

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 REDDiT: REHASHING NOISE FOR DISCRETE VISUAL GENERATION

Anonymous authors

Paper under double-blind review

ABSTRACT

In the visual generative area, discrete diffusion models are gaining traction for their efficiency and compatibility. However, pioneered attempts still fall behind their continuous counterparts, which we attribute to noise (absorbing state) design and sampling heuristics. In this study, we propose a rehashing noise approach for discrete diffusion transformer (termed **ReDDiT**), with the aim to extend absorbing states and improve expressive capacity of discrete diffusion models. ReDDiT enriches the potential paths that latent variables traverse during training with randomized multi-index corruption. The derived rehash sampler, which reverses the randomized absorbing paths, guarantees high diversity and low discrepancy of the generation process. These reformulations lead to more consistent and competitive generation quality, mitigating the need for heavily tuned randomness. Experiments show that ReDDiT significantly outperforms the baseline model (reducing gFID from 6.18 to **1.61**) and is on par with the continuous counterparts. The code and models will be publicly available.

1 INTRODUCTION

Diffusion has been a competitive approach for generative workloads (Dhariwal & Nichol, 2021; Rombach et al., 2022b; Li et al., 2024), offering strong bidirectional perception and well-structured mechanisms Zhang et al. (2023) for global control over content. Within the continuous domain, diffusion transformers (DiTs) Peebles & Xie (2023), which progressively refine image latents from Gaussian noise, have achieved impressive and scalable results. Recently, the community shows a growing interest in discrete diffusion models (Hu & Ommer, 2024; Swerdlow et al., 2025), which is based on their practical advantages, *e.g.*, compatibility with language models for the indexable codebook, and efficiency for predicting multiple tokens at each inference. Early endeavors Chang et al. (2022; 2023); Gu et al. (2022) pursue efficiency through integrating visual tokenizers and BERT-style [mask] tokens Devlin et al. (2019). Recent studies Bai et al. (2025); Yang et al. (2025) improved the generation quality, demonstrating great potential of discrete diffusion.

Despite the progress, the performance of discrete diffusion methods remains lagging behind their continuous counterparts. Representative approaches, *e.g.*, masked visual token models (MVTMs) Chang et al. (2022); Yu et al. (2023), are puzzled by the mask design and confidence-based re-mask sampler (Hur et al., 2024), which restricts model's expressive capacity and makes prediction sensitive to adaptions given extensive training, Fig. 1(upper). Moreover, when paired with large-vocabulary codebooks from high-fidelity modern tokenizers, they encounter challenges such as slower sampling speeds and numerical inaccuracy (Zheng et al., 2024).

To address these limitations, we first propose two hypotheses. First, while discrete methods learn to recover plausible tokens from a monotonous [mask] canvas, the used noise design may not be well-suited for discrete visual generation. In continuous diffusion, Gaussian noise is used to progressively degrade the input to learn a smooth distribution shift (Ho et al., 2020; Lu et al., 2022). Discrete masking mimics this paradigm by collapsing all masked tokens to a single absorbing state, which, however, lacks the variability of Gaussian noise, in terms of both vocabulary richness and latent diversity. Consequently, the discrete process offers a far coarser signal, which limits its ability to represent diverse data distributions (Santos et al., 2023; Austin et al., 2021). Moreover, while continuous diffusion models introduce stochasticity at every inference step through noise injection,

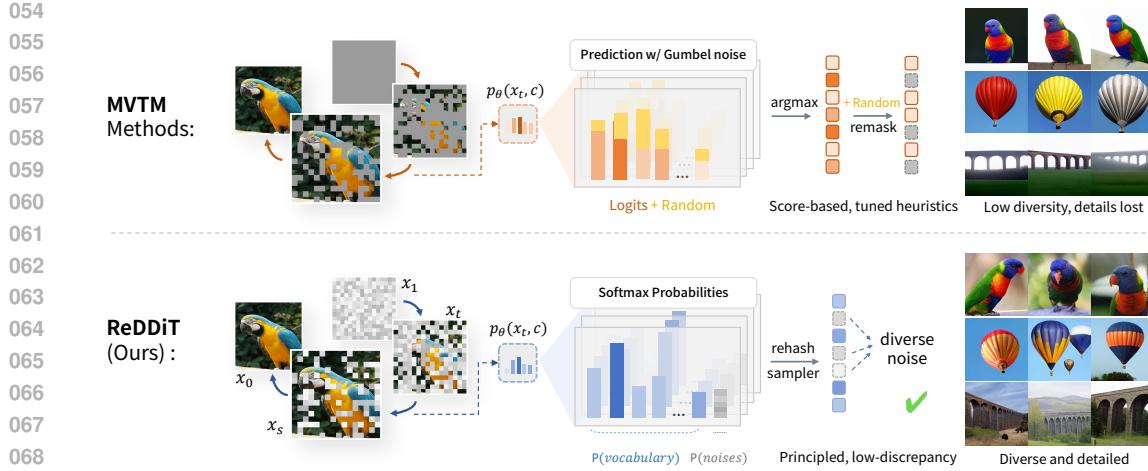


Figure 1: **Comparison the baseline discrete model (MVTM) with ReDDiT.** MVTMs rely on score-based remasking strategies with Gumbel-max to sample from logits, which leads to lower token diversity and suboptimal token selection. In contrast, ReDDiT introduces a systematic, low-discrepancy rehashing mechanism that leverages softmax-based probabilities, enabling diverse, high-quality sampling through a learned distribution. (This figure is best viewed in color)

discrete unmasking is inherently binary: tokens are either masked or deterministically decoded, Fig. 1(upper). This rigid mechanism constrains the flexibility of sample refinement during generation.

Second, the confidence-based re-mask sampler of MVTMs introduces a form of handcrafted randomness, which is implemented through Gumbel-max, to approximate sampling diversity. Unfortunately, this sampler compromises the probabilistic fidelity of generation, and the need to carefully balance token numbers decoded per step (for mitigating accumulation errors) leads to redundant sampling passes. As a result, Gumbel-max has evolved to a heavily tuned time variant trick with unstable performance, particularly when scaled to large-vocabulary codebooks. The above factors, rather than quantization alone, induce the performance gap between discrete and continuous models.

In this study, we propose a discrete diffusion model with an elaborate rehashing noise design, Fig. 1(lower). Our approach, termed **ReDDiT**, addresses the limitations of the uni-mask design by redefining absorbing states towards larger representational capacity, through enriching the potential paths that latent variables can traverse during diffusion. Specifically, we expand the masks to multiple indices along with the codebook and randomize them during data corruption. A rehash sampler is also derived with principled discrete diffusion theories to reverse the diffusion path for generation, guaranteeing high diversity and low discrepancy of the sampling process. We demonstrate that this rehashed noise facilitates learning a superior and regularized expressiveness, while eliminating reliance to hyper-parameterized randomness during sampling.

We further revisit the commonly used discrete diffusion objective and update it with empirical modifications. By adopting an improved ELBO Sahoo et al. (2024); Shi et al. (2024) with representation alignment (RepA) Yu et al. (2025) loss, we optimize the training efficiency and substantially improve the generation quality of discrete generative models. Moreover, ReDDiT aligns with recent advances in discrete flow matching Gat et al. (2024); Shaul et al. (2024), enabling token refreshment during sampling without training post-correction models (Lezama et al., 2022).

2 METHODOLOGY

For self-containment, we first review the DDM theory in Sec. 2.1. We then reformulate its diffusion dynamics and introduce rehashing noise for ReDDiT in Sec. 2.2. We finally discuss connection and comparison with other discrete diffusion models in Sec. 2.3.

108
109

2.1 PRELIMINARY: DISCRETE DIFFUSION MODEL

110 DDM defines a forward process over discrete variables by gradually corrupting the image tokens to
 111 absorbing states (masks) through a continuous-time Markov process. Assume that the data consists
 112 of tokens from a finite vocabulary \mathcal{V} . $x \in \mathcal{V}^L$ is a sequence of tokens (e.g., an image tokenized into
 113 indices) with length L . We denote the clean data as $x_{t=0}$ (x_0 for short), and noise it gradually as
 114 $t \rightarrow 1$. DDM defines an absorbing token $\mathbf{m} \in \mathcal{V}$, such that once a token is noised to \mathbf{m} it remains
 115 unchanged. At the terminal time $t = 1$, x_t fully transits to \mathbf{m}^L , which means $x_1^{i=1 \sim L} = \mathbf{m}$.

116 Let α_t be the noise scheduler (a monotonically decreasing survival function that satisfies $\alpha_0 =$
 117 $1, \alpha_1 = 0$). For $0 \leq s < t \leq 1$, the forward corruption process is governed by a continuous-time
 118 transition kernel $q(x_t^i | x_s^i)$ at the i -th element, as

$$119 \quad q(x_t^i | x_s^i) = \begin{cases} 1 - \alpha_{t|s}, & \text{if } x_t^i = \mathbf{m}, x_s^i \neq \mathbf{m} \\ \alpha_{t|s}, & \text{if } x_t^i = x_s^i, x_s^i \neq \mathbf{m} \\ 1, & \text{if } x_t^i = x_s^i, x_s^i = \mathbf{m} \\ 0, & \text{otherwise} \end{cases}, \quad \alpha_{t|s} = \frac{\alpha_t}{\alpha_s}. \quad (1)$$

124 Denoting q as the transition kernel and $\text{Cat}(\cdot; \pi)$ the categorical distribution determined by probability
 125 π , the corrupted data distribution at time t is written as

$$126 \quad x_t \sim q(x_t | x_0), q(x_t | x_0) = \text{Cat}(x_t; \alpha_t x_0 + (1 - \alpha_t) \mathbf{m}^L). \quad (2)$$

127 The generative model learns the reverse process $p_\theta(x_s | x_t)$, which denoises sample x_t at arbitrary
 128 time $t \in (0, 1]$ to a less noised state x_s at time $s < t$. Denoting $\delta(x_t^i, \mathbf{m})$ as the indicator function that
 129 only computes on masked tokens, and $\alpha'_t = \frac{d\alpha_t}{dt}$, the learning objective is derived Shi et al. (2024) as

$$131 \quad \mathcal{L}_{\text{DDM}} = -\mathbb{E}_{x_0, x_t} \int_{t=0}^{t=1} \left[\frac{\alpha'_t}{1 - \alpha_t} \sum_{i=1}^L \delta(x_t^i, \mathbf{m}) \log p_\theta(x_0^i | x_t) \right] dt. \quad (3)$$

134 For a linear scheduler, Eq. 3 is simplified via variable substitution Sahoo et al. (2024) to an equivalent
 135 form, as

$$136 \quad \mathcal{L}_{\text{DDM-linear}} = -\mathbb{E}_{t, x_0, x_t} \left[\frac{1}{t} \sum_{i=1}^L \delta(x_t^i, \mathbf{m}) \log p_\theta(x_0^i | x_t) \right]. \quad (4)$$

138 For conditional generation, class information c (e.g., labels or text prompts) is introduced to the
 139 denoising model as additional input. Following classifier-free guidance Ho & Salimans (2022), the
 140 model is trained with a random drop of labels, and the prediction is interpolated at inference, as

$$141 \quad \hat{p}_\theta(x_t, c) = p_\theta(x_t, \emptyset) + w \cdot (p_\theta(x_t, c) - p_\theta(x_t, \emptyset)), \quad (5)$$

143 where \emptyset is the dropped label and $w \geq 0$ controls the guidance strength.

