

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 CONTINUOUS SPECULATIVE DECODING FOR AUTOREGRESSIVE IMAGE GENERATION

005 **Anonymous authors**

006 Paper under double-blind review



022 Figure 1: Continuous speculative decoding accelerates the inference speed while maintaining the
023 original generation quality. For each panel, left: image generated by MAR; right: image generated
024 by MAR with continuous speculative decoding. Speed-up ratio of each single image is presented.

025 ABSTRACT

028 Continuous visual autoregressive (AR) models have demonstrated promising per-
029 formance in image generation. However, the heavy autoregressive inference bur-
030 den imposes significant overhead. In Large Language Models (LLMs), specula-
031 tive decoding has effectively accelerated discrete autoregressive inference. How-
032 ever, the absence of an analogous theory for continuous distributions precludes its
033 use in accelerating continuous AR models. To fill this gap, this work presents con-
034 tinuous speculative decoding, and addresses challenges from: 1) low acceptance
035 rate, caused by inconsistent output distribution between target and draft models,
036 and 2) modified distribution without analytic expression, caused by complex inte-
037 gral. To address challenge 1), we introduce an approximate acceptance criterion to
038 address the inefficiency in sampling. Furthermore, we propose denoising trajec-
039 tory alignment and token pre-filling strategies. To address challenge 2), we intro-
040 duce acceptance-rejection sampling algorithm with an appropriate upper bound,
041 thereby avoiding explicitly calculating the integral. Furthermore, our denoising
042 trajectory alignment is also reused in acceptance-rejection sampling, effectively
043 avoiding repetitive diffusion model inference. Extensive experiments demonstrate
044 that our proposed continuous speculative decoding achieves over 2 \times speedup on
045 off-the-shelf models, while maintaining the generation quality.

046 1 INTRODUCTION

048 Autoregressive (AR) models have demonstrated significant potential and achieved competitive per-
049 formance in image generation tasks (Van Den Oord et al., 2016; Van den Oord et al., 2016; Esser
050 et al., 2021; Yu et al., 2023a; Tian et al., 2025). These models predict next token sequentially based
051 on previously generated tokens, outputting discrete categorical distributions. Typically, the input
052 image is mapped from the pixel space to tokens through vector quantization (VQ), after which the
053 AR model generates images by predicting discrete distributions for subsequent tokens. While this
process shows substantial promise in image generation, VQ can lead to training instability and might

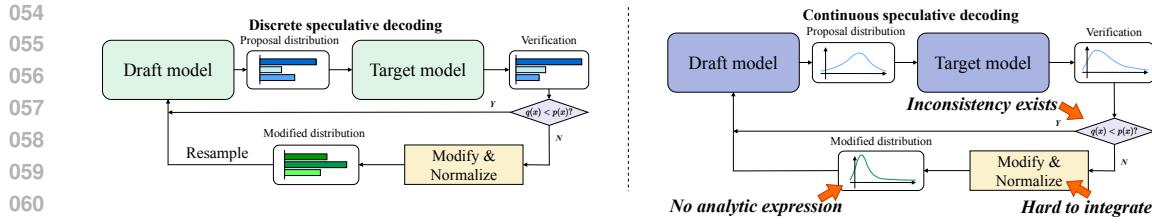


Figure 2: Comparison between discrete and continuous speculative decoding. Discrete situation offers the convenience of directly computing probabilities and simply sampling from modified distributions. In contrast, continuous situation faces challenges in the inconsistency of output distributions, leading to low acceptance criterion as well as low acceptance rate, and the modified distributions without analytic expression, caused by complex integral.

not adequately capture the nuanced image details (Yu et al., 2023a; Mentzer et al., 2023). Recently, continuous visual AR models have been proposed to predict visual tokens from continuous probability distributions (Tschanne et al., 2025; Li et al., 2025; Xu et al., 2024; Sucheng et al., 2025). In this framework, the model predicts next token’s continuous distribution based on prior outputs, often implemented via denoising processes (Ho et al., 2020; Nichol & Dhariwal, 2021). This approach overcomes the limitations of VQ, offering a promising solution for autoregressive image generation.

However, continuous visual AR models suffer from slow inference speed due to the sequential decoding process, a limitation shared with LLMs. LLMs commonly employ speculative decoding (Leviathan et al., 2023; Chen et al., 2023) to accelerate autoregressive inference procedure. This algorithm employs a draft-and-verification mechanism, where a smaller draft model generates several draft tokens, and then a more accurate yet often larger target model verifies them. Inspired by this, previous works have investigated the efficacy of speculative decoding for accelerating discrete visual AR models (Jang et al., 2024; Teng et al., 2024). However, the related methods are formulated for discrete output distributions, whereas a continuous formulation remains undeveloped.

