, x_{\sigma(t)}^i)$
 end for
 Sample $k \sim \text{Categorical}(N, \{w_D^i\}_{i=1}^N)$
return $\mathbf{x} = \mathbf{x}_\sigma^k$

as a tuning parameter that can be used to trade computational cost for improved accuracy. Importantly, even in the extreme case $N = 1$, the two versions of SMC that we propose below will reduce to standard ARDM and ARDG, respectively. We therefore argue that even a small N can be used to improve the performance compared to the respective baselines.

In the derivation below we condition on a fixed generation order σ , but do not include this in the notation for brevity. SMC is a class of methods for sampling from a sequence of "target distributions" $\{\pi_t(\mathbf{x}_{\sigma(\leq t)})\}_{t=1}^D$. This is done by sequentially generating a collection of N interacting samples (or particles) that can be seen as approximate draws from the targets. The target distributions are assumed to be known only up to normalization, such that

$$\pi_t(\mathbf{x}_{\sigma(\leq t)}) = \frac{\gamma_t(\mathbf{x}_{\sigma(\leq t)})}{Z_t}$$

where $\gamma_t(\mathbf{x}_{\sigma(\leq t)})$ can be evaluated but Z_t may be intractable. In many cases—including the application studied in this paper—the actual distribution of interest corresponds to the final target distribution $\pi_D(\mathbf{x}_{\sigma(\leq D)})$ and the *intermediate* target distributions, $\{\pi_t\}_{t=1}^{D-1}$ can be viewed as auxiliary quantities. Importantly, as mentioned above, the algorithm is consistent under weak assumptions regardless of the choice of these intermediate targets, which can thus be seen as design variables.

In what follows we will derive two versions of SMC-based discriminator guidance. Both are based on the same sequence of unnormalized target distributions,

$$\gamma_t(\mathbf{x}_{\sigma(\leq t)}) = W_t(\mathbf{x}_{\sigma(\leq t)})p_\theta(\mathbf{x}_{\sigma(\leq t)}) \quad (16)$$

for $t = \{1, \dots, D\}$, where W_t is defined in Equation (10). This means that, at iterations $t = D$, samples are approximately from $p_{\text{data}}(\mathbf{x})$ according to Equation (11).

From the definition of γ_t in Equation (16) it follows that we can write

$$\gamma_t(\mathbf{x}_{\sigma(\leq t)}) = \frac{W_t(\mathbf{x}_{\sigma(\leq t)})p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})}{W_{t-1}(\mathbf{x}_{\sigma(<t)})}\gamma_{t-1}(\mathbf{x}_{\sigma(<t)}), \quad (17)$$

which can be interpreted as follows:

1. The conditional distribution $p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$ takes the role of a prior transition from iteration $t - 1$ to t ,
2. The ratio $W_t(\mathbf{x}_{\sigma(\leq t)})/W_{t-1}(\mathbf{x}_{\sigma(<t)})$ takes the role of a likelihood term correcting for the discrepancies between the target at iteration $t - 1$ and t .

Using this interpretation we can readily obtain the expressions for two commonly used versions of SMC adapted to our setting: the bootstrap SMC and the fully adapted SMC, respectively. Full derivations are given in the appendix.

3.2.1 Bootstrap Discriminator Guidance

We initialize a set of N particles $\mathbf{x}_0^i = \emptyset$ (completely masked) and corresponding *importance weights* $w_0^i = 1/N$. For a bootstrap SMC we use the prior transition, i.e., our pretrained generative model, as a proposal distribution to propagate samples from one iteration to the next. That is, at iteration t we sample

$$x_{\sigma(t)}^i \sim p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}^i), \quad i = \{1, \dots, N\}. \quad (18)$$

The samples are extended as $\mathbf{x}_{\sigma(\leq t)}^i = (\mathbf{x}_{\sigma(<t)}^i, x_{\sigma(t)}^i)$.

The importance weights are then updated using the "likelihood", i.e.,

$$w_t^i \propto \frac{W_t(\mathbf{x}_{\sigma(\leq t)}^i)}{W_{t-1}(\mathbf{x}_{\sigma(<t)}^i)}w_{t-1}^i, \quad (19)$$

and are normalized to sum to one (hence the proportionality sign).

A key concept in SMC is *resampling*, which allows the algorithm to focus on the more promising (high weight) samples and discard the less promising (low

weight) ones. This is done by monitoring the Effective Sample Size (ESS), defined as

$$\text{ESS} := \frac{1}{\sum_{i=1}^N (w_t^i)^2}. \quad (20)$$

If the ESS becomes too low (say, $\text{ESS} < N/2$), then we generate a new set of particles by sampling with replacement from the current set of particles, with probabilities given by the importance weights. After resampling, the importance weights are reset to be $1/N$.

The Bootstrap Discriminator Guidance (BSDG) procedure for discriminator guidance is summarized in Algorithm 2. Note that, in the extreme case $N = 1$, the procedure reduces to simply applying the pre-trained ARDM without any guidance. N can thus be seen as a tuning-parameter that can be used to improve the quality of generated samples at the cost of increased generation order, and even for $N \gtrsim 1$ we obtain a functional method.

3.2.2 Fully Adapted Discriminator Guidance

An alternative to bootstrap SMC is to "adapt" the transition probability for extending the particles as well as the resampling probabilities to the current target distribution. If this is done in a locally optimal way, the method is referred to as fully adapted SMC (Pitt and Shephard, 1999; Naeseth et al., 2019).