144
145 2.2 DISCRETE DIFFUSION WITH REHASHING NOISE

146 The ordinal structure inherent in discrete data provides a valuable inductive bias for designing
 147 transition kernels in diffusion dynamics. Prior studies Austin et al. (2021); Campbell et al. (2022)
 148 show that assigning higher transition probabilities to neighboring pixel values—forming a *discrete*
 149 *Gaussian-like noise*—outperforms the single absorbing state approach on pixel-level datasets like
 150 CIFAR-10. However, when using visual tokenizers, the structure of discretized latents is learned
 151 rather than pre-defined, making such ordinal assumptions inapplicable. This insight motivates us
 152 to extend conventional mask tokens to a set of indices, and reverse the diffusion path with noise
 153 rehashing. This design allows the model to optimize its embedding space during training, enhancing
 154 its ability to model flexible and data-driven noise structures. We visualize the learned distributions in
 155 Fig. 2 (right).

156 **Reformulation.** Given d categories, let $\mathbf{e}_i \in \mathbb{R}^d$ be its one-hot vector where the i -th value is 1. We
 157 denote $\mathcal{E} = \{\mathbf{e}_i \in \mathbb{R}^d \mid i = 1, \dots, d\}$ as the basis of a categorical distribution (known as d -simplex),
 158 and a basis for absorbing states with capacity m : $\mathcal{M} = \{\mathbf{m}_j \in \mathbb{R}^m \mid j = 1, \dots, m\}$. With subscript
 159 $i, j = 0$ for pure visual or mask space, the sum of \mathcal{E} and \mathcal{M} can be denoted as

$$161 \quad \mathcal{V}_{(d,m)} \triangleq \left\{ \mathbf{v}_{(i,j)} \in \mathbb{R}^{d+m} \mid \mathbf{v}_{(i,j)} = \begin{cases} \mathbf{e}_i \oplus \mathbf{0}_m, & \text{for } i = 1, \dots, d, j = 0 \\ \mathbf{0}_d \oplus \mathbf{m}_j, & \text{for } j = 1, \dots, m, i = 0 \end{cases} \right\}. \quad (6)$$

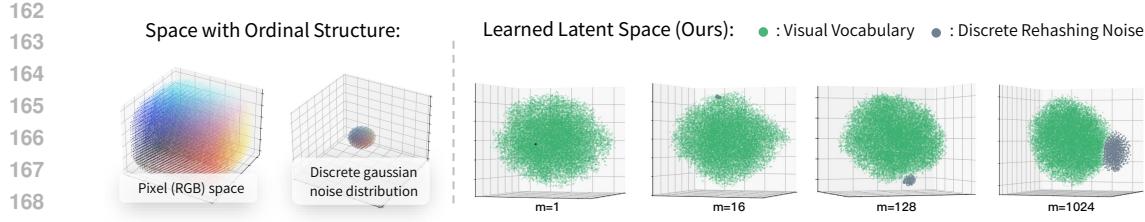


Figure 2: **Visualization of pixel and latent spaces.** m denotes the number of enriched noise indices. Note that the 3D t-SNE plot (right) is used **solely for clustering with no isotropic features**.

We further denote the subspace \mathcal{E}_d , $\mathcal{M}_m \in \mathcal{V}_{(d,m)}$ which contain *either* valid or mask tokens, as

$$\mathcal{E}_d = \{\mathbf{v}_{(i,0)} \in \mathcal{V}_{(d,m)} \mid i = 1, \dots, d\}, \quad \mathcal{M}_m = \{\mathbf{v}_{(0,j)} \in \mathcal{V}_{(d,m)} \mid j = 1, \dots, m\}. \quad (7)$$

To exploit visits across all the possible paths, we rewrite the transition kernel defined by Eq. 1 as

$$q(x_t^i | x_s^i) = \begin{cases} 1 - \alpha_{t|s}, & \text{if } x_t^i \in \mathcal{M}_m, x_s^i \notin \mathcal{M}_m \\ \alpha_{t|s}, & \text{if } x_t^i = x_s^i, x_s^i \notin \mathcal{M}_m \\ 1/m, & \text{if } x_t^i \in \mathcal{M}_m, x_s^i \in \mathcal{M}_m \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

With above definitions, we reformulate the diffusion process of x as a **transition from \mathcal{E}_d to \mathcal{M}_m** . We train the model by feeding it with corrupted data, of which the distribution is inferred as $x_t \sim \text{Cat}(x_t; \alpha_t x_0 + (1 - \alpha_t) \text{U}(\mathcal{M}_m^L))$, where $\text{U}(\mathcal{M}_m^L)$ is the uniform distribution upon \mathcal{M}_m^L .

Rehash Sampling. To generate a sequence of length L , the reverse process starts with $x_1 \sim \text{U}(\mathcal{M}_m^L)$. The subsequent latents x_t are generated by discretizing the reverse timeline T to K steps. We denote this schedule as $T^{1:K+1}$ such that $T^1 = 1$ and $T^{K+1} = \varepsilon$, with ε being an arbitrarily small positive constant. Let δ^i indicate the i -th token's value, the reverse process is deduced as

$$q_{s|t}^i = q(x_s^i | x_t) = \begin{cases} 1, & \text{if } x_s^i = x_t^i, x_t^i \notin \mathcal{M}_m \\ \frac{1 - \alpha_s}{m(1 - \alpha_t)}, & \text{if } x_s^i \in \mathcal{M}_m, x_t^i \in \mathcal{M}_m \\ \frac{\alpha_s - \alpha_t}{1 - \alpha_t} \delta^i p_\theta(x_t), & \text{if } x_s^i \notin \mathcal{M}_m, x_t^i \in \mathcal{M}_m \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

Comparing with MVTM sampler in Alg. 1, our rehash sampler is shown in Alg. 2. **Our algorithm shares the similar idea as MDLM Sahoo et al. (2024), but applies `torch.multinomial` (Multnm. in step 10) for low-discrepancy¹ categorical sampling.**

Algorithm 1 MVTM Sampling

- 1: **Inputs:** label c , scheduler α_t , length L ,
- 2: **Settings:** number of steps K , $G(t)$, \mathcal{G}
- 3: Initialize: $x_1 \leftarrow \mathcal{M}_1^L$, $t \leftarrow 1$.
- 4: **for** $k = 1$ to K **do**
- 5: $t \leftarrow \frac{K-k+1}{K}$, $s \leftarrow \frac{K-k}{K}$
- 6: logits $\leftarrow f_\theta(x_t, c)$
- 7: $p_{\text{score}} \leftarrow \text{logits} + G(t) \cdot \mathcal{G}$
- 8: $x_{\text{pred}} \leftarrow \text{argmax}(p_{\text{score}})$ \triangleright Predict-all
- 9: $x_s \leftarrow \text{where}(x_t = [m], x_{\text{pred}}, x_t)$
- 10: $p_{\text{conf}} \leftarrow p_{\text{score}} + G(t) \cdot \mathcal{G}$
- 11: $m_{\text{re}} \leftarrow \text{argsort}(p_{\text{conf}})[1 : L \cdot (1 - \alpha_s)]$
- 12: $x_s \leftarrow \text{where}(m_{\text{re}}, [m], x_s)$ \triangleright Re-mask
- 13: **end for**
- 14: **Return:** fully unmasked sequence x_0

Algorithm 2 Rehash Sampling (Ours)

- 1: **Inputs:** label c , scheduler α_t , length L .
- 2: **Settings:** number of steps K .
- 3: Initialize: $x_1 \sim \text{U}(\mathcal{M}_m^L)$, $t \leftarrow 1$, $T^{1:K}$.
- 4: **for** $k = 1$ to K **do**
- 5: $t \leftarrow T^k$, $s \leftarrow T^{k+1}$
- 6: **# the rehash operation:**
- 7: $x_t \leftarrow \text{where}(x_t \in \mathcal{M}_m, \text{U}(\mathcal{M}_m^L), x_t)$
- 8: logits $\leftarrow f_\theta(x_t, c)$
- 9: $p \leftarrow \text{Softmax}(\text{logits})$
- 10: $q_{s|t} \leftarrow \frac{\alpha_s - \alpha_t}{1 - \alpha_t} \cdot p + \delta_{m[0]} \cdot \frac{1 - \alpha_s}{1 - \alpha_t}$
- 11: $x_{\text{pred}} \leftarrow \text{Multnm.}(q_{s|t})$ \triangleright w/ masks
- 12: $x_s \leftarrow \text{where}(x_t \in \mathcal{M}_m, x_{\text{pred}}, x_t)$
- 13: **end for**
- 14: **Return:** fully unmasked sequence x_0

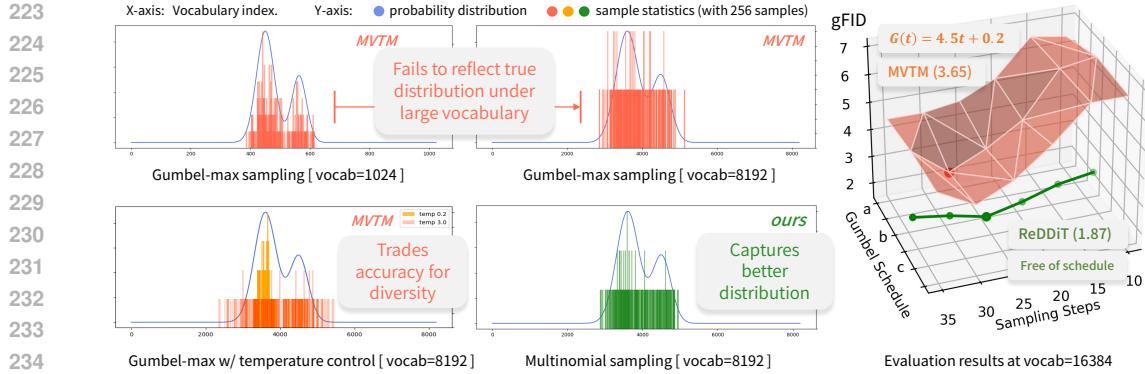
¹MDLM uses gumbel-max for sampling, which may incur inaccuracy (Zheng et al., 2024). Besides, we deliberately merge the probabilities at step 9 to keep an overall noise sampling probability, as small values might be truncated, which also worsens sampling accuracy.