To bridge this gap, we present **Continuous Speculative Decoding**, designed to accelerate the inference of continuous visual AR models. However, as illustrated in Figure 2, the following two challenges can be identified while developing this method: a) **Inconsistent output distribution**. There is significant inconsistency in diffusion and autoregressive procedures, resulting in distinct draft and target output distributions. This leads to a very low acceptance criterion, thus lowering the acceptance rate. b) **Modified distribution without analytic expression**. When an output token is rejected, a new token will be drawn from the modified distribution. In continuous distributions, there is no analytic expression for this due to the complex integral for normalizing the probability distribution, making it impossible to directly sample from this distribution.

To address the aforementioned challenges, this work proposes a series of technical optimizations. **First**, to improve the acceptance rate, we introduce an approximate acceptance criterion to address the inefficiency in sampling. Furthermore, we develop denoising trajectory alignment method based on the proposed theorem of *proximity in the reparameterization*. This method aims to reduce the distance between two output distributions, thereby resolving the inconsistency in the diffusion procedure. Additionally, we introduce a token pre-filling strategy to improve the early low acceptance rate. This strategy leverages the principle that the output of AR models depends on previous states, addressing inconsistencies within the autoregressive procedure. **Second**, to effectively sample from the complex modified distribution, we introduce the acceptance-rejection sampling algorithm (Casella et al., 2004) by deriving a proper upper bound, which helps to eliminate the need for computing complex integral. Besides, to avoid repetitive use of diffusion model for inference during acceptance-rejection sampling, we derive an easy-to-compute rejection threshold by reusing the denoising trajectory alignment.

Our continuous speculative decoding can be integrated seamlessly into many of existing models, as shown in Figure 1. We validate the effectiveness of our algorithm on three continuous visual AR models (Li et al., 2025; Sucheng et al., 2025; Wu et al., 2025) at two resolutions (256 & 512) through qualitative and quantitative evaluations. Specifically, we measure wall-time improvements and report image generation performance using Fréchet Inception Distance (FID) (Heusel et al., 2017)

108 and Inception Score (IS) (Salimans et al., 2016). Extensive experiments show that our algorithm
 109 achieves over $2\times$ inference speedup while maintaining generation quality.
 110

111 Our contributions can be summarized as follows:

- 112 • We are the first to propose continuous speculative decoding, bridging the gap of speculative
 113 decoding to continuous distributions and enabling substantial acceleration of continuous
 114 visual AR models.
- 115 • We address the distinct problem of low acceptance rate in continuous speculative decoding
 116 with an approximated acceptance criterion by introducing two novel techniques: denoising
 117 trajectory alignment, which reduces the draft-target distribution distance, and token pre-
 118 filling, which mitigates the early low acceptance rate problem.
- 119 • We solve the hard-to-sample problem of modified distribution via acceptance-rejection
 120 sampling with a proper upper bound to avoid the complex integral and an easy-to-compute
 121 rejection threshold with denoising trajectory alignment to avoid repetitive model inference.
- 122 • We seamlessly integrate our algorithm into three existing continuous visual AR models
 123 without extra training or architectural changes. Extensive experiments show that it achieves
 124 over $2\times$ inference speedup while maintaining generation quality.

127 2 RELATED WORK

128 2.1 AUTOREGRESSIVE IMAGE GENERATION

131 Autoregressive (AR) models are widely used in image generation. Early works perform generation
 132 at pixel level using CNN (Van den Oord et al., 2016; Chen et al., 2018), RNN (Van Den Oord
 133 et al., 2016) and Transformer (Parmar et al., 2018; Chen et al., 2020). VAR (Tian et al., 2025)
 134 modifies the autoregressive paradigm into next-scale prediction to gradually increase the scale of
 135 predictions. Similar to the autoregressive language model, AR image generation through discrete
 136 token prediction is scalable to text-conditioned image generation (Liu et al., 2024; Sun et al., 2024;
 137 Yu et al., 2023b). However, training discrete image tokenizer is difficult, and its ability to convey
 138 detailed visuals is still questionable (Yu et al., 2023a; Mentzer et al., 2023). GIVT (Tschannen
 139 et al., 2025) represents continuous tokens via Gaussian mixture models. MAR (Li et al., 2025) and
 140 DisCo-Diff (Xu et al., 2024) generate tokens via diffusion process (Ho et al., 2020) conditioned
 141 by the autoregressive model. HART (Tang et al., 2024) employs discrete and continuous tokenizer
 142 to generate images, with classification for discrete tokens and denoising for the residual between
 143 primitive visual tokens and discrete tokens. xAR (Sucheng et al., 2025) extends the conception
 144 of token and reformulates discrete token classification as continuous entity regression. However,
 145 autoregressive models suffer from heavy inference overhead. The inference speed is slowed down
 by step-by-step generation.