Similarly to how Equation (18) is viewed as sampling from the prior, the locally optimal transition corresponds to sampling from a posterior. With our interpretation of the weight ratio as a likelihood, this means that we sample $x_{\sigma(t)}^i$ from a distribution proportional to

$$\frac{W_t(\{\mathbf{x}_{\sigma(<t)}^i, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(<t)}^i)}p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}^i), \quad (21)$$

i.e., the product of the likelihood and the prior. Note that this is the same transition probability as used by ARDG; cf Equation (14).

Using this transition kernel, the importance weights are updated as

$$w_t^i \propto w_{t-1}^i \times \sum_{x_{\sigma(t)}} \frac{W_t(\{\mathbf{x}_{\sigma(<t)}^i, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(<t)}^i)}p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}^i), \quad (22)$$

where the sum is the normalizing constant ("marginal likelihood") of Equation (21). Similarly to above, the weights are normalized to sum to one.

Considering Equation (22) it can be noted that the importance weights are independent of the realized values

$\{x_{\sigma(t)}^i\}$ and only depend on particles up to iteration $t - 1$. This is a consequence of the perfect adaptation of the transition kernel to the current target γ_t . However, this also opens up for a one-step look-ahead in the resampling step of the fully adapted SMC. That is, we resample the particles from iteration $t - 1$ using the weights in Equation (22) *before* extending the particles to iteration t according to Equation (21). When the resampling is performed in this way, we obtain new importance weights $w_t^i = 1/N$; if we do not resample at some iteration we instead compute weights according to Equation (22).

The resulting algorithm for Fully Adapted Discriminator Guidance (FADG) can be found in Algorithm 3. In the extreme case $N = 1$ the method reduces to the basic ARDG procedure (Algorithm 1). Hence, similarly to how BSDG can be seen as a generalization of (unguided) ARDM, the FADG can be seen as a generalization of ARDG.

3.3 Generating Samples of Varying Dimension

We will use ARDM with discriminator guidance for generating graphs. These can have different number of nodes, n , which in turn means that the dimension D of generated samples is random. Hence, we are interested in generating from the joint distribution $p(\mathbf{x}, D)$. Similarly to DiGress (Vignac et al., 2023), we enable this by using the factorization $p_{\theta}(\mathbf{x}, D) = p(D)p_{\theta}(\mathbf{x}|D)$, where we set $p(D)$ to the empirical distribution computed from the training set. Specifically, in the case of graphs we sample first the number of nodes n from its empirical distribution, and then, for the case of undirected graphs without self-loops, compute $D = n + n(n - 1)/2$, corresponding to the total number of nodes and edges (see Section 4 for details). Formally we condition on D in the generative procedures presented above, i.e., we target the distribution $p(\mathbf{x}|D)$ instead of $p(\mathbf{x})$, but we omit this conditioning for notational simplicity.

Furthermore, in the SMC algorithms we condition on the permutation σ , so that the generation order is shared between particles. That is, we first simulate $D \sim p(D)$ and $\sigma \sim p(\sigma|D)$ and apply SMC to target $p_{\theta}(\mathbf{x}|D, \sigma)$. To generate a single sample we then randomly select one of the N particles with probabilities given by their importance weights.³ To generate multiple samples we repeat this entire procedure, sampling a new pair (D, σ) for each round.

³Alternatively, depending on the application, we can keep all N samples and average over them when computing expectations w.r.t. the generative distribution, resulting in a reduction of Monte Carlo variance.

Table 1: Evaluation metrics on the QM9 dataset, when using different generation orders and types of Discriminator Guidance (DG). ARDM is a standard autoregressive diffusion model (no guidance) and ARDG, BSDG, and FADG are the three methods proposed in this paper. ARDM* is a standard ARDM that has been trained for twice as long to match the extra training time required for training the discriminator.

	Model	Val. (%) \uparrow	Uniq. (%) \uparrow	Atm.S (%) \uparrow	Mol.S (%) \uparrow
	Test data	97.9	100	98.6	87.4
	DiGress	95.4	97.6	98.1	79.8
Uniform	ARDM	88.0	99.9	96.1	55.7
	ARDM*	89.7	99.6	96.0	55.6
	ARDG	88.3	99.7	96.7	66.3
	BSDG	97.4	99.4	98.6	87.8
	FADG	96.7	99.5	98.7	88.1
	NEsN	ARDM	95.4	99.9	97.3
ARDM*		96.8	99.9	97.6	78.5
ARDG		95.3	99.9	97.4	76.9
BSDG		98.2	99.9	98.6	88.4
FADG		98.1	99.5	98.6	87.8
NsEs	ARDM	92.5	99.6	97.4	77.3
	ARDM*	95.4	99.7	98.0	82.4
	ARDG	95.0	99.9	97.7	81.0
	BSDG	97.4	99.7	98.7	88.6
	FADG	97.2	99.4	98.6	88.0

3.4 Computational Complexity

Adding discriminator guidance will inevitably add an additional computational cost. ARDM requires D evaluations of the generative model. ARDG requires evaluating the discriminator for each possible value of $x_{\sigma(t)}$, involving $(1 + d)D$ network evaluations, with d the dimension of $x_{\sigma(t)}$ (i.e., for each t , one evaluation of the generative model, and then d evaluations of the discriminator). FADG, which samples from the same proposal as ARDG, will add to the computational complexity by a factor of the number of particles, N , i.e., it requires $N(1 + d)D$ forward passes through the networks.

BSDG on the other hand use the standard ARDM as proposal, and just a single forward pass through the discriminator per sample step t for weighting the particles. This means that it requires $2ND$ evaluations of the networks. For high dimensional $x_{\sigma(t)}$, the number of network evaluations in BSDG can thus be lower than that of standard ARDG if $N < \frac{(1+d)}{2}$.