216 The random nature of absorbing states inspires a rehash operation: we shuffle these tokens at the
 217 beginning of each step by $x_t \leftarrow \text{where}(x_t \in \mathcal{M}_m, \text{U}(\mathcal{M}_m^L), x_t)$. Proof to Eq.9 is included in
 218 Appendix. B.

219

220 2.3 DISCUSSION

221



235
 236 **Figure 3: Sampler comparison.** *Left:* Gumbel-max is theoretically equivalent to our method, yet
 237 it struggles to reflect the true distribution under limited sample passes. The multinomial approach
 238 captures the distribution more accurately. *Right:* our model achieves lower gFID across different
 239 sampling steps without tuning Gumbel-max, indicating more efficient and faithful sampling. *a, b, c*
 240 refer to three uniformly sampled $G(t)$ set for MVTM sampling. See supplementary for experimental
 241 codes. (This figure is best viewed in color)

242

243

244 **Comparison with MVTM.** Masked visual token models (MVTMs) borrow the objective

245

$$\mathcal{L}_{\text{MVTM}} = -\mathbb{E}_{t, x_0, x_t} \sum_{i=1}^L \delta(x_t^i, \mathbf{m}) \log p_{\theta}(x_t^i | x_t), \quad (10)$$

246

247 from masked language models Devlin et al. (2019) and predict on masked tokens with a maximum
 248 likelihood. Besides the reformulated corruption (Eq. 8) and reverse process (Eq. 9), ReDDiT differs
 249 in the following aspects: *(i)* the training objective (Eq. 4), which is derived from DDM, providing
 250 better theoretical and empirical results. *(ii)* it can easily sample with a arbitrarily discretized timeline,
 251 while MVTM couples training and inference, restricting its sampling flexibility; *(iii)* the rehash
 252 sampler (Alg. 2) includes absorbing states in categorical sampling with lower discrepancy, different
 253 from MVTM’s predict-remask sampler with time variant intensity $G(t)$ over Gumbel noise \mathcal{G} (Alg. 1)
 254 ². Gumbel-max suffers from numerical inaccuracy (Zheng et al., 2024) and we noitice that it becomes
 255 worse on large vocabulary (Fig. 1, 3 with our reproduced results), which limits MVTM’s potential.

256

257

258 **Relationship to DFM.** Discrete flow matching (DFM) Gat et al. (2024) introduces a transition
 259 process based on masked tokens. Its training objective was initially designed as the masked token
 260 loss (10), and evolved to a time-weighted cross-entropy loss (Shaul et al., 2024) for generalized
 261 diffusion paths, which is similar to ours. The similarity enables a direct comparison between the DFM
 262 sampler and our rehash sampler using the same trained model weights. We notice that it generally
 263 requires more steps to reach optimal results, as the DFM sampler offers a refinement mechanism via
 264 token-wise updates. Since the gradual decoding method is shared, we can integrate certain DFM steps
 265 into our sampling procedure for refinement. This leads to ~ 0.1 gFID improvement on ImageNet-1K.
 266 Refer to Appendix D for details.

267

268

269

²The logits corresponding to previously restored tokens’ indices are manually set to infinity for both methods, so that they will not be noised again in the following steps. This leads to an implementation of any-order auto-regressive model (Ou et al., 2024) if DDM’s decoded tokens per step is limited to 1.

270 3 EXPERIMENT
271272 3.1 IMPLEMENTATION
273274 **Datasets.** The experiments are conducted on ImageNet-1K Deng et al. (2009), which consists of
275 1000 categories, 1281167 images and are cropped to resolution 256×256 for training. The generation
276 quality is evaluated using Fréchet Inception Distance (FID) Heusel et al. (2017) and the Inception
277 Score (IS) Salimans et al. (2017). FID measures the distance between the distributions of generated
278 and real images in the feature space of a pre-trained Inception network, while IS evaluates both the
279 confidence and diversity of generated images by analyzing predicted label distribution. We compute
280 generation FID (gFID \downarrow)³ and IS \uparrow on 50k generated samples.
281282 **Pre-processing.** Following the setting in LlamaGen Sun et al. (2024), we apply the ten-crop
283 augmentation on images, and use pre-trained tokenizers to convert them to discrete tokens. We
284 pick IBQ-f16 Shi et al. (2025) tokenizer as default for its scalable and promising performance in
285 generation tasks, which uses a 16×16 downsampling ratio and converts a 256×256 image into
286 256 discrete tokens. The tokenizer has a codebook with 16384 entries. The LlamaGen-f16 (used in
287 Tab. 2) and LlamaGen-f8 tokenizer Sun et al. (2024) (used in Tab. 1) are also used for comparison
288 with recent discrete generation methods. All tokenizers are used out-of-the-box without modification.
289290 **Representation Alignment.** Recent study Yu et al. (2025) has shown that the alignment of in-
291 termediate representations between diffusion transformers and vision encoders accelerates training
292 convergence of diffusion models. Accordingly, the alignment is designed as a regularization term
293 with $\lambda = 0.5$. We extract diffusion transformer’s 8-th layer intermediate feature $\mathbf{h}^{[n]}(x_t)$ and align it
294 with the original image’s dinov2-b Oquab et al. (2023) encoded features $f_{\text{enc}}(x_0^{\text{ori}})$. The intermediate
295 features are projected by a small trainable MLP h_φ . The $\text{sim}(\cdot, \cdot)$ computes the mean of element-wise
296 cosine similarity between embeddings, as
297

298
$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{DDM-linear}} + \lambda \mathcal{L}_{\text{RepA}}, \quad \mathcal{L}_{\text{RepA}} = -\mathbb{E}_{x, t} [\text{sim}(f_{\text{enc}}(x_0^{\text{ori}}), h_\varphi(\mathbf{h}^{[n]}(x_t)))]. \quad (11)$$

300 This alignment was proposed for continuous diffusion models, and we firstly validate that it’s also
301 suitable for training discrete models. However, from our observation, as a training acceleration
302 technique, RepA **does not** provide relative performance gain if training sufficiently for discrete
303 latents. We only use RepA to improve training efficiency and probe the inner dynamics through
304 training as in Fig. 4. See Appendix F for a detailed discussion.
305306 **Training and Evaluation.** The proposed model is based on DiT Peebles & Xie (2023) architecture,
307 with reference to its discrete prediction version Sahoo et al. (2024). 2D-RoPE Su et al. (2024) and
308 min-SNR Zhang & Sennrich (2019) are applied for training efficiency. The model is optimized
309 using the AdamW optimizer with a cosine decay. Class-conditional training is enabled using class
310 embeddings and a drop-rate of 0.1 for generation with CFG. Details are provided in Appendix E.
311312 3.2 PERFORMANCE AND COMPARISON
313314 We compare the proposed ReDDiT model with other generative models on the ImageNet-1K 256×256
315 in Tab. 1. The IBQ tokenizer is used for the default L and XL models. We also utilize LlamaGen-f8
316 with 128 noise capacity to evaluate its high-resolution potentials (noted as ReDDiT-XL_{f8}). We use a
317 linear increasing guidance following the common practice of Gao et al. (2023).
318319 **Generation Quality.** As shown in Tab. 1, ReDDiT achieves the best performance among the
320 compared discrete models. It outperforms the baseline (MaskGIT Chang et al. (2022)) with signif-
321 icant margins (gFID: 2.13 vs 6.18 and IS: 294.7 vs. 182.1). It also outperforms the recent DDM
322 method Hu & Ommer (2024) and TiTok-S-128 Yu et al. (2024), which is extensively fine-tuned on
323 quantized latents. Compared with continuous diffusion models, ReDDiT exhibits on-par efficiency
324 and performance, showing great potential for discrete generation.
325326 ³The gFID is used as the quality metric for generative models’ performance, while rFID refers to the
327 reconstruction quality of a visual tokenizer.
328

324
 325 **Table 1: Performance comparison on class-conditional ImageNet 256×256.** Look-up free quan-
 326 tizers are beyond the scope of this paper. *ft.*(in gray) indicates that the decoder is fine-tuned for
 327 quantized latents. Wall-clock inference time relative to ReDDiT-XL is reported.

328 329 Type	330 331 332 333 Model	334 335 336 337 338 Tokenizer		339 340 341 342 343 344 Generator				
		#tokens	345 346 347 codebook	348 349 gFID↓	350 351 IS↑	352 353 354 355 #Params	356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 #Steps	371 372 373 374 375 Time
330 331 332 333 Diff.	LDM-4 Rombach et al. (2022a)	4096×3	-	3.60	247.7	400M	250	-
	DiT-XL/2 Peebles & Xie (2023)	1024×4	-	2.27	278.2	675M	250	18
	MDTv2 Gao et al. (2023)	1024×14	-	1.58	314.7	676M	256	18
	SiT-XL Ma et al. (2024)	1024×4	-	2.42	238.5	675M	30	2
	SiT-XL w/ Solver Wang et al. (2025)	1024×4	-	2.24	244.1	730M	15	1.2
334 335 336 337 338 AR	Taming-VQGAN Esser et al. (2021)	256	1024	15.78	74.3	1.4B	256	8
	RQ-Transformer Huang et al. (2023)	256	16384	7.55	134.0	3.8B	64	8.5
	ViT-VQGAN Yu et al. (2022)	1024	8192	4.17	175.1	1.7B	1024	>10
	LlamaGen-3B Sun et al. (2024)	576	16384	2.18	263.3	3.1B	576	20
	RandAR-XXL Pang et al. (2024)	512	16384	2.15	322.0	1.4B	88	4
339 340 341 MVTM	VAR-d30 Tian et al. (2024)	680	4096	1.97	334.7	2.0B	10	0.5
	MaskGIT Chang et al. (2022)	256	1024	6.18	182.1	227M	8	0.2
	MaskGIL-XXL Xin et al. (2025)	256	16384	3.71	303.4	1.4B	8	0.8
342 343 344 DDM	TiTok-S-128 _{ft} . Yu et al. (2024)	128	4096	1.97	281.8	287M	64	1.6
	ITM Hu & Ommer (2024)	1024	16384	5.30	183.0	546M	100	3
	ReDDiT-L (ours)	256	16384	2.13	294.7	346M	20	0.5
	ReDDiT-XL (ours)	256	16384	1.74	313.6	675M	32	1
	ReDDiT-XL _{ft} (ours)	1024	16384	1.61	318.5	675M	64	2

345
 346 **Table 2: Comparison of models with the same tokenizer.** Reconstruction FID (rFID) indicates the
 347 tokenizer’s reconstruction quality from its quantized codes. Dim denotes codebook dimension.