146 2.2 SPECULATIVE DECODING

148 Speculative decoding (Leviathan et al., 2023; Chen et al., 2023) achieves lossless acceleration by
 149 verifying the draft model with the target model. Following this, previous works mainly focus on re-
 150 ducing draft model overhead and strengthening the consistency between the draft and target models.
 151 SpecInfer (Miao et al., 2023) employs multiple small draft models and aggregates their predictions
 152 into a tree structure to be verified through tree-based parallel decoding. Eagle (Li et al., 2024b;c)
 153 improves the draft accuracy through the prediction at the feature level instead of the token level to
 154 tackle the feature uncertainty problem. Jacobi iteration is also employed to reduce inference over-
 155 head in the decoding process (Santilli et al., 2023; Zhao et al., 2024b;a; Kou et al., 2024). Online
 156 Speculative Decoding (Liu et al., 2023) and DistillSpec (Zhou et al., 2023) align the output from the
 157 draft model with the target model with more training.

158 Recent works are now using speculative decoding to improve the efficiency of autoregressive im-
 159 age generation. SJD (Teng et al., 2024) improves the Jacobi iteration process by adding speculative
 160 decoding while keeping the variety of image generation. LANTERN series (Jang et al., 2024; Park
 161 et al., 2025) looks at distribution ambiguity and uses relaxation to add more flexible candidate to-
 162 kens, maintaining high image quality. These studies have greatly speeds up the process by reducing

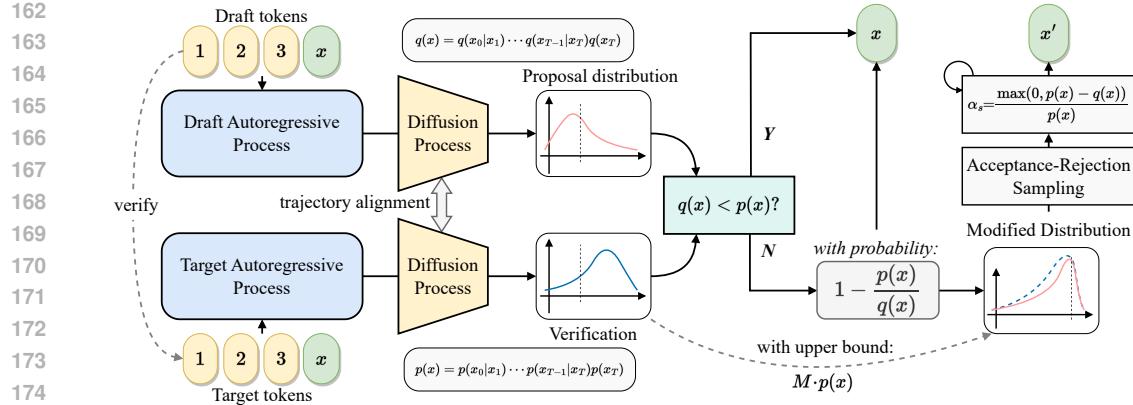


Figure 3: The overview of continuous speculative decoding. The diffusion model component of continuous AR models is leveraged. Tokens 1 ~ 3 are prefix tokens, and token x is to be verified. In this method, x is accepted if $q(x) < p(x)$. Otherwise, x is rejected with probability $1 - p(x)/q(x)$, followed by sampling x' from the modified distribution via acceptance-rejection sampling.

the inference steps needed for generating visual tokens. However, they are only applicable to discrete space. By contrast, our work extends speculative decoding to continuous AR models.

3 METHODOLOGY

3.1 FROM DISCRETE TO CONTINUOUS SPECULATIVE DECODING

We first introduce the discrete form of speculative decoding. It utilizes a draft model M_q with output distribution $q(x)$ and a target model M_p with $p(x)$. M_q generates a sequence of draft tokens $x_{i:j} = \{x_i, \dots, x_j\}$ where $x \sim q(x)$, which are then verified by M_p in parallel. x is accepted if the acceptance criterion $p(x)/q(x) > 1$; otherwise, it is rejected with probability $1 - p(x)/q(x)$ and resampled from a modified distribution $p'(x) = \text{norm}(\max(0, p(x) - q(x))) = \frac{\max(0, p(x) - q(x))}{\sum_{x'} \max(0, p(x') - q(x'))}$.

Then, we discuss continuous speculative decoding (see Figure 3). In continuous visual AR models (Li et al., 2025), the output distribution is typically modeled via diffusion models, namely:

$$p(x_{N,0:T} | x_{1:N-1}) = p(x_{N,T}) \prod_{t=1}^T p(x_{N,t-1} | x_{N,t}, x_{1:N-1}), \quad (1)$$

where N represents N -th autoregressive step. $t \in [0, T]$ is the diffusion timestep. $x_{N,t}$ represents the state of x at AR step N and diffusion timestep t . $x_{1:N-1} = \{x_1, x_2, \dots, x_{N-1}\}$ denotes the variables before N . For simplicity, we omit N and $x_{1:N-1}$ and let $Y = x_{0:T}$ to obtain $p(Y) = p(x_T) \prod_{t=1}^T p(x_{t-1} | x_t)$.