Table 2: Evaluation metrics on the MOSES dataset, when using different types of Discriminator Guidance (DG). For results for different generation orders, see the appendix.

Order	DG	Validity (%) \uparrow	Uniqueness (%) \uparrow	Novelty (%) \uparrow	Filters (%) \uparrow	FCD \downarrow	SNN \uparrow	Frag \uparrow	Scaf \uparrow
Uniform	ARDM	82.2	100	97.2	94.9	2.937	0.483	0.993	0.067
	ARDM*	82.6	100	97.0	95.6	3.153	0.494	0.991	0.050
	ARDG	80.5	100	95.0	95.3	2.705	0.502	0.994	0.127
	BSDG	85.9	100	92.4	98.4	2.609	0.560	0.993	0.059
	FADG	90.1	100	91.7	98.4	2.537	0.541	0.995	0.105

4 EXPERIMENTS

The main objective with our experiments is to highlight how incorporating discriminator guidance improves the generative performance of the ARDM. Therefore, we have not put any effort into improving any architectural aspects of the backbone neural network, but use the same graph transformer as DiGress (Vignac et al., 2023) for both the generative model and the discriminator. Details on the architecture as well as details about training are given in the appendix.

As our model is working on partially masked graphs, we cannot compute the extra features used by Vignac et al. (2023), as we will not have access to, e.g., the graph Laplacian. On the other hand, we evaluate incorporating a recent approach by Ekström Kelvinius and Lindsten (2023) where the order in the ARDM is not drawn from a uniform distribution. Instead, the graphs are generated either by first assigning the values of the nodes in a random order, and then the edges (called NsEs), or by always assigning the edges connecting the most recently assigned node with the already assigned nodes (called NEsN).

The evaluation metrics are computed on a sample of 1 000 molecules generated by the models. In the SMC algorithms, we use $N = 10$ particles per sample, meaning we effectively have to generate 10 000 molecules, but only keep 1 000 for evaluation. We include a model ARDM*, which is standard ARDM with a training time comparable with

4.1 QM9

We first evaluate our methods on the QM9 dataset (Wu et al., 2018) and compare our results to DiGress (Vignac et al., 2023), which is a recent diffusion model for graphs and hence a very competitive baseline. As the metrics on the standard QM9 setup (no hydrogens) are already very good with a standard ARDM (see appendix), we turn our attention to the more difficult task of explicitly modeling hydrogens. We evaluate our model using the metrics

Validity (Val., fraction on molecules for which RDKit can obtain a valid SMILES string), Uniqueness (Uniq., fraction of molecules with a different SMILES representation), Atom stable (Atm.S, fraction of atoms with correct valency) and Molecule stable (Mol.S, fraction of molecules with 100 % atom stable). As can be seen in Table 1, standard ARDM with a uniform order performs slightly worse on all metrics when compared to DiGress, but using a non-uniform order (NEsN or NsEs) improves performance. Most importantly, however, adding discriminator guidance improves over ARDM for all choices of orders, and in particular the SMC algorithms obtain close to ideal metrics when compared to the dataset.

4.2 MOSES

Next, we evaluate ARDM with discriminator guidance on MOSES (Polykovskiy et al., 2020), a more challenging dataset of almost two million small drug-like molecules. This dataset is a lot larger than QM9, and we therefore do not generate a new dataset of the same size as the original dataset, but instead only 200 000 molecules, and then use random subsampling of these during training of the discriminator. As we do not see any direct differences in performance when using a different generation order, we present the results for the uniform order in Table 2, and refer to the appendix for results when using different generation orders.

MOSES is a standard benchmark with its own metrics, and apart from Validity and Uniqueness used for QM9, it also uses Novelty (fraction of the valid molecules not part of the training set), Filters (fraction of generated molecules which pass filters used in dataset construction), Fréchet ChemNet Distance (FCD) (Preuer et al., 2018) (distance between hidden activations of a pretrained ChemNet for generated and test set), SNN (average similarity between molecules in generated set and their nearest neighbor in test set), Fragment Similarity (Frag, comparison of distributions of BRICS fragments (Degen et al., 2008) in generated and test set), and Scaffold similarity (Scaff, comparison of frequencies of Bemis-Murcko scaffolds (Bemis and Mur-

cko, 1996) in generated and test set).

We observe that the different discriminator guidance methods improve many of the metrics over standard ARDM. However, it seems that they do so at a slight cost of novelty, i.e., they are slightly more prone to generating molecules from the training set. This could be interpreted as if the metrics are improved merely as a result of sampling more from the training set. However, this is not the case, since rerunning the evaluation using only novel molecules (see appendix), we observe the same improvements. It should also be noted that the metrics are computed over valid molecules, so even if the novelty is lower, the total number of valid *and* novel molecules can be higher.

5 DISCUSSION & CONCLUSIONS

In this paper, we have derived three discriminator guidance methods for ARDMs. First, we have a version (ARDG) that can be seen as a discrete counterpart of the continuous discriminator guidance by Kim et al. (2023). Second, we have two SMC versions (BSDG and FADG) that can further improve the generation quality by generating multiple interacting samples in parallel. Our empirical results when generating molecular graphs show how discriminator guidance improves the generative performance over standard ARDMs.

This improved performance naturally comes with an increased computational cost as the discriminator needs to be evaluated at each generation step. Additionally, for SMC we sample multiple particles in parallel for each generated sample. However, in contrast to other diffusion models, ARDMs have a fixed number of generation steps (D , one per variable). Choosing the number of particles N can therefore be seen as a tuning knob that allows trading generation quality for computational cost, similarly to the number of diffusion steps in standard diffusion models (both continuous and discrete models like D3PM and DiGress). In particular, for the extreme case $N = 1$, BSDG and FADG reduce to ARDM (no guidance) and ARDG, respectively, meaning that even for small N we obtain well-functioning methods (which is confirmed by our empirical results with $N = 10$).