348 349 Model	350 351 352 353 354 355 VQ Tokenizer Info.			356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 Generator		
	371 372 373 374 375 Identity		rFID	376 377 378 379 380 dim	#Params	381 382 383 384 385 gFID↓
LlamaGen-L _{AR} Sun et al. (2024)					343M	3.80
RandAR-L _{AR} Pang et al. (2024)	LlamaGen-f16 Sun et al. (2024)	2.19	8	343M	2.55	
Ours _{DDM(ReDDiT-L)}					346M	2.41
IBQ-B _{AR} Shi et al. (2025)	IBQ-tokenizer Shi et al. (2025)	1.37	256	343M	2.88	
Ours _{DDM(ReDDiT-L)}					346M	2.13

358 **Efficiency.** ReDDiT is born with the high-efficiency advantage of discrete diffusion models, com-
 359 paring with AR models. As shown in Tab. 1, the inference time of ReDDiT is slightly longer
 360 than MaskGIT, while the performance is overwhelming. Without acceleration techniques, ReDDiT
 361 achieves a competitive performance which AR and traditional diffusion models use more than 250
 362 steps to achieve. Notably, when armed with recent efforts that tailored KV-Cache Liu et al. (2025) for
 363 discrete diffusion models, ReDDiT’s inference can be further boosted (not included in the main paper
 364 for fair comparison). See Appendix G for details.

365 Besides the major comparison, we also conduct an experiment that utilizes the identical tokenizer
 366 in previous AR models and validate our method’s effectiveness. As can be seen in Tab. 2, ReDDiT
 367 outperforms AR methods in generation tasks across different tokenizers.

369 3.3 DETERMINING NOISE CAPACITY

371 The reformulated discrete diffusion dynamics defines transitioning from \mathcal{E}_d to \mathcal{M}_m . Under this
 372 setting, it is necessary to empirically determine the optimal value of m for a fixed tokenizer with
 373 vocabulary size d , as the latent representations learned by VAEs are variant. We keep the training
 374 setup fixed and conduct experiments *w.r.t.* the noise capacity m . We also visualize $\mathcal{L}_{\text{RepA}}$, which
 375 captures the degree of representation alignment Yu et al. (2025) within the transformer.

376 The alignment loss visualization shows that increasing the number of absorbing states introduces
 377 greater randomness, initially making predictions more difficult due to confusion with valid tokens.
 378 However, this gap narrows as training progresses, and the model converges to a similar alignment

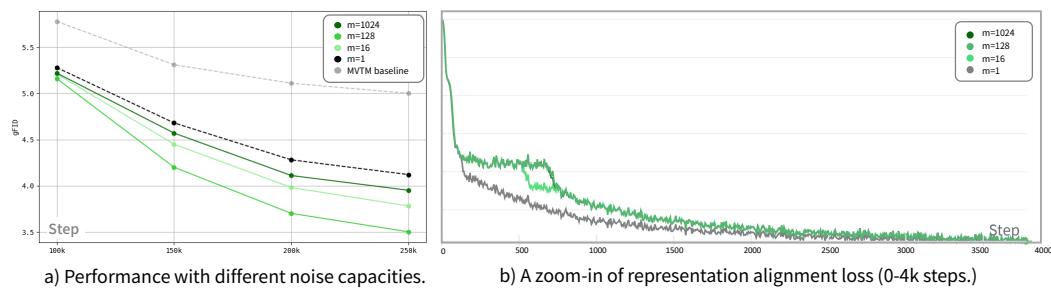


Figure 4: **Comparison of noise capacities.** We re-implemented training with the same training recipe. The generation quality and representation alignment trends are visualized.

lower bound, suggesting effective representation learning across different configurations. As shown in Fig. 4 (left), generation quality improves with increasing noise capacity initially. The LlamaGen-f16 tokenizer achieves peak performance at $m = 128$, while the IBQ tokenizer performs best at $m = 1024$. We attribute this to the codebook design: the lower dimensional LlamaGen-f16 codebook produces more compact latents, which also determines its smaller noise endurance.

3.4 ABLATION STUDY

Unless specified, all the models are trained on ImageNet 256×256 under the default settings for 100k iterations. We use a constant guidance scale of 2.0 and 20 steps for generation, and report gFID \downarrow computed on 50K samples. Precision (Prec. \uparrow) and Recall (Rec. \uparrow) are also reported.

Sampling Timeline. Recovering complete information from noise remains critical to diffusion-based models (Lu et al., 2022; Wu et al., 2024). Recent work shows MVTM’s non-linear scheduler for training is less critical when using high-capacity tokenizers. Evidence of time-invariance in DDMs (Sahoo et al., 2024; Shi et al., 2024) further supports decoupling training from sampling. In our experiments, a linear scheduler with constant signal-to-noise ratio decay, yields optimal training dynamics. Among the timeline discretization tested, Fig. 5, the *cosine* schedule is employed for our ReDDiT model for best performance in Tab. 3.

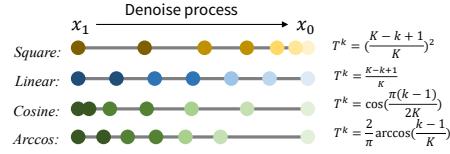


Figure 5: Illustration of discretized timeline with $K = 7$. The slow-to-fast sampling works better than linear schedules.

Table 3: **Ablated Design Choices.** ReDDiT-L is trained for 100k iters. Final setting denoted in gray.

(a) General model design					(b) Sampling timeline		
Train Config	Sample Config	gFID	Prec.	Rec.	Steps	Timeline	gFID
MVTM + RepA loss	MVTM sampler	6.83	0.75	0.39	20	linear	7.18
Switch to objective (11)	MVTM sampler	6.23	0.77	0.41	32	linear	6.43
same as above	Rehash sampler	5.75	0.78	0.45	20	arccos	5.04
+ 2D-RoPE + min-SNR	Rehash sampler	5.51	0.79	0.45	20	square	7.39
same as above	+ DFM refine	5.40	0.81	0.52	20	cosine	4.91

General Design. We ablate the general choices of ReDDiT, which starts with a re-trained MVTM baseline methods (with LlamaGen-f16 and RepA for faster convergence as default) in Tab. 3. The applied techniques like 2D-RoPE are also ablated with re-training. As shown, through the revised objective and our proposed sampler, ReDDiT alone improves FID by ~ 1.0 compared to the baseline model. When combined with modern modification on transformers, it can further improve the performance, showing its complementaryness with main-stream efforts.

432 **Impact of Rehash Operation.** To validate the rehash
 433 operation for encouraging path diversity by resampling
 434 noise tokens, we compare noise capacities $m = 1$ and
 435 $m = 128$. As shown in Table 4, increasing capacity with a
 436 fixed absorbing state actually degrades performance com-
 437 pared to the baseline. While enabling random initialization
 438 improves gFID, the full rehash mechanism is essential to
 439 unlock the model’s capacity, confirming that active resam-
 440 pling is required to prevent overly deterministic sampling.

Table 4: Ablation on Rehashing.

Setting	gFID
$m = 1$ (Baseline)	4.13
$m = 128$ (Fixed State)	4.25
$m = 128$ (No Rehash)	4.07
$m = 128$ (Full Rehash)	3.92

442 3.5 QUALITATIVE RESULT

443 **Class-conditional Generation.** Figure 6 presents representative class-conditional samples gener-
 444 ated by the proposed ReDDiT model. The outputs across diverse image classes consistently exhibit
 445 high fidelity and diversity. [Please refer to Appendix H for more samples](#).

446 **Image Editing.** We further demonstrate ReDDiT’s editing capability in Figure 6, highlighting its
 447 bi-directional perceptual competence. Following MaskGIT Chang et al. (2022), we replace a region
 448 of the input image with noise tokens and employ the same generation pipeline to inpaint the missing
 449 content, conditioned on a class label c .

453 4 RELATED WORK

454 **Diffusion Models.** Diffusion models Ho et al. (2020); Song et al. (2020) have emerged as a powerful
 455 class of generative methods that learn data distributions by reversing a gradual noising process over
 456 time. These models are primarily designed for continuous domains such as images Dhariwal & Nichol
 457 (2021); Gao et al. (2023); Peebles & Xie (2023), defining a forward process that transforms data x_0
 458 into noise x_1 : $x_t \sim \mathcal{N}(\sqrt{\alpha_t}x_0; (1 - \alpha_t)\mathbf{I})$ where α_t controls the noise schedule. The generative
 459 (reverse) process learns a denoising model $p_\theta(x_s | x_t)$, often parameterized via a neural network θ to
 460 predict either noise or clean data.

461 **Discrete Diffusion Models.** Discrete diffusion has been previously governed by masked visual
 462 token models (MVTMs) Chang et al. (2022; 2023); Gu et al. (2022); Yu et al. (2023; 2024); Hur et al.
 463 (2024). This model leverages a BERT-style [mask] token to corrupt the tokenized image sequence
 464 and trained the network with a simple cross-entropy loss on masked tokens, resulting in a score-based
 465 prediction. It generates tokens in a non-autoregressive fashion, by remasking the tokens with least
 466 scores at each inference as depicted in Alg. 1.