Acceptance Criterion The acceptance criterion is defined as the ratio of its probability under target distribution to the one under draft distribution, that is, $p(x_{0:T})/q(x_{0:T})$. In discrete form, the probability can be directly obtained. But in continuous form, the probability is usually obtained via diffusion process (Ho et al., 2020; Song et al., 2020). Specifically, $x_{0:T}$ is sampled from draft model via reverse diffusion (denoising) process. So the ratio is given by:

$$\frac{p(Y)}{q(Y)} = \frac{p(x_T) \prod_{t=1}^T p(x_{t-1} | x_t)}{q(x_T) \prod_{t=1}^T q(x_{t-1} | x_t)}, \quad (2)$$

where x_T is a Gaussian noise and $p(x_{t-1} | x_t)$ is approximated as a Gaussian distribution through a neural network θ (Nichol & Dhariwal, 2021), and $\mu_\theta(x_t, t)$ and $\Sigma_\theta(x_t, t)$ are mean and variance predicted by θ , that is:

$$p(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)). \quad (3)$$

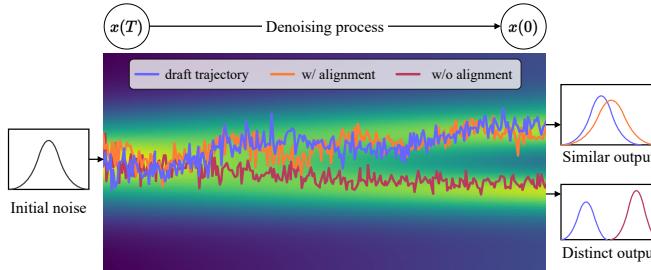
216
217
218
219
220
221
222
223
224
225

Figure 4: Illustration of denoising trajectory alignment. The denoising process maps the noise distribution to data distribution through gradual denoising. These denoising steps form a trajectory. The aligned trajectory (orange curve) leads to a similar output distribution, while the unaligned one (red curve) produces a far-away one, obtaining low $p(Y_p)/q(Y_q)$.

232
233
234
235
236
237

However, this approach is impractical because the acceptance rate is low. As discussed in Appendix A.1.1, the draft model’s trajectory $x_{0:T}$ inherently diverges from the target model’s expected trajectory. In each denoising step, samples drawn from the draft model’s distribution q are unlikely to fall near μ of the target distribution p , which results in a low single-step ratio p/q . As the multi-step denoising process proceeds, the overall $p(Y)/q(Y)$ becomes extremely small.

238
239
240

Therefore, our work adopts a practical approximation: the ratio of the joint probabilities $p(Y_p)/q(Y_q)$, where both Y_p and Y_q share the same x_0 . This ratio serves as a surrogate for the shared path ratio. Then, the ratio is calculated through:

241
242
243

$$\frac{p(Y_p)}{q(Y_q)} := \frac{p(x_T^p) \prod_{t=1}^T p(x_{t-1}^p | x_t^p)}{q(x_T^q) \prod_{t=1}^T q(x_{t-1}^q | x_t^q)}, \quad x_0^p = x_0^q = x_0, \quad (4)$$

244
245
246
247
248

Modified Distribution When a draft token is rejected, a new one is sampled from the modified distribution $p'(x)$. $p'(x)$ is derived from the normalization of $\max(0, p(x) - q(x))$. Therefore, replacing the summation in the normalization denominator of the discrete form with integral yields the continuous form, namely:

249
250
251

$$p'(Y) = \frac{\max(0, p(Y) - q(Y))}{Z}, \text{ where } Z = \int_{Y'} \max(0, p(Y') - q(Y')) dY'. \quad (5)$$

252
253
254
255
256

However, in practical operation, an extremely low acceptance rate emerges, as shown in Table 4, thereby leading to poor acceleration performance. To address this, we propose the following technical optimizations.

257

3.2 MITIGATING DISTRIBUTION INCONSISTENCY

258
259
260

Under the acceptance criterion given by Equation 4, the acceptance rate α is extremely small (nearly 0%). We attribute this to two sources of inconsistency in continuous visual AR models: **1) inconsistency in diffusion process**, and **2) inconsistency in autoregressive process**.

261
262
263
264
265

Inconsistency in Diffusion Process Significant inconsistency exists in the denoising process. As illustrated in Figure 4, draft and target denoising trajectories diverge to different outputs. The output distance between the two distributions is large. Therefore, for a draft output x_0 , $q(Y_q)$ is large while $p(Y_p)$ is quite small, leading to a low $p(Y_p)/q(Y_q)$, consequently leads to low acceptance rate α .