We think that this trade-off can be of importance in applications or situations where it is more important that the obtained samples are of high quality, than being able to generate a large number of diverse samples with varying quality. For example, if the generated samples are to be further screened for some downstream application and this screening is expensive to perform, it is desirable that the samples that are chosen for screening are of higher quality, even if they are

relatively few. SMC provides the possibility to optimize for this requirement by increasing the number of particles and lowering the number of generated samples within a fixed computational budget.

Acknowledgments

This research is financially supported by the Swedish Research Council via the project *Handling Uncertainty in Machine Learning Systems* (contract number: 2020-04122), the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, and the Excellence Center at Linköping–Lund in Information Technology (ELLIIT). The computations were enabled by the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre.

References

- Austin, J., Johnson, D. D., Ho, J., Tarlow, D., and van den Berg, R. (2021). Structured Denoising Diffusion Models in Discrete State-Spaces. In *Advances in Neural Information Processing Systems*.
- Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. (2018). Discriminator Rejection Sampling. In *International Conference on Learning Representations*.
- Bemis, G. W. and Murcko, M. A. (1996). The Properties of Known Drugs. 1. Molecular Frameworks. *Journal of Medicinal Chemistry*, 39(15):2887–2893.
- Cardoso, G., el Idrissi, Y. J., Moulines, E., and Corff, S. L. (2023). Monte Carlo guided Denoising Diffusion models for Bayesian linear inverse problems. In *The Twelfth International Conference on Learning Representations*.
- Che, T., ZHANG, R., Sohl-Dickstein, J., Larochelle, H., Paull, L., Cao, Y., and Bengio, Y. (2020). Your GAN is Secretly an Energy-based Model and You Should Use Discriminator Driven Latent Sampling. In *Advances in Neural Information Processing Systems*, volume 33, pages 12275–12287. Curran Associates, Inc.
- Chen, T., Zhang, R., and Hinton, G. (2023). Analog Bits: Generating Discrete Data using Diffusion Models with Self-Conditioning. In *International Conference on Learning Representations*.
- Degen, J., Wegscheid-Gerlach, C., Zaliani, A., and Rarey, M. (2008). On the art of compiling and using ‘drug-like’ chemical fragment spaces. *ChemMedChem*, 3(10):1503–1507.
- Dieleman, S., Sartran, L., Roshannai, A., Savinov, N., Ganin, Y., Richemond, P. H., Doucet, A., Strudel,

- R., Dyer, C., Durkan, C., Hawthorne, C., Leblond, R., Grathwohl, W., and Adler, J. (2022). Continuous diffusion for categorical data.
- Ekström Kelvinius, F. and Lindsten, F. (2023). Autoregressive Diffusion Models with non-Uniform Generation Order. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc.
- Hoogeboom, E., Gritsenko, A. A., Bastings, J., Poole, B., van den Berg, R., and Salimans, T. (2022a). Autoregressive Diffusion Models. In *International Conference on Learning Representations*.
- Hoogeboom, E., Nielsen, D., Jaini, P., Forré, P., and Welling, M. (2021). Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions. In *Advances in Neural Information Processing Systems*.
- Hoogeboom, E., Satorras, V. G., Vignac, C., and Welling, M. (2022b). Equivariant Diffusion for Molecule Generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, pages 8867–8887. PMLR.
- Jo, J., Lee, S., and Hwang, S. J. (2022). Score-based Generative Modeling of Graphs via the System of Stochastic Differential Equations. In *Proceedings of the 39th International Conference on Machine Learning*, pages 10362–10383. PMLR.
- Kim, D., Kim, Y., Kwon, S. J., Kang, W., and Moon, I.-C. (2023). Refining Generative Process with Discriminator Guidance in Score-based Diffusion Models. In *Proceedings of the 40th International Conference on Machine Learning*, pages 16567–16598. PMLR.
- Naesseth, C. A., Lindsten, F., and Schön, T. B. (2019). Elements of sequential Monte Carlo. *Foundations and Trends in Machine Learning*, 12(3):307–392.
- Niu, C., Song, Y., Song, J., Zhao, S., Grover, A., and Ermon, S. (2020). Permutation Invariant Graph Generation via Score-Based Generative Modeling. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 4474–4484. PMLR.
- Pitt, M. K. and Shephard, N. (1999). Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association*, 94(446):590–599.
- Polykovskiy, D., Zhebrak, A., Sanchez-Lengeling, B., Golovanov, S., Tatanov, O., Belyaev, S., Kurbanov, R., Artamonov, A., Aladinskiy, V., Veselov, M., Kadurin, A., Johansson, S., Chen, H., Nikolenko, S., Aspuru-Guzik, A., and Zhavoronkov, A. (2020). Molecular Sets (MOSES): A Benchmarking Platform for Molecular Generation Models. *Frontiers in Pharmacology*, 11.
- Preuer, K., Renz, P., Unterthiner, T., Hochreiter, S., and Klambauer, G. (2018). Fréchet ChemNet Distance: A Metric for Generative Models for Molecules in Drug Discovery. *Journal of Chemical Information and Modeling*, 58(9):1736–1741.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. (2015). Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265. PMLR.
- Song, Y. and Ermon, S. (2019). Generative Modeling by Estimating Gradients of the Data Distribution. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2020). Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Turner, R., Hung, J., Frank, E., Saatchi, Y., and Yosinski, J. (2019). Metropolis-Hastings Generative Adversarial Networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6345–6353. PMLR.
- Uria, B., Murray, I., and Larochelle, H. (2014). A Deep and Tractable Density Estimator. In *Proceedings of the 31st International Conference on Machine Learning*, pages 467–475. PMLR.
- Vignac, C., Krawczuk, I., Siraudin, A., Wang, B., Cevher, V., and Frossard, P. (2023). DiGress: Discrete Denoising diffusion for graph generation. In *The Eleventh International Conference on Learning Representations*.
- Wu, L., Trippe, B. L., Naesseth, C. A., Blei, D., and Cunningham, J. P. (2023). Practical and Asymptotically Exact Conditional Sampling in Diffusion Models. In *Thirty-Seventh Conference on Neural Information Processing Systems*.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande,