467 Recent studies unlocked the principled discrete diffusion model (DDM) Sahoo et al. (2024); Shi et al.
 468 (2024) and discrete flow-matching (DFM) Gat et al. (2024); Shaul et al. (2024), which adapt the
 469 Markov chain theory, enabling generation over text Ou et al. (2024); Nie et al. (2025), moleculesShaul
 470 et al. (2024), and other discrete representations Austin et al. (2021); Nisonoff et al. (2024). Unlike
 471 MVTMs, the principled DDM and DFM mostly derive a time-weighted cross-entropy loss to supervise
 472 the training procedure and apply a gradual unmasking method based on probabilities.

476 5 CONCLUSION

477 We proposed ReDDiT, a discrete visual generative model built upon a discrete diffusion architecture
 478 with novel noise designs and efficient sampling strategies. Our key contribution lies in the integra-
 479 tion of rehashing noise with samplers, which together ensure both diversity and low discrepancy
 480 throughout the generative process. By introducing rehashing noise, ReDDiT enriches the potential
 481 paths that latent variables can traverse during training, regularize training dynamics and enhances
 482 model’s representational capacity. Extensive experiments demonstrate that discrete generative models
 483 can achieve performance on par with their continuous counterparts while offering top-tier efficiency.
 484 This study paves a promising way for discrete generative modeling and offers fresh insights toward
 485 unifying visual and language generation—a path we leave for future exploration.

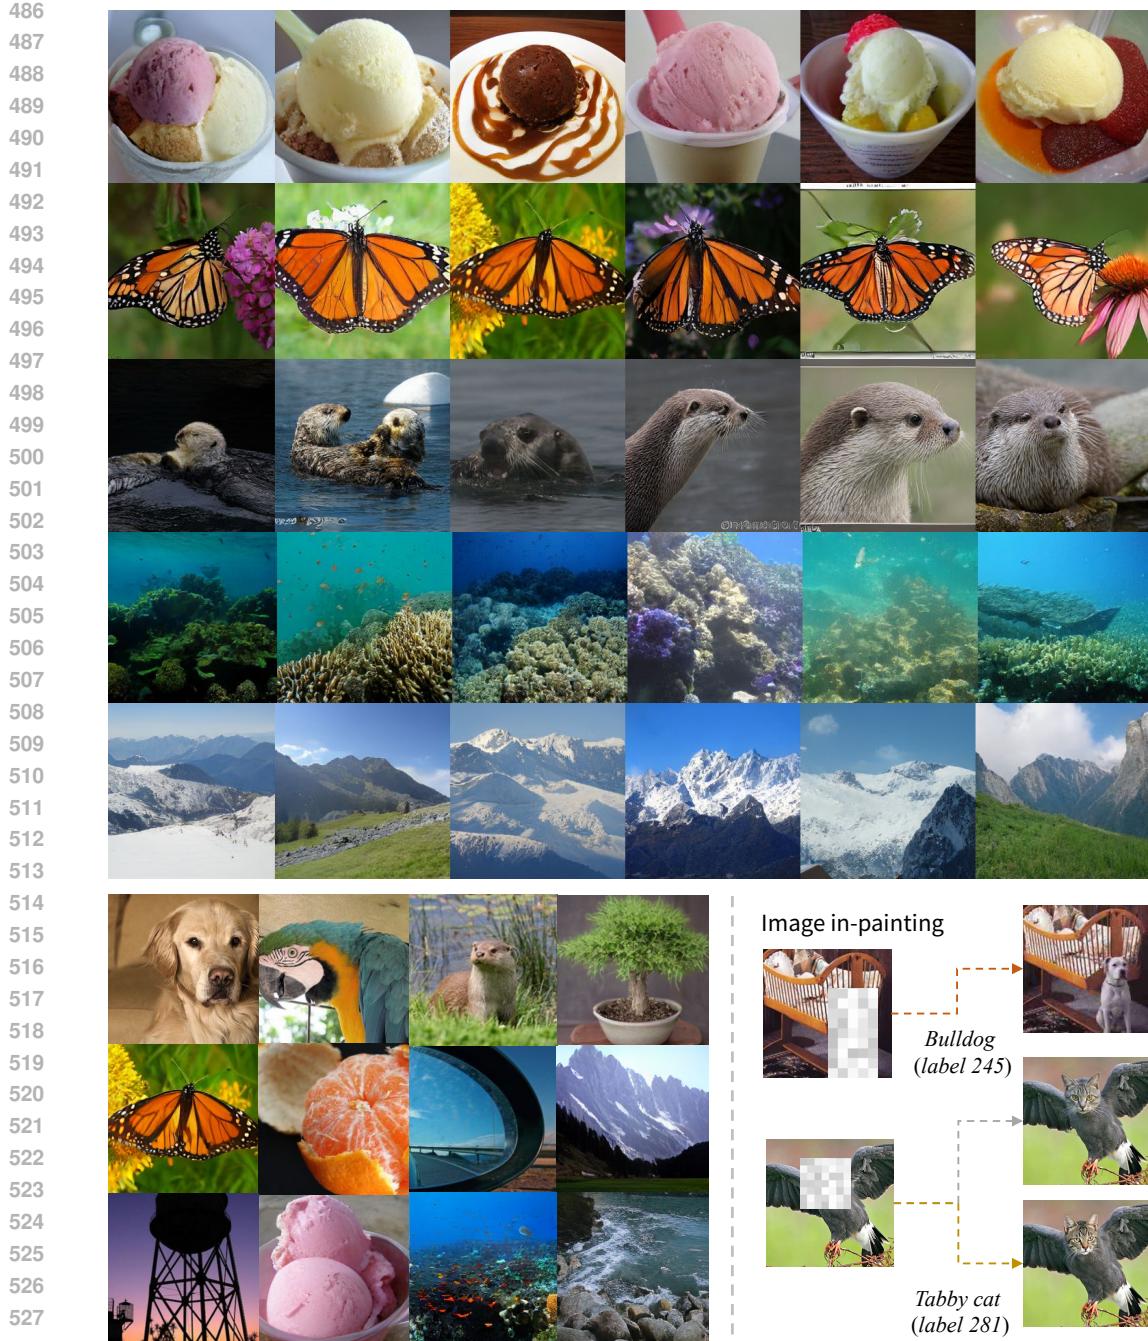


Figure 6: Class-conditional generation and in-painting samples of ReDDiT on ImageNet 256 × 256.

Reproducibility statement We provide key algorithms in the main text. Further details are available in the code implementation in the supplementary materials. The dataset used is publicly available, while the experiment process is carried out following the common practice of generative models.

REFERENCES

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *NeurIPS*, pp. 17981–17993, 2021.

540 Jinbin Bai, Tian Ye, Wei Chow, Enxin Song, Qing-Guo Chen, Xiangtai Li, Zhen Dong, Lei Zhu,
 541 and Shuicheng Yan. Meissonic: Revitalizing masked generative transformers for efficient high-
 542 resolution text-to-image synthesis. In *ICLR*, 2025.

543 Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and
 544 Arnaud Doucet. A continuous time framework for discrete denoising models. *NeurIPS*, 35:
 545 28266–28279, 2022.

546 Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative
 547 image transformer. In *CVPR*, pp. 11315–11325, 2022.

548 Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang,
 549 Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation
 550 via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023.

551 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale
 552 hierarchical image database. In *IEEE CVPR*, pp. 248–255, 2009.

553 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
 554 bidirectional transformers for language understanding. In *NAACL*, pp. 4171–4186, 2019.

555 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*,
 556 34:8780–8794, 2021.

557 Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image
 558 synthesis. In *CVPR*, pp. 12873–12883, 2021.

559 Shanghua Gao, Pan Zhou, Ming-Ming Cheng, and Shuicheng Yan. Mdtv2: Masked diffusion
 560 transformer is a strong image synthesizer. *arXiv preprint arXiv:2303.14389*, 2023.

561 Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and
 562 Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:
 563 133345–133385, 2024.

564 Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and
 565 Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the
 566 IEEE/CVF conference on computer vision and pattern recognition*, pp. 10696–10706, 2022.

567 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans
 568 trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, pp.
 569 6626–6637, 2017.

570 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*,
 571 2022.

572 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 573 neural information processing systems*, 33:6840–6851, 2020.

574 Vincent Tao Hu and Björn Ommer. [mask] is all you need. *arXiv preprint arXiv:2412.06787*, 2024.

575 Mengqi Huang, Zhendong Mao, Zhuowei Chen, and Yongdong Zhang. Towards accurate image
 576 coding: Improved autoregressive image generation with dynamic vector quantization. In *CVPR*,
 577 pp. 22596–22605, 2023.

578 Jiwan Hur, Dong-Jae Lee, Gyojin Han, Jaehyun Choi, Yunho Jeon, and Junmo Kim. Unlocking the
 579 Capabilities of Masked Generative Models for Image Synthesis via Self-Guidance. In *NeurIPS*,
 580 2024.

581 José Lezama, Huiwen Chang, Lu Jiang, and Irfan Essa. Improved masked image generation with
 582 token-critic. In *ECCV*, pp. 70–86. Springer, 2022.

583 Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xinchi Deng, Yingfang Zhang,
 584 Xingchao Liu, Minbin Huang, Zedong Xiao, et al. Hunyuan-dit: A powerful multi-resolution
 585 diffusion transformer with fine-grained chinese understanding. *arXiv preprint arXiv:2405.08748*,
 586 2024.

594 Zhiyuan Liu, Yicun Yang, Yaojie Zhang, Junjie Chen, Chang Zou, Qingyuan Wei, Shaobo Wang,
 595 and Linfeng Zhang. dLLM-Cache: Accelerating Diffusion Large Language Models with Adaptive
 596 Caching. *arXiv preprint arXiv:2506.06295*, 2025.

597 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast
 598 solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*,
 599 2022.

600 Nanye Ma, Mark Goldstein, Michael S. Albergo, Nicholas M. Boffi, Eric Vanden-Eijnden, and Saining
 601 Xie. *SiT: Exploring Flow and Diffusion-Based Generative Models with Scalable Interpolant
 602 Transformers*, pp. 23–40. Springer Nature Switzerland, 2024.

603 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-
 604 Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*,
 605 2025.

606 Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance
 607 for discrete state-space diffusion and flow models. *arXiv preprint arXiv:2406.01572*, 2024.

608 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 609 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
 610 robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

611 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li.
 612 Your absorbing discrete diffusion secretly models the conditional distributions of clean data. *arXiv
 613 preprint arXiv:2406.03736*, 2024.

614 Ziqi Pang, Tianyuan Zhang, Fujun Luan, Yunze Man, Hao Tan, Kai Zhang, William T Freeman, and
 615 Yu-Xiong Wang. Randar: Decoder-only autoregressive visual generation in random orders. *arXiv
 616 preprint arXiv:2412.01827*, 2024.