266
267
268
269

We propose denoising trajectory alignment to enhance the consistency of the output distributions. Note that in Equation 3, x_{t-1}^p is obtained via reparameterization given by $x_{t-1}^p = \sqrt{\Sigma_\theta^p(x_t^p, t)} \cdot \varepsilon_t^p + \mu_\theta^p(x_t^p, t)$, where $\varepsilon_t^p \sim \mathcal{N}(0, I)$ (same for x_{t-1}^q). Denoising trajectory alignment can reduce the expected distance between x_{t-1}^p and x_{t-1}^q , promised by the following theorem.

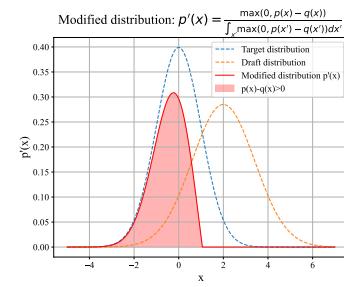


Figure 5: Illustration of distributions. Dashed lines: draft and target distributions. Red area: modified distribution (unnormalized, omitting Z for simplicity), whose integral lack an analytic expression.

270 **Theorem 1 (Proximity in the Reparameterization)** *Setting $\varepsilon_t^p = \varepsilon_t^q$ in reparameterization reduces the expected distance $\mathbb{E} [\|x_{t-1}^q - x_{t-1}^p\|^2]$ between x_{t-1}^p and x_{t-1}^q by $2 \cdot \text{tr} [\sqrt{\Sigma_t^q \Sigma_t^p}]$.*

273 Detailed proofs can be found in Appendix A. We emphasize that reducing the distance of output
274 samples aims to increase $p(x)$, thereby increasing $p(x)/q(x)$ for higher acceptance rate. Note that:
275

$$\begin{aligned} 276 \quad p(x_{t-1}|x_t) &= \frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma_\theta(x_t, t)|}} \exp \left\{ \frac{1}{2} [x_{t-1} - \mu_\theta(x_t, t)]^T \Sigma_\theta^{-1}(x_t, t) [x_{t-1} - \mu_\theta(x_t, t)] \right\} \\ 277 \\ 278 \quad &= \frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma_\theta(x_t, t)|}} \exp \left\{ \frac{1}{2} \varepsilon_t^T \varepsilon_t \right\}, \quad x_t \in \{x_t^p, x_t^q\}, \varepsilon_t \in \{\varepsilon_t^p, \varepsilon_t^q\}. \end{aligned} \quad (6)$$

281 The ratio $p(x_{t-1}^p|x_t^p)/q(x_{t-1}^q|x_t^q)$ given $\varepsilon_t = \varepsilon_t^p = \varepsilon_t^q$ is:
282

$$\begin{aligned} 283 \quad \frac{p(x_{t-1}^p|x_t^p)}{q(x_{t-1}^q|x_t^q)} &= \frac{\frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma_t^p|}} \exp \left\{ \frac{1}{2} \varepsilon_t^T \varepsilon_t \right\}}{\frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma_t^q|}} \exp \left\{ \frac{1}{2} \varepsilon_t^T \varepsilon_t \right\}} = \frac{\sqrt{|\Sigma_t^q|}}{\sqrt{|\Sigma_t^p|}}. \end{aligned} \quad (7)$$

286 For simplicity, we define:
287

$$\Sigma = \prod_{t=2}^T \sqrt{|\Sigma_t^q|} / \prod_{t=2}^T \sqrt{|\Sigma_t^p|}. \quad (8)$$

290 Substituting Σ into Equation 4 makes (assuming $p(x_T^p) = q(x_T^q)$):
291

$$\frac{p(Y_p)}{q(Y_q)} = \frac{p(x_T^p)p(x_0|x_1^p) \prod_{t=2}^T p(x_{t-1}^p|x_t^p)}{q(x_T^q)q(x_0|x_1^q) \prod_{t=2}^T q(x_{t-1}^q|x_t^q)} = \frac{p_\theta(x_0|x_1^p)}{q_\theta(x_0|x_1^q)} \cdot \Sigma. \quad (9)$$

295 Therefore, at each step, employing the same ε_t reduces the expected distance by $2 \cdot \text{tr} [\Sigma_t^q \Sigma_t^p]$, which
296 enables the two models to generate closer samples and finally increases the acceptance rate.
297

298 **Inconsistency in Autoregressive Process** Inconsistency also arises in autoregressive steps. Figure 299 shows that acceptance rate α is very low (5%) at the initial AR steps, and it increases pro-
300gressively as AR steps grow. This stems from the different draft and target prefix embeddings (Li
301 et al., 2025). Owing to this, the AR models naturally yield divergent predictions, which in turn
302 leads to low α . As the AR steps increase, the inputs of the two models gradually converge, thereby
303 improving consistency between their outputs and finally raising the α .