V. (2018). MoleculeNet: A benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530.

Checklist

1. For all models and algorithms presented, check if you include:

- (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model.
Yes
- (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm.
Yes High-level description in main paper. Details in appendix.
- (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries.
After acceptance, code has been published on GitHub (see link on page 1)

2. For any theoretical claim, check if you include:

- (a) Statements of the full set of assumptions of all theoretical results.
Yes
- (b) Complete proofs of all theoretical results.
Not applicable. We do not provide any theorem statements in need of proofs. All theoretical results are derived in the main text.
- (c) Clear explanations of any assumptions.
Yes. Assumptions explained in main paper, with additional details in appendix.

3. For all figures and tables that present empirical results, check if you include:

- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL).
Yes
- (b) All the training details (e.g., data splits, hyperparameters, how they were chosen).
Yes, in appendix and code repository.
- (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times).
Yes, we provide additional results to quantify the variability of the obtained results in the appendix

(d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider).

Yes, in appendix

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:

- (a) Citations of the creator If your work uses existing assets.
Yes
- (b) The license information of the assets, if applicable.
Yes, when applicable in appendix
- (c) New assets either in the supplemental material or as a URL, if applicable.
Not Applicable
- (d) Information about consent from data providers/curators.
Yes, when applicable in appendix
- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content.
Not Applicable

5. If you used crowdsourcing or conducted research with human subjects, check if you include:

- (a) The full text of instructions given to participants and screenshots.
Not Applicable
- (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable.
Not Applicable
- (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation.
Not Applicable

A SMC DETAILS

In this section, we give a more formal derivation of the two SMC algorithms that we use for discriminator guidance.

As mentioned in the main text, SMC samples from a sequence of target distributions, $\{\pi_t(\mathbf{x}_{\sigma(\leq t)})\}_{t=1}^D$, where π_t is known up to a normalizing constants, i.e.,

$$\pi_t(\mathbf{x}_{\sigma(\leq t)}) = \frac{\gamma_t(\mathbf{x}_{\sigma(\leq t)})}{Z_t}. \quad (23)$$

To obtain these samples, SMC consists of three steps: resampling, propagating, and weighting. Each of the three steps are associated with its own additional SMC ingredient: the *ancestor probabilities* ν_{t-1} (used in the resampling step), the *proposal distributions* $q_t(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$ (used in the propagation step), and the *importance weights* w_t (used in the weighting step). The latter can be obtained by first computing the unnormalized importance weights \tilde{w}_t as

$$\tilde{w}_t = \omega_t(\mathbf{x}_{\sigma(\leq t)}) \frac{w_{t-1}}{\nu_{t-1}}, \quad (24)$$

where

$$\omega_t(\mathbf{x}_{\sigma(\leq t)}) = \frac{\gamma_t(\mathbf{x}_{\sigma(\leq t)})}{\gamma_{t-1}(\mathbf{x}_{\sigma(<t)})q(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})}, \quad (25)$$

and then normalizing by summing over all particles.

To design our SMC algorithms, we hence need to specify the target distributions γ_t , determine a proposal distribution, $q_t(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$, and how to compute the ancestor probabilities ν_{t-1} . From this, the importance weights will follow according to Equation (24).

We use $\gamma_t(\mathbf{x}_{\sigma(\leq t)}) = W_t(\mathbf{x}_{\sigma(\leq t)})p_\theta(\mathbf{x}_{\sigma(\leq t)})$, as this leads to, at iteration $t = D$, samples that are approximately from $p_{\text{data}}(\mathbf{x})$ (see Equation (11)).

A.1 Bootstrap Particle Filter (BSPF)

For designing a bootstrap particle filter for discriminator guidance, we can use our pretrained model as the proposal distribution, i.e.,

$$q_t(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}) = p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)}). \quad (26)$$

As in a standard BSPF, we set the ancestor sample weights equal to the importance weights, i.e.,

$$\nu_{t-1} = w_{t-1}. \quad (27)$$

With this choice of proposal and ancestor probabilities, we have that

$$\begin{aligned} \omega_t(\mathbf{x}_{\sigma(\leq t)}) &= \frac{\gamma_t(\mathbf{x}_{\sigma(\leq t)})}{\gamma_{t-1}(\mathbf{x}_{\sigma(<t)})q(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})} \\ &= \frac{W_t(\mathbf{x}_{\sigma(\leq t)})p_\theta(\mathbf{x}_{\sigma(\leq t)})}{W_{t-1}(\mathbf{x}_{\sigma(<t)})p_\theta(\mathbf{x}_{\sigma(<t)})p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})} \\ &= \frac{W_t(\mathbf{x}_{\sigma(\leq t)})p_\theta(\mathbf{x}_{\sigma(\leq t)})}{W_{t-1}(\mathbf{x}_{\sigma(<t)})p_\theta(\mathbf{x}_{\sigma(\leq t)})} \\ &= \frac{W_t(\mathbf{x}_{\sigma(\leq t)})}{W_{t-1}(\mathbf{x}_{\sigma(<t)})}. \end{aligned} \quad (28)$$