617 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *CVPR*, pp.
 618 4195–4205, 2023.

619 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 620 resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10674–10685, 2022a.

621 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 622 resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022b.

623 Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu,
 624 Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
 625 models. *NeurIPS*, 37:130136–130184, 2024.

626 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.
 627 Improved techniques for training gans. In *NeurIPS*, pp. 2226–2234, 2017.

628 Javier E Santos, Zachary R Fox, Nicholas Lubbers, and Yen Ting Lin. Blackout diffusion: generative
 629 diffusion models in discrete-state spaces. In *ICML*, pp. 9034–9059. PMLR, 2023.

630 Neta Shaul, Itai Gat, Marton Havasi, Daniel Severo, Anuroop Sriram, Peter Holderith, Brian Karrer,
 631 Yaron Lipman, and Ricky TQ Chen. Flow matching with general discrete paths: A kinetic-optimal
 632 perspective. *arXiv preprint arXiv:2412.03487*, 2024.

633 Fengyuan Shi, Zhuoyan Luo, Yixiao Ge, Yujiu Yang, Ying Shan, and Limin Wang. Scalable Image
 634 Tokenization with Index Backpropagation Quantization. *arXiv preprint arXiv:2412.02692*, 2025.

635 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized
 636 masked diffusion for discrete data. *NeurIPS*, 37:103131–103167, 2024.

637 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv
 638 preprint arXiv:2010.02502*, 2020.

639 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced
 640 transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

648 Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan.
 649 Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint*
 650 *arXiv:2406.06525*, 2024.

651

652 Alexander Sverdlov, Mihir Prabhudesai, Siddharth Gandhi, Deepak Pathak, and Katerina Fragki-
 653 adaki. Unified multimodal discrete diffusion. *arXiv preprint arXiv:2503.20853*, 2025.

654

655 Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling:
 656 Scalable image generation via next-scale prediction. *NeurIPS*, 37:84839–84865, 2024.

657

658 Shuai Wang, Zexian Li, Qipeng zhang, Tianhui Song, Xubin Li, Tiezheng Ge, Bo Zheng, and Limin
 659 Wang. Differentiable Solver Search for Fast Diffusion Sampling. *arXiv preprint arXiv:2505.21114*,
 660 2025.

661

662 Xiaoping Wu, Jie Hu, and Xiaoming Wei. Rdpm: Solve diffusion probabilistic models via recurrent
 663 token prediction. *arXiv preprint arXiv:2412.18390*, 2024.

664

665 Yi Xin, Le Zhuo, Qi Qin, Siqi Luo, Yuwen Cao, Bin Fu, Yangfan He, Hongsheng Li, Guangtao
 666 Zhai, Xiaohong Liu, and Peng Gao. Resurrect Mask AutoRegressive Modeling for Efficient and
 667 Scalable Image Generation. *arXiv preprint arXiv:2507.13032*, 2025.

668

669 Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada:
 670 Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*, 2025.

671

672 Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong
 673 Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan.
 674 In *ICLR*, 2022.

675

676 Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong
 677 Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion-
 678 tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023.

679

680 Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen.
 681 An image is worth 32 tokens for reconstruction and generation. *Advances in Neural Information
 682 Processing Systems*, 37:128940–128966, 2024.

683

684 Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and
 685 Saining Xie. Representation alignment for generation: Training diffusion transformers is easier
 686 than you think. In *ICLR*, 2025.

687

688 Biao Zhang and Rico Sennrich. Root mean square layer normalization. *NeurIPS*, 32, 2019.

689

690 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 691 diffusion models. In *ICCV*, pp. 3836–3847, 2023.

692

693 Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Masked
 694 diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical
 695 sampling. *arXiv preprint arXiv:2409.02908*, 2024.

696

697

698

699

700 **A USE OF LARGE LANGUAGE MODELS**

701

702 In the process of drafting this paper, Large Language Models (LLMs) were solely utilized for **writing**
polish (e.g., optimizing sentence structure, enhancing expression fluency). No LLM was involved in
 703 core academic work such as conceptualization, literature review, data analysis, argument construction,
 704 or conclusion formulation of this study.

705

706 As the human authors of this paper, we bear full and sole responsibility for the paper’s content,
 707 including the accuracy of research data, validity of academic arguments, integrity of research methods,
 708 and compliance with academic ethics.

702 B DISCRETE DIFFUSION WITH REHASHING NOISE

704 **Complete Definition and Deduction.** We provide a full theoretical discussion on the corrupted
 705 distribution and reverse process defined in the main paper.

706 Given d categories, let $\mathbf{e}_i \in \mathbb{R}^d$ be its one-hot vector where the i -th value is 1. We denote $\mathcal{E} = \{\mathbf{e}_i \in \mathbb{R}^d \mid i = 1, \dots, d\}$ as the basis of a categorical distribution, and a basis for absorbing states with
 707 capacity m : $\mathcal{M} = \{\mathbf{m}_j \in \mathbb{R}^m \mid j = 1, \dots, m\}$. The sum of \mathcal{E} and \mathcal{M} can be denoted as
 708

$$710 \quad \mathcal{V}_{(d,m)} \triangleq \left\{ \mathbf{v}_{(i,j)} \in \mathbb{R}^{d+m} \mid \mathbf{v}_{(i,j)} = \begin{cases} \mathbf{e}_i \oplus \mathbf{0}_m, & \text{for } i = 1, \dots, d, j = 0 \\ \mathbf{0}_d \oplus \mathbf{m}_j, & \text{for } j = 1, \dots, m, i = 0 \end{cases} \right\}. \quad (12)$$

712 We further denote the subspace \mathcal{E}_d , $\mathcal{M}_m \in \mathcal{V}_{(d,m)}$ which contain either valid or mask tokens, as

$$713 \quad \mathcal{E}_d = \{\mathbf{v}_{(i,0)} \in \mathcal{V}_{(d,m)} \mid i = 1, \dots, d\}, \quad \mathcal{M}_m = \{\mathbf{v}_{(0,j)} \in \mathcal{V}_{(d,m)} \mid j = 1, \dots, m\}. \quad (13)$$

714 To exploit visits across all the possible paths, for $0 \leq s < t \leq 1$, we write the transition kernel as⁴

$$716 \quad q(x_t^i \mid x_s^i) = \begin{cases} 1 - \alpha_{t|s}^{\leftarrow}, & \text{if } x_t^i \in \mathcal{M}_m, x_s^i \notin \mathcal{M}_m, \\ \alpha_{t|s}^{\leftarrow}, & \text{if } x_t^i = x_s^i, x_s^i \notin \mathcal{M}_m, \\ 1/m, & \text{if } x_t^i \in \mathcal{M}_m, x_s^i \in \mathcal{M}_m, \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

721 **Proof of the Corrupted Distribution.** The presentation in the main paper simplifies the theory
 722 without specifying the transition matrix Q_t due to page limitation. We make a detailed version with
 723 important yet basic matrix calculation in this section.

724 Let $\mathbf{I}_{(d,m)}$, $\mathbf{M}_{(d,m)}$ and $\boldsymbol{\pi}_{(d,m)}$ be matrices in $\mathbb{R}^{(d+m) \times (d+m)}$, defined as

$$726 \quad \mathbf{I}_{(d,m)} = \begin{bmatrix} \mathbf{I}_d & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{M}_{(d,m)} = \begin{bmatrix} 0 & \frac{1}{m} \mathbf{1}_{d \times m} \\ 0 & 0 \end{bmatrix}, \quad \boldsymbol{\pi}_{(d,m)} = \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{m} \mathbf{1}_m \mathbf{1}_m^\top \end{bmatrix} \quad (15)$$

729 where \mathbf{I}_d is the $d \times d$ identity matrix, and $\mathbf{1}_m \in \mathbb{R}^m$ is a vector of ones.

730 The transition matrix $Q_{t|s} \in \mathbb{R}^{(d+m) \times (d+m)}$ is defined as:

$$732 \quad Q_{t|s} = \alpha_{t|s}^{\leftarrow} \mathbf{I}_{(d,m)} + (1 - \alpha_{t|s}^{\leftarrow}) \mathbf{M}_{(d,m)} + \boldsymbol{\pi}_{(d,m)} \quad (16)$$

733 which can be demonstrated intuitively:

$$735 \quad Q_{t|s} = \underbrace{\begin{bmatrix} \alpha_{t|s}^{\leftarrow} & 0 & \cdots & 0 & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} & \cdots & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} \\ 0 & \alpha_{t|s}^{\leftarrow} & \cdots & 0 & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} & \cdots & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha_{t|s}^{\leftarrow} & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} & \cdots & \frac{1-\alpha_{t|s}^{\leftarrow}}{m} \\ 0 & 0 & \cdots & 0 & \frac{1}{m} & \frac{1}{m} & \cdots & \frac{1}{m} \\ 0 & 0 & \cdots & 0 & \frac{1}{m} & \frac{1}{m} & \cdots & \frac{1}{m} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & \frac{1}{m} & \frac{1}{m} & \cdots & \frac{1}{m} \end{bmatrix}}_{\times d} \underbrace{\times m}_{\times m}$$

747 The corrupted data distribution is a direct derivative of Eq. 16 by setting $s = 0$:

$$748 \quad x_t = x_0 Q_{t|0} \\ 749 \quad = \alpha_t x_0 \mathbf{I}_{(d,m)} + (1 - \alpha_t) x_0 \mathbf{M}_{(d,m)} + x_0 \boldsymbol{\pi}_{(d,m)} \\ 750 \quad = \alpha_t x_0 + (1 - \alpha_t) x_0 \mathbf{M}_{(d,m)} \\ 751 \quad \sim \alpha_t x_0 + (1 - \alpha_t) \mathbf{U}(\mathcal{M}_m^L) \quad (17)$$

753 where $\mathbf{U}(\mathcal{M}_m^L)$ is the uniform distribution on \mathcal{M}_m^L .

755 ⁴To maintain simplicity, we use $\alpha_{t|s}^{\leftarrow} = \frac{\alpha_t}{\alpha_s}$ and $\alpha_{t|s}^{\rightarrow} = \frac{1-\alpha_s}{1-\alpha_t}$ to denote transition rate for the corruption and
 reverse process, respectively.