304 To address this, we propose pre-filling a portion (e.g., 5%) of tokens from the target model to ensure
305 a consistent prefix. This does not increase inference latency, as speculative decoding at a low ac-
306 ceptance rate is functionally equivalent to the target model step-by-step decoding (Leviathan et al.,
307 2023). Furthermore, pre-filling improves the overall acceptance rate.

308 Finally, $p(Y_p)/q(Y_q)$ can be computed with the help of denoising trajectory alignment and token
309 pre-filling to obtain a considerable acceptance rate.
310

3.3 RESAMPLE FROM THE MODIFIED DISTRIBUTION

313 The simplified illustration of Equation 5 is shown in Figure 5. Unlike discrete form, where
314 $\sum \max(0, p(x) - q(x))$ can be directly computed, the analytic expression of Z can't be computed
315 because it involves the integral of the product of a series of Gaussian distributions given by Equa-
316 tion 4 and 3. Therefore, $p'(Y)$ cannot be directly sampled.

317 To tackle this problem, a viable approach is acceptance-rejection sampling (Casella et al., 2004),
318 which first samples from a proposal distribution, and then calculates a pre-defined rejection threshold
319 α_s and samples $\epsilon \sim U(0, 1)$. If $\epsilon < \alpha_s$, the sample is accepted, otherwise it is rejected and sampling
320 from the proposal distribution will be repeated until the sample is accepted.

321 To realize this sampling method, we first define α_s as:
322

$$\alpha_s = \frac{p'(Y)}{M \cdot p(Y)}, \quad (10)$$



Figure 6: Qualitative results of images generated by MAR with continuous speculative decoding.

where $p(Y)$ is the target distribution, and M is the upper bound factor that holds $M \cdot p(Y) \geq p'(Y)$ for any Y . Given $\max(0, p(Y) - q(Y)) \leq p(Y)$, M is set to $1/Z$ considering that:

$$p'(Y) = \frac{\max(0, p(Y) - q(Y))}{Z} \leq \frac{p(Y)}{Z} \mapsto M \cdot p(Y). \quad (11)$$

We substitute $M = 1/Z$ into Equation 10 to eliminate Z :

$$\alpha_s = \frac{\max(0, p(Y) - q(Y))/Z}{p(Y)/Z} = \frac{\max(0, p(Y) - q(Y))}{p(Y)}. \quad (12)$$

Equation 12 gives the analytic expression of α_s without calculating Z . However, in its naive implementation, repetitive diffusion model inference is needed to sample $p(Y)$. It brings heavy extra overhead. To tackle this problem, denoising trajectory alignment is introduced to simplify Equation 12. Accordingly, we have derived the following corollary.

Corollary 1 (Easy-to-Compute Rejection Threshold) *With denoising trajectory alignment introduced from Equation 9, the rejection threshold α_s has the easy-to-compute form:*

$$\alpha_s = \frac{\max(0, \Sigma \cdot p_\theta(x_0|x_1^p) - q_\theta(x_0|x_1^q))}{\Sigma \cdot p_\theta(x_0|x_1^p)}. \quad (13)$$

See Appendix A for detailed proofs. In this form, x_0 is sampled from $p_\theta(x_0|x_1^p)$, which is merely a Gaussian distribution defined by Equation 3, thus avoiding extra model inference. We can compute α_s by gathering Σ and sampling x_0 from the Gaussian distribution, thereby completing the acceptance-rejection sampling to obtain samples equivalent to those derived from sampling $p'(Y)$.

4 EXPERIMENT

4.1 IMPLEMENTATION DETAILS

We systematically conduct experiments with open-source continuous visual AR model MAR (Li et al., 2025), xAR (Sucheng et al., 2025) (trained on ImageNet (Deng et al., 2009)) and Harmon (Wu et al., 2025) (trained on more data sources). MAR and xAR are evaluated under 256×256 resolution. Harmon is evaluated under both 256×256 and 512×512 resolutions. MAR-B (208M), xAR-B (172M) and Harmon-0.5B are selected as draft models. MAR-H (943M), xAR-H (1.1B) and Harmon-1.5B are selected as target models. We employ Fréchet Inception Distance (FID) (Heusel et al., 2017), Inception Score (IS) (Salimans et al., 2016), and wall-time speedup on a single NVIDIA A100 GPU as evaluation metrics. More details of experiment settings, ablation studies and quantitative results can be found in Appendix.