Then, the importance weights become

$$\tilde{w}_t = \frac{W_t(\mathbf{x}_{\sigma(\leq t)})}{W_{t-1}(\mathbf{x}_{\sigma(<t)})} \frac{w_{t-1}}{\nu_{t-1}}, \quad (29)$$

which, if resampling is performed with $\nu_{t-1} = w_{t-1}$, gives $\tilde{w}_t = \frac{W_t(\mathbf{x}_{\sigma(\leq t)})}{W_{t-1}(\mathbf{x}_{\sigma(\leq t)})}$. However, if resampling is not performed (i.e., if ESS is sufficiently large), this can be viewed as instead setting $\nu_{t-1} = 1/N$ and using a low variance sampler (for example, systematic). The expression in Equation (29) covers both cases.

A.2 Fully Adapted Particle Filter (FAPF)

In a fully adapted particle filter (Pitt and Shephard, 1999; Naesseth et al., 2019), the proposal should be the *locally optimal proposal*

$$q_t(x_{\sigma(t)}|\mathbf{x}_{\sigma(\leq t)}) = \frac{\gamma_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})}{\int \gamma_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})dx_{\sigma(t)}}. \quad (30)$$

For notational convenience when deriving the locally optimal proposal, we define

$$C_t(\mathbf{x}_{\sigma(\leq t)}) = \frac{\int \gamma_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})dx_{\sigma(t)}}{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})}. \quad (31)$$

By multiplying the right-hand side of the locally optimal proposal by $\frac{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})}{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})}$ and rearranging terms, we find that the locally optimal proposal is the same as the proposal in ARDG (Section 3.1), namely

$$\begin{aligned} q_t(x_{\sigma(t)}|\mathbf{x}_{\sigma(\leq t)}) &= \frac{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})}{\int \gamma_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})dx_{\sigma(t)}} \frac{\gamma_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})}{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})} \\ &= \frac{1}{C_t(\mathbf{x}_{\sigma(\leq t)})} \frac{W_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})p_\theta(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(\leq t)})p_\theta(\mathbf{x}_{\sigma(\leq t)})} \\ &= \frac{1}{C_t(\mathbf{x}_{\sigma(\leq t)})} \frac{W_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(\leq t)})} p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(\leq t)}). \end{aligned} \quad (32)$$

As $C_t(\mathbf{x}_{\sigma(\leq t)})$ is part of the normalizing constant of the proposal, and as $x_{\sigma(t)}$ is discrete, we can compute $C_t(\mathbf{x}_{\sigma(\leq t)})$ as

$$C_t(\mathbf{x}_{\sigma(\leq t)}) = \sum_{x_{\sigma(t)}} \frac{W_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})}{W_{t-1}(\mathbf{x}_{\sigma(\leq t)})} p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(\leq t)}). \quad (33)$$

In addition to using the locally optimal proposal, a fully adapted particle filter uses un-normalized ancestor probabilities which take into account possible values of future $x_{\sigma(t)}$ as (Pitt and Shephard, 1999)

$$\tilde{w}_{t-1} = w_{t-1} \frac{\int \gamma_t(\{\mathbf{x}_{\sigma(\leq t)}, x_{\sigma(t)}\})dx_{\sigma(t)}}{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})} = w_{t-1} C_t(\mathbf{x}_{\sigma(\leq t)}), \quad (34)$$

which can be normalized by summing over all particles.

Finally, the un-normalized importance weights are computed by using $\omega_t(\mathbf{x}_{\sigma(\leq t)})$, which with the locally optimal proposal becomes

$$\begin{aligned} \omega_t(\mathbf{x}_{\sigma(\leq t)}) &= \frac{\gamma_t(\mathbf{x}_{\sigma(\leq t)})}{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})q(x_{\sigma(t)}|\mathbf{x}_{\sigma(\leq t)})} \\ &= \frac{\gamma_t(\mathbf{x}_{\sigma(\leq t)})}{\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})} \frac{C_t(\mathbf{x}_{\sigma(\leq t)})\gamma_{t-1}(\mathbf{x}_{\sigma(\leq t)})}{\gamma_t(\mathbf{x}_{\sigma(\leq t)})} \\ &= C_t(\mathbf{x}_{\sigma(\leq t)}), \end{aligned} \quad (35)$$

resulting in

$$\tilde{w}_t = \frac{w_{t-1}}{\nu_{t-1}} C_t(\mathbf{x}_{\sigma(\leq t)}). \quad (36)$$

It can be noted that if resampling is performed, this is done using ν_t obtained by normalizing $\tilde{\nu}_t$ from Equation (34), leading to $\tilde{w}_t = 1$ and hence, $w_t = 1/N$. But just as for a BSPF, if ESS is still high enough and resampling is not performed in a step, this can be viewed as setting $\nu_t = 1/N$, meaning ν_t is not a normalization of $\tilde{\nu}_t$. We cover both these cases in Equation (36).

A.3 Experimental settings

In all SMC experiments, we used $N = 10$. We used systematic resampling and resampled when $\text{ESS} < 0.7N$.

B NETWORK AND TRAINING DETAILS

For both the generator and discriminator, we have used the graph transformer developed for DiGress. We have used the code from the the public DiGress code repository⁴. However, we do not use the extra features used for DiGress, as we do not have access to a full graph.