Proof of the Reverse Process. To generate a sequence of length L , the reverse process starts with $x_1 \sim U(\mathcal{M}_m^L)$. Let $a \odot b$ denote the Hadamard product between two vectors a and b , the reverse process is inferred as:

$$\begin{aligned}
q(x_s \mid x_t) &= \frac{Q_{t|s}x_t \odot Q_{s|0}^\top x_0}{x_t^\top Q_{t|0}^\top x_0} \quad (\text{D3PM deduction}) \\
&= \frac{[\alpha_{t|s}^\leftarrow \mathbf{I}_{(d,m)} x_t + (1 - \alpha_{t|s}^\leftarrow) \mathbf{M}_{(d,m)} x_t + \boldsymbol{\pi}_{(d,m)} x_t] \odot [\alpha_s x_0 + (1 - \alpha_s) \mathbf{M}_{(d,m)}^\top x_0]}{x_t^\top [\alpha_t x_0 + (1 - \alpha_t) \mathbf{M}_{(d,m)}^\top x_0 + \boldsymbol{\pi}_{(d,m)}^\top x_0]} \\
&= \frac{[\alpha_{t|s}^\leftarrow \mathbf{I}_{(d,m)} x_t + (1 - \alpha_{t|s}^\leftarrow) \mathbf{M}_{(d,m)} x_t + \boldsymbol{\pi}_{(d,m)} x_t] \odot [\alpha_s x_0 + (1 - \alpha_s) \mathbf{M}_{(d,m)}^\top x_0]}{\alpha_t x_t^\top x_0 + (1 - \alpha_t) x_t^\top \mathbf{M}_{(d,m)}^\top x_0} \quad (18)
\end{aligned}$$

We consider the separate cases: $x_t^i = x_0^i$ and $x_t^i \in \mathcal{M}_m$.

Case 1. For $x_t^i = x_0^i$, Eq. 18 is simplified as

$$q(x_s^i \mid x_t^i = x_0^i) = \frac{\alpha_{t|s}^{\leftarrow} x_0^i \odot \alpha_s x_0^i}{\alpha_t x_0^i \top x_0^i} = 1 \quad (19)$$

Case 2. For $x_t^i \in \mathcal{M}_m$, we have

$$\begin{aligned}
q(x_s^i \mid x_t^i \in \mathcal{M}_m) &= \frac{[(1 - \alpha_{t|s}^{\leftarrow}) \mathbf{M}_{(d,m)} x_t^i + \boldsymbol{\pi}_{(d,m)} x_t^i] \odot [\alpha_s x_0 + (1 - \alpha_s) \mathbf{M}_{(d,m)}^\top x_0]}{(1 - \alpha_t) x_t^{i \top} \mathbf{M}_{(d,m)}^\top x_0} \\
&= \frac{[(1 - \alpha_{t|s}^{\leftarrow}) \alpha_s \mathbf{M}_{(d,m)} x_t^i \odot x_0 + \boldsymbol{\pi}_{(d,m)} (1 - \alpha_s) x_t^i \odot \mathbf{M}_{(d,m)}^\top x_0]}{(1 - \alpha_t) x_t^{i \top} \mathbf{M}_{(d,m)}^\top x_0} \\
&= \frac{(\alpha_s - \alpha_t) \mathbf{M}_{(d,m)} x_t^i \odot x_0 + (1 - \alpha_s) \boldsymbol{\pi}_{(d,m)} x_t^i \odot \mathbf{M}_{(d,m)}^\top x_0}{(1 - \alpha_t) x_t^{i \top} \mathbf{M}_{(d,m)}^\top x_0} \tag{20}
\end{aligned}$$

Notice that $\alpha_{t|s}^{\rightarrow} = \frac{1-\alpha_s}{1-\alpha_t}$, and we have

$$q(x_s^i \in \mathcal{M}_m \mid x_t^i \in \mathcal{M}_m) = \frac{1 - \alpha_s}{m(1 - \alpha_t)} = \frac{\alpha_t^{\rightarrow}}{m} \quad (21)$$

$$q(x_s^i \notin \mathcal{M}_m \mid x_t^i \in \mathcal{M}_m) = \frac{\alpha_s - \alpha_t}{1 - \alpha_s} = 1 - \alpha_{t|s}^{\rightarrow} \quad (22)$$

Combining case 1 with case 2, we have

$$q(x_s^i \mid x_t^i) = \begin{cases} 1, & \text{if } x_s^i = x_t^i, x_t^i \notin \mathcal{M}_m, \\ \alpha_{t|s}^i / m, & \text{if } x_s^i \in \mathcal{M}_m, x_t^i \in \mathcal{M}_m, \\ 1 - \alpha_{t|s}^i, & \text{if } x_s^i \notin \mathcal{M}_m, x_t^i \in \mathcal{M}_m, \\ 0, & \text{otherwise.} \end{cases} \quad (23)$$

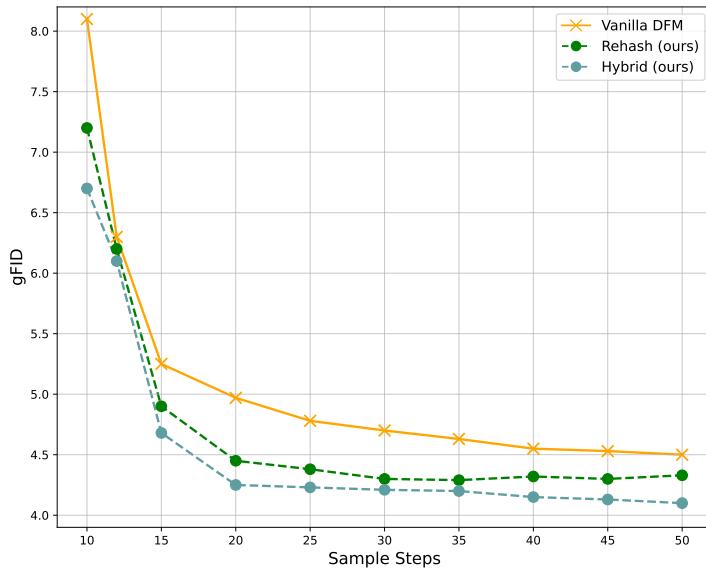
Following MDLM's deduction, assume that the denoising network can reconstruct x_0 perfectly, we use $p_\theta(x_t)$ to approximate this reverse process for complex sequences, and get

$$q(x_s^i|x_t) = \begin{cases} 1, & \text{if } x_s^i = x_t^i, x_t^i \notin \mathcal{M}_m, \\ \overrightarrow{\alpha_{t|s}}/m, & \text{if } x_s^i \in \mathcal{M}_m, x_t^i \in \mathcal{M}_m, \\ (1 - \overrightarrow{\alpha_{t|s}})p_\theta^i(x_t), & \text{if } x_s^i \notin \mathcal{M}_m, x_t^i \in \mathcal{M}_m, \\ 0, & \text{otherwise.} \end{cases} \quad (24)$$

810
811
C COMPLEXITY AND LIMITATIONS

812 Our rehash sampler follows DDM-theoretic principles and is implemented using `torch.multinomial`,
 813 which internally relies on Gumbel-max (Zheng et al., 2024). As a result, its asymptotic computational
 814 complexity is comparable to that of the original MaskGIT sampler. Practical efficiency gains,
 815 however, stem from two key factors. First, discrete visual tokens exhibit high correlation and
 816 redundancy, allowing multiple tokens to be predicted simultaneously and reducing the number of
 817 required operations compared with fully autoregressive or diffusion-based approaches. Second, DDM
 818 decouples the scheduler during training and sampling (Sahoo et al., 2024), enabling high-quality
 819 reconstruction from noise using arbitrarily defined timesteps. Using a cosine scheduler, ReDDiT
 820 achieves improved sample quality within fewer steps, as validated empirically.

821 Despite its effectiveness, the rehash sampler has limitations that motivate future work. The current
 822 rehashing strategy is applied at every step, but the impact on diversity and convergence is not fully
 823 characterized, suggesting that optimized rehash frequency or intensity could improve performance.
 824 Additionally, integrating ReDDiT with dynamic token-update mechanisms in DFM frameworks
 825 may further enhance sample quality and diversity, particularly for complex multimodal generation
 826 tasks. These considerations highlight potential directions for extending and refining discrete diffusion
 827 sampling methods.

828
829 **D SAMPLING FROM LEARNED NETWORKS**
830
831
832
833

853 Figure 7: Generation quality comparison with DFM methods. The experiments are conducted on
 854 ReDDiT-L with a constant classifier-free guidance ($\text{cfg} = 2.0$).

856 We present a detailed version of discrete flow matching (DFM) sampler 3, and discuss the integration
 857 of it with ours. Fig. 7 presents a quantitative comparison of the vanilla DFM sampler, our proposed
 858 rehash sampler, and a hybrid strategy that combines both approaches by incorporating selected DFM
 859 steps into the rehash trajectory. All methods are evaluated using identical model weights, as the
 860 training objectives are compatible due to their shared time-weighted loss formulation.

862 The rehash sampler exhibits stronger overall performance than DFM, especially in the 15–32 step
 863 range, where it achieves low and stable gFID scores. This suggests that our modification enables
 more efficient decoding trajectories without sacrificing sample quality. The hybrid variant, which

864	Algorithm 3 DFM Sampling Stepwise Pseudo Code
865	
866	Require: x_t , labels, timestep t , step size Δt
867	1: Compute jump probabilities: $j_t \leftarrow 1 - \alpha_t, j_s \leftarrow 1 - \alpha_{t-\Delta t}$
868	2: Determine guidance scale ω from schedule
869	3: Obtain logits $\text{logits}_{\text{cond}}$, $\text{logits}_{\text{uncond}}$ via forward pass
870	4: $\text{logits}_{x_0} \leftarrow \text{logits}_{\text{uncond}} + \omega \cdot (\text{logits}_{\text{cond}} - \text{logits}_{\text{uncond}})$
871	5: $p_{x_0} \leftarrow \text{softmax}(\text{logits}_{x_0})$
872	6: Sample $\hat{x}_0 \sim p_{x_0}$ using categorical sampling
873	7: Construct one-hot encodings: $\delta_{x_0}, \delta_{x_t}$
874	8: corrective $\leftarrow \frac{j_s}{j_t} \cdot \delta_{x_t}$
875	9: $u \leftarrow \frac{j_t - j_s}{j_t} \cdot \delta_{x_0}$
876	10: Overwrite u in masked range with corrective terms
877	11: Mask entries already present in x_t from u
878	12: Compute total transition intensity: $\lambda \leftarrow \sum u$, elementwise
879	13: Draw Bernoulli mask: $M \sim \text{Bernoulli}(1 - \exp(-\lambda))$
880	14: For each masked position in M , sample from categorical u to obtain updated x_s
881	15: return x_s
882	

integrates only the middle and final steps of the DFM update into the rehash schedule, also delivers consistent gains over the vanilla DFM, suggesting that partial refinement from DFM is beneficial even when the majority of the trajectory is governed by our rehash dynamics.