	M_p	M_q	γ	α	Speedup ratio			
					bs=1	bs=8	bs=128	bs=256
381	MAR-H	MAR-B	32	0.19	1.44 ×	1.61 ×	2.17 ×	2.33 ×
382	MAR-H	MAR-B	16	0.26	1.37×	1.51×	2.07×	2.20×
383	MAR-H	MAR-B	8	0.27	1.26×	1.44×	1.88×	1.96×
384	MAR-H	MAR-B	4	0.30	1.11×	1.20×	1.56×	1.62×
385	xAR-H	xAR-B	32	0.22	1.77 ×	2.10 ×	2.52 ×	2.72 ×
386	xAR-H	xAR-B	16	0.26	1.58×	2.06×	2.31×	2.61×
387	xAR-H	xAR-B	8	0.29	1.61×	1.86×	2.07×	2.18×
388	xAR-H	xAR-B	4	0.36	1.30×	1.62×	1.92×	2.11×

Table 1: Results of speedup ratio and acceptance rate α on MAR and xAR under different draft lengths and batch sizes. The bs refers to batch size.

Resolution	M_p	M_q	γ	α	Speedup ratio			
					bs=1	bs=8	bs=16	bs=32
256	Harmon-H	Harmon-B	32	0.17	1.47 ×	1.67 ×	1.88 ×	2.05 ×
	Harmon-H	Harmon-B	16	0.21	1.41×	1.58×	1.78×	1.93×
	Harmon-H	Harmon-B	8	0.25	1.29×	1.44×	1.60×	1.72×
	Harmon-H	Harmon-B	4	0.33	1.11×	1.22×	1.33×	1.42×
512	Harmon-H	Harmon-B	32	0.15	1.63 ×	1.94 ×	2.23 ×	2.54 ×
	Harmon-H	Harmon-B	16	0.23	1.55×	1.83×	2.09×	2.35×
	Harmon-H	Harmon-B	8	0.25	1.41×	1.65×	1.85×	2.05×
	Harmon-H	Harmon-B	4	0.38	1.20×	1.37×	1.50×	1.63×

Table 2: Results of speedup ratio and acceptance rate α on Harmon under different resolutions, draft lengths, and batch sizes. The bs refers to batch size. Due to CUDA out-of-memory, this set of experiments opt for smaller batch sizes.

M_p	M_q	w/o CFG		w/ CFG	
		FID \downarrow	IS \uparrow	FID \downarrow	IS \uparrow
MAR-L		2.60	221.4	1.78	296.0
MAR-L	MAR-B	2.59 ± 0.04	218.4 ± 3.4	1.81 ± 0.05	303.7 ± 4.3
MAR-H		2.35	227.8	1.55	303.7
MAR-H	MAR-B	2.36 ± 0.05	228.5 ± 2.2	1.60 ± 0.05	301.6 ± 2.6
MAR-H	MAR-L	2.34 ± 0.04	228.9 ± 2.8	1.57 ± 0.04	301.4 ± 2.5

Table 3: Evaluation of FID and IS on unconditional and conditional generation, compared with original MAR-L and MAR-H models. Our method achieves acceleration while maintaining performance within a reasonable interval.

4.2 MAIN RESULTS

Speedup results. Table 1 and 2 show the speedup ratio and the overall acceptance rate under different batch sizes, where draft lengths range from 8 to 32. Overall, as the batch size grows, the efficacy of speculative decoding becomes more evident. More specifically, our algorithm achieves an impressive speedup of up to $2.33 \times$ on MAR, $2.72 \times$ on xAR and $2.54 \times$ on Harmon.

Quantitative results. The class-conditioned and unconditioned FID and IS metrics for our continuous speculative decoding evaluated on MAR models are shown in Table 3. We conduct multiple experiments and report the average performance and the standard deviation to ensure the reliability of our conclusions. These results demonstrate that our algorithm significantly preserves the quality of generated images, which will be further discussed in subsequent sections. Thus, our approach offers a robust solution for efficient and reliable model inference.

M_p	M_q	γ	α	$\mathbb{E} [\ x^q - x^p\ ^2]$
			w/o align w/ align	w/o align w/ align
MAR-H	MAR-B	32	0.07 0.30	2.56 1.13
MAR-H	MAR-B	16	0.07 0.33	2.36 0.91
MAR-H	MAR-B	8	0.13 0.31	2.22 0.82
MAR-H	MAR-B	4	0.14 0.32	2.17 0.80

Table 4: Ablation study on the acceptance rate α and the average distance of each draft and target token with and without denoising trajectory alignment under different draft length γ .

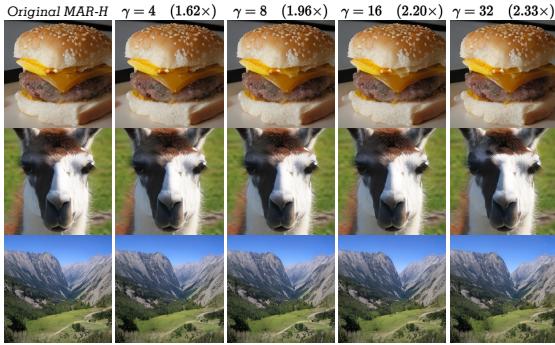


Figure 7: Qualitative Comparison Results. We show the generated images using different draft length γ .