B.1 Generator

For the generator, we tried to closely follow the hyperparameters (e.g., number of layers and hidden dimensions) chosen for DiGress. When training, we used the Adam optimizer (while DiGress use AdamW), and tweaked the learning rate slightly depending on the dataset and generation order (Uniform, NEsN, or NEN). For QM9, we used 75 % of the full dataset as training data (roughly 98 000 molecules), slightly less than DiGress which used 100 000 molecules.

B.2 Discriminator

Architecture The discriminator has a backbone architecture that is identical to the generator (i.e., a graph transformer with the same number of layers and hidden dimensions). This backbone produces hidden features for each node and edge in the graph. In the generator, these features are used to predict the class logits in the generative distribution $p_\theta(x_{\sigma(t)}|\mathbf{x}_{\sigma(<t)})$. For the discriminator, however, we use mean pooling to obtain a pooled node feature, \mathbf{h}_{node} , and edge feature, \mathbf{h}_{edge} . That is, if letting V and E represent the set of nodes and set of edges, respectively, we obtain

$$\mathbf{h}_{\text{node}} = \frac{1}{|V|} \sum_{i \in V} \mathbf{h}_i \quad (37)$$

$$\mathbf{h}_{\text{edge}} = \frac{1}{|E|} \sum_{(i,j) \in E} \mathbf{h}_{(i,j)}. \quad (38)$$

These are concatenated and processed by an MLP to predict a single logit l , i.e.,

$$l = \text{MLP}(\mathbf{h}_{\text{node}} \oplus \mathbf{h}_{\text{edge}}), \quad (39)$$

with \oplus denoting concatenation of vectors. We used a single hidden layer in the MLP, with the hidden dimension being the same as the dimension of \mathbf{h}_{node} .

The logit can be converted into the probability of the sample being from the real dataset by applying a sigmoid. However, as shown in Equation (12), when computing the weights $W_t(\mathbf{x}_{\sigma(\leq t)})$, we use the logit directly as $W_t(\mathbf{x}_{\sigma(\leq t)}) = \exp(l)$ which works better numerically.

Training The already trained generative model is used to generate another dataset which is combined with the real dataset to train the discriminator. This generated dataset is of the same size as the original dataset in the case of QM9, and in the case of MOSES consists of 200 000 samples, which is $\sim 10\%$ the size of the original dataset.

We used the AdamW optimizer and tweaked the learning rate by observing cross entropy on a validation set which consisted of both real and generated data. The validation set was created by setting aside 15% of the dataset. We find empirically that initializing the discriminator with the weights of the generator helps with training, especially in the initial phase. We therefore used this strategy when initializing all discriminators.

⁴<https://github.com/cvignac/DiGress/>

B.3 Computational Resources

Both generators and discriminators were trained on single NVIDIA A100 40 GB GPUs.

C ADDITIONAL RESULTS

C.1 QM9 without hydrogen

For the task of generating molecules from QM9 when not modelling hydrogens, we provide the Validity and Uniqueness for standard ARDM, compared to DiGress, in Table 3.

Table 3: Evaluation metrics on the QM9 dataset, without modeling hydrogens. The metrics are very good, as we have comparable results with DiGress (Vignac et al., 2023).

Model	Validity (%) \uparrow	Uniqueness (%) \uparrow
DiGress	99.0	96.2
ARDM	99.1	96.9

C.2 MOSES with non-Uniform Generation Order

Just as for QM9, we evaluated the approach of Ekström Kelvinius and Lindsten (2023) where the distribution of generation orders, $p(\sigma)$, is not uniform over all possible orders. However, in contrast to QM9 where this approach gave significant boost in performance, this was not the case for MOSES. We do see, however, the same qualitative results that discriminator guidance consistently provides improved performance, given a choice of $p(\sigma)$. The full MOSES results can be found in Table 4

Table 4: Evaluation metrics on the MOSES dataset, when using different generation orders and types of Discriminator Guidance (DG).

Order	DG	Validity (%) \uparrow	Uniqueness (%) \uparrow	Novelty (%) \uparrow	Filters (%) \uparrow	FCD \downarrow	SNN \uparrow	Frag \uparrow	Scaf \uparrow
Uniform	ARDM	82.2	100	97.2	94.9	2.937	0.483	0.993	0.067
	ARDM*	82.6	100	97.0	95.6	3.153	0.494	0.991	0.050
	ARDG	80.5	100	95.0	95.3	2.705	0.502	0.994	0.127
	BSDG	85.9	100	92.4	98.4	2.609	0.560	0.993	0.059
	FADG	90.1	100	91.7	98.4	2.537	0.541	0.995	0.105
NEsN	ARDM	81.6	100	98.3	88.5	3.452	0.463	0.989	0.072
	ARDM*	82.1	100	98.6	94.5	2.942	0.486	0.986	0.065
	ARDG	85.0	100	97.2	95.3	3.138	0.495	0.990	0.072
	BSDG	87.2	100	96.7	96.7	2.681	0.519	0.993	0.064
	FADG	88.9	100	93.3	98.7	2.440	0.533	0.995	0.081
NsEs	ARDM	81.3	100	97.0	93.5	2.993	0.494	0.986	0.039
	ARDM*	76.9	100	97.4	92.1	3.159	0.485	0.990	0.060
	ARDG	84.2	100	95.4	94.7	2.776	0.509	0.992	0.037
	BSDG	89.0	100	92.4	98.4	2.778	0.532	0.994	0.057
	FADG	91.0	100	93.7	98.8	2.343	0.539	0.995	0.075

C.3 MOSES metrics on only novel molecules

In Table 5, we present the corresponding metrics for MOSES but only on novel molcules, and see the same pattern with DG improving over standard ARDM. As mentioned in the paper, all metrics are computed on only valid