By leveraging shared gradual decoding infrastructure, the hybrid approach enables practical integration of DFM refinement into the ReDDiT framework with minimal overhead. As noted in the main paper, this leads to a ~ 0.1 improvement in gFID on ImageNet-1K, reinforcing the complementary strengths of the two samplers. We leave the comprehensive study on the optimal integration of different samplers for future exploration.

E EXPERIMENT DETAILS

We provide detailed training and generation configurations for ReDDiT in Table 5. Our method incorporates DINOv2-B for representation alignment, which requires computing image features during the forward pass (only activated during training). This introduces an overhead, making training roughly $1.2\times$ slower than solely on discrete tokens. However, this additional cost is offset by faster convergence and improved stability, particularly in early training stages.

The use of quantized latents allows for larger batch sizes under limited GPU memory, making our approach more accessible for low-resource settings. Additionally, aligning discrete codes with semantic features improves the quality and diversity of learned representations. Overall, our design balances computational efficiency with model performance, making it a practical choice for both research and deployment.

F DISCUSSION ON REPRESENTATION ALIGNMENT

Representation Alignment (RepA) introduces a similarity-based auxiliary loss that aligns intermediate features of the diffusion model with pretrained DINOv2 embeddings. Although originally proposed for continuous diffusion, its effectiveness naturally extends to the discrete setting. In ReDDiT, discrete tokens are first mapped into a continuous embedding space, after which the architecture is identical to transformer-based continuous diffusion models (e.g., DiT). This means that the absence of continuous input does not fundamentally alter the structure of the model’s internal representations. However, discrete tokenizers restrict direct gradient flow from pixels to the codebook, making it more difficult for the model to organize high-level semantics during early training.

RepA provides an external semantic scaffold that compensates for this difficulty. By encouraging the network to match DINOv2’s robust visual features, RepA helps establish meaningful structure in the latent representations before the denoising objective becomes sufficiently informative. Empirically,

918 Table 5: Experiment details for ReDDiT on ImageNet-1K. *Vari.* refers to a time-variant growing
 919 guidance scale following MDTv2, which is a common practice for diffusion models.
 920

921 Setting	922 ReDDiT-L (Ablation)	923 ReDDiT-L	924 ReDDiT-XL	925 ReDDiT-XL_{f8}
926 Hidden Size	1024	1024	1280	1280
927 Transformer Block	24	24	28	28
928 Attention Head	16	16	20	20
929 Image Tokenizer	LlamaGen-f16	IBQ-f16	IBQ-f16	LlamaGen-f8
930 Codebook Size	16384	16384	16384	16384
931 Noise Capacity	128	1024	1024	128
932 Sequence Length	256	256	256	1024
933 RepA Latent Size	16×16	16×16	16×16	32×32
934 Batch Size	64	64	32	16
935 Global Batch Size	1024	1024	1024	1024
936 LR scheduler	Cosine Decay	Cosine Decay	Cosine Decay	Cosine Decay
937 Learning Rate	3e-4	3e-4	3e-4	4e-4
938 Minimal LR	1e-5	1e-5	1e-5	1e-5
939 Warmup Steps	2k	2k	2k	2k
940 Training Steps	500k	500k	500k	500k
941 Training Time	~1 day	~1 day	~2 days	~3 days
942 Generation CFG (<i>Vari.</i>)	1.0-5.0	1.0-6.5	1.0-6.5	1.0-5.5

937
938 Table 6: **Acceleration of ReDDiT using response cache K_r .**
939

940 Model	941 Config		942 Performance	
	943 Steps	944 K_r	945 Relative Speed	946 gFID↓
947 ReDDiT-L	32	2	$\times 1.33$	2.28 ($\Delta = 0.15$)
			$\times 1.52$	1.88 ($\Delta = 0.14$)
948 ReDDiT-XL	64	4	$\times 2.17$	1.83 ($\Delta = 0.09$)
			$\times 2.56$	1.71 ($\Delta = 0.10$)

948 we observe that removing RepA leads to an early training plateau, whereas with RepA the alignment
 949 loss rapidly decreases and stabilizes, indicating improved organization of semantic information. This
 950 behavior mirrors observations in continuous diffusion models and supports the view that RepA offers
 951 a general mechanism for accelerating convergence, independent of whether the base diffusion process
 952 is discrete or continuous.

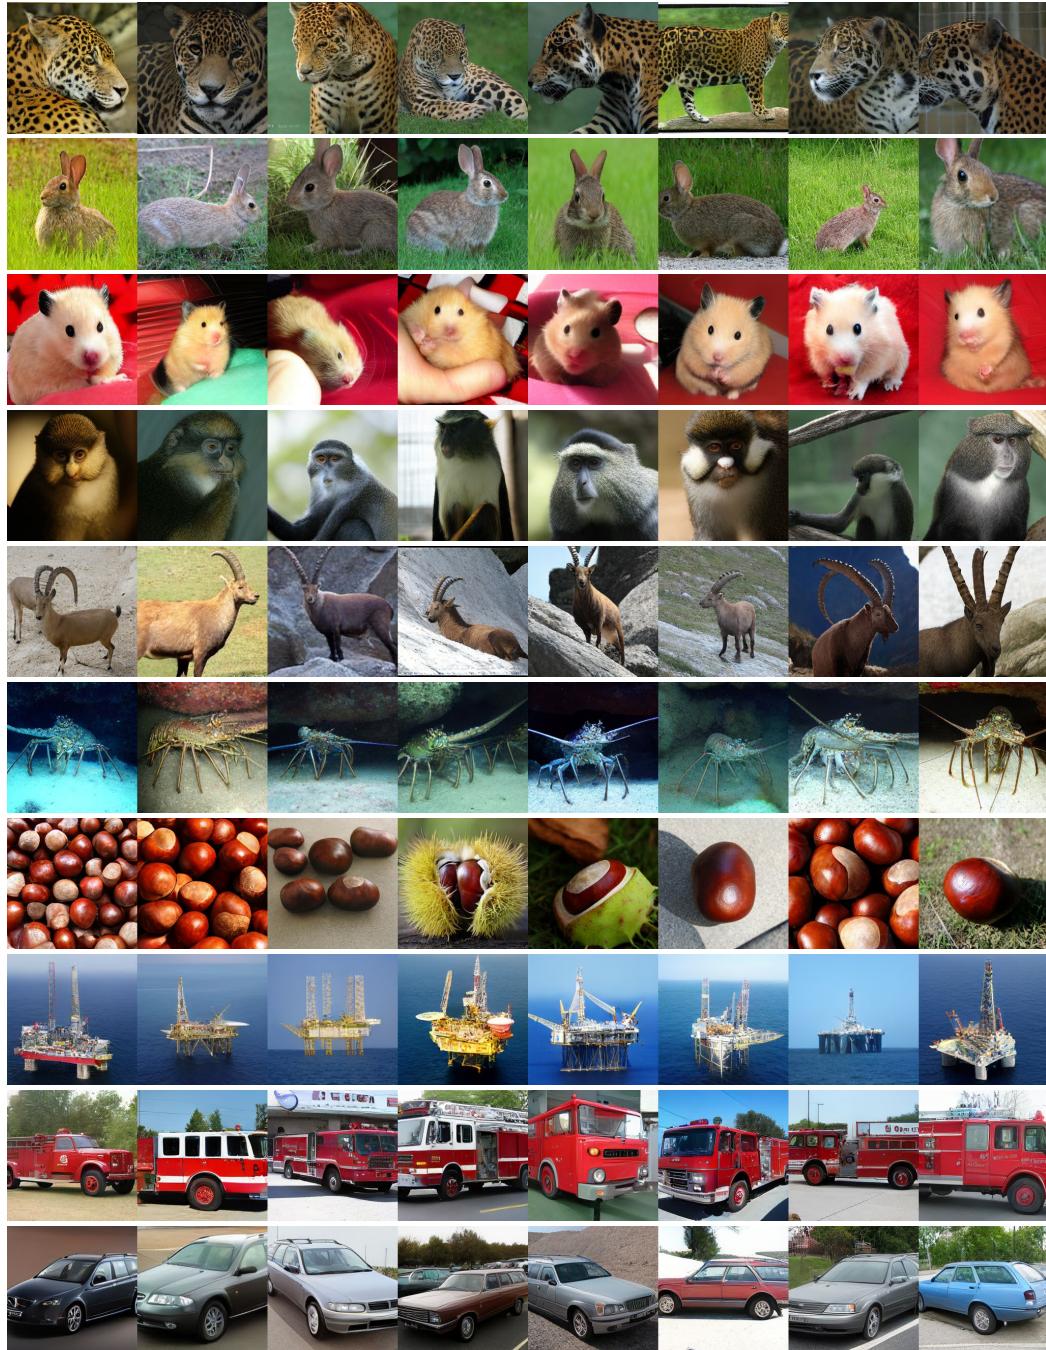
953 G ACCELERATING REDDIT

954 Recent efforts on scaling and accelerating discrete diffusion models are making this generative
 955 paradigm more practical than theoretical attempts. We adapt the dLLM-Cache Liu et al. (2025) design
 956 into our framework, which efficiently reuses intermediate computations without compromising model
 957 performance. Since the condition is modulated using AdaLN and introduces minimal calculation,
 958 we do not activate K_p (cache for prompt). As the decoding of visual sequence varies with time
 959 more quickly than language decoding, we implement the cache for response with small values like
 960 $K_r = 2$ or 4 , which means the K and V of transformer layer is updated every 2 or 4 decoding steps
 961 instead of per step. As shown in Tab. 6, the inference speed is boosted up to 2 times with minimal
 962 performance drop, which makes our largest model ReDDiT-XL_{f8} comparable to diffusion models
 963 with accelerated solvers.

964 H QUALITATIVE RESULTS

965 We provide more samples of ReDDiT’s generation in Fig. 8.

972
973
974
975
976



1020
1021
1022
1023
1024
1025

Figure 8: Class-conditional generation samples of ReDDiT on ImageNet 256×256 .