Table 5: Metrics for MOSES, but computed only on novel molecules

		Valid	Uniqueness	Novelty	Filters	FCD	SNN	Frag	Scaf
Uniform	ARDM	82.2	100	100	94.7	2.972	0.480	0.993	0.068
	ARDG	80.5	100	100	95.0	2.797	0.496	0.993	0.131
	BSDG	85.9	100	100	98.2	2.699	0.524	0.994	0.063
	FADG	90.1	100	100	98.3	2.607	0.533	0.995	0.114

molecules, meaning that even if the novelty is lower, the *total number* of novel (and valid) molecules generated with desired properties can be higher. For example, according to the numbers in Table 5 and Table 2, the fraction of the generated molecules by ARDM that are **valid**, **novel** and pass the **filters** is $0.822 \cdot 0.972 \cdot 0.947 = 75.7\%$, while the same number for FADG is $0.901 \cdot 0.917 \cdot 0.983 = 81.2\%$.

C.4 Generation with different seeds

To evaluate how much the results can vary, we generate new sets of 1000 molecules using a different seed, and present the results in Table 6 and Table 7. The same pattern can be seen in these tables as those in the main text, with discriminator guidance improving most metrics.

Table 6: Evaluation metrics on the QM9 dataset when using a different seed

Order	Model	Validity (%) \uparrow	Uniqueness (%) \uparrow	Atom stable (%) \uparrow	Mol stable (%) \uparrow
	Test data	97.9	100	98.6	87.4
	DiGress	95.4	97.6	98.1	79.8
Uniform	ARDM	87.5	100	96.0	55.3
	ARDG	88.0	99.4	97.0	67.8
	BSDG	97.0	99.7	98.9	89.0
	FADG	97.6	99.2	98.7	88.3
NEsN	ARDM	95.1	99.7	97.1	73.0
	ARDG	93.2	99.8	97.3	74.5
	BSDG	98.8	99.7	98.8	89.4
	FADG	97.0	99.6	98.6	88.0
NsEs	ARDM	92.7	99.9	97.3	77.0
	ARDG	94.3	99.7	97.8	81.6
	BSDG	97.9	99.7	98.5	87.1
	FADG	97.3	99.1	98.4	85.8

C.5 Varying number of particles

In Table 8, we present the results on MOSES for BSDG and FADG when varying the number of particles, N . We conclude that $N = 10$ is a reasonable compute-quality trade-off for this case.

D LICENSE INFORMATION

QM9 is provided by the original authors on Figshare⁵, and was accessed through Pytorch Geometric⁶. The license is unspecified.

The MOSES dataset was accessed via the official open source library `moses`⁷ with an MIT License.

⁵https://springernature.figshare.com/collections/Quantum_chemistry_structures_and_properties_of_134_kilo_molecules/978904

⁶https://pytorch-geometric.readthedocs.io/en/latest/generated/torch_geometric_datasets.QM9.html

⁷<https://github.com/molecularsets/moses/>

Table 7: Evaluation metrics on the MOSES dataset with a different seed.

Order	DG	Validity (%) \uparrow	Uniqueness (%) \uparrow	Novelty (%) \uparrow	Filters (%) \uparrow	FCD \downarrow	SNN \uparrow	Frag \uparrow	Scaf \uparrow
Uniform	ARDM	80.4	100	97.4	93.8	2.942	0.487	0.992	0.095
	ARDG	84.6	100	95.0	94.9	2.503	0.509	0.995	0.106
	BSDG	89.6	100	92.4	98.2	2.789	0.527	0.992	0.050
	FADG	88.3	100	91.4	98.2	2.532	0.537	0.993	0.082
NEsN	ARDM	81.9	100	98.2	90.2	3.110	0.468	0.990	0.083
	ARDG	85.6	100	97.1	94.2	2.887	0.501	0.993	0.075
	BSDG	86.6	100	95.4	97.2	2.746	0.518	0.994	0.046
	FADG	89.2	100	93.4	98.1	2.749	0.527	0.993	0.035
NsEs	ARDM	78.8	100	98.6	93.3	3.159	0.487	0.987	0.091
	ARDG	83.7	100	96.3	95.5	2.736	0.505	0.994	0.055
	BSDG	88.0	100	93.2	98.0	2.636	0.526	0.995	0.079
	FADG	91.0	100	92.7	98.7	2.565	0.535	0.993	0.095

Table 8: Evaluation metrics on the MOSES dataset when using the uniform generation order and varying the number of particles, N .

	N	Validity (%) \uparrow	Uniqueness (%) \uparrow	Novelty (%) \uparrow	Filters (%) \uparrow	FCD \downarrow	SNN \uparrow	Frag \uparrow	Scaf \uparrow
BSDG	5	87.9	100	92.0	97.5	2.606	0.528	0.993	0.073
	10	85.9	100	92.4	98.4	2.609	0.560	0.993	0.059
	25	87.9	100	92.4	98.2	2.650	0.538	0.992	0.079
	50	89.8	100	92.9	98.4	2.459	0.537	0.993	0.090
FADG	5	88.1	100	91.5	98.4	2.504	0.533	0.994	0.082
	10	90.1	100	91.7	98.4	2.537	0.541	0.995	0.105
	25	90.1	100	92.1	99.1	2.589	0.537	0.991	0.036
	50	87.5	100	90.5	99.1	2.809	0.532	0.993	0.035

Parts of the code used in this work (e.g., the graph transformer and code for evaluation) were taken from the public DiGress code repository⁸ and modified. This code is licensed with an MIT License.

⁸<https://github.com/cvignac/DiGress/>