M_p	M_q	γ	α/Speed
			0% 5% 15%
MAR-H	MAR-B	32	0.25/1.63 <u>0.30/1.63</u> 0.33/1.61
MAR-H	MAR-B	16	0.32/1.53 <u>0.33/1.52</u> 0.34/1.51
MAR-H	MAR-B	8	0.33/1.47 <u>0.31/1.47</u> 0.34/1.44
MAR-H	MAR-B	4	0.31/1.21 <u>0.32/1.21</u> 0.34/1.20

Table 5: Ablation study on pre-filling ratio. Underline indicates the highest speedup. **Bold** means the highest α .

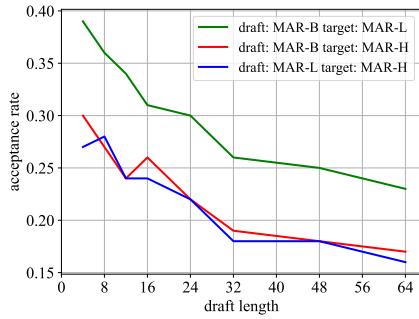


Figure 8: Plots of acceptance rate α along with draft lengths γ on MAR.

Qualitative results. Visualization results are shown in Figure 6 and 7 to showcase the image quality generated by our algorithm. Figure 6 demonstrates the results with our algorithm. Figure 7 showcases the results of the original MAR-H with autoregressive step 256 and varying draft lengths γ . In addition to a significant acceleration, our method can maintain the quality of the generated images, which is consistent with the theoretical proof.

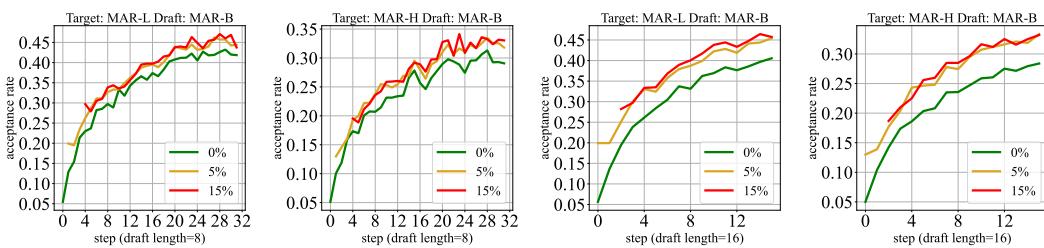


Figure 9: Plots of per-step acceptance rate α under different pre-filling ratios, along with different draft length γ , averaged on 1000 samples from MAR.

4.3 ABLATION STUDY

The α vs. γ . The relationship between acceptance rate α and draft length γ on MAR model is depicted in Figure 8. As the length of the draft increases, the acceptance rate tends to decline. This observation suggests that while longer drafts can substantially mitigate inference overhead, they are intrinsically constrained by the capabilities of the draft model itself. Consequently, an increase in the number of draft lengths is associated with greater deviations from the target model’s distribution, ultimately leading to reduced acceptance rates.

486 **Effectiveness of denoising trajectory alignment.** Table 4 shows the acceptance rate and the av-
 487 erage distance of each draft and target token with and without our aligned trajectory. The results
 488 show that, denoising trajectory alignment reduces the distance between the draft tokens and the tar-
 489 get tokens. Before alignment, the distance between them reaches > 2 , leading to quite small α .
 490 As the distance is reduced to ~ 1 , $p(Y)/q(Y)$ increases, further increasing the α to $> 30\%$. Also,
 491 denoising trajectory alignment helps to simplify the calculation of $p(Y)/q(Y)$.
 492

493 **Influence of pre-filled tokens.** The ablation study of pre-filling ratios at 0%, 5%, and 15% on
 494 MAR model is illustrated in Figure 9. Pre-filling can compensate for the low acceptance rates
 495 observed during the initial stages of autoregressive sampling and enhance the overall acceptance rate,
 496 as shown in Table 5. As the pre-filling ratio increases, the advantages conferred by this approach
 497 exhibit diminishing returns.
 498

5 CONCLUSION

501 We present continuous speculative decoding to accelerate continuous visual AR models. We propose
 502 denoising trajectory alignment based on proximity in the reparameterization theorem and token pre-
 503 filling to enhance the acceptance rate. Acceptance-rejection sampling is introduced with a proper
 504 upper bound to sample the modified distribution without analytic expression. The repetitive diffu-
 505 sion model inference is tackled by reusing denoising trajectory alignment. Extensive experiments
 506 show that our algorithm achieves over $2\times$ speedup while maintaining the output distribution. We
 507 expect our work will provide more thoughts and insights into the inference acceleration with con-
 508 tinuous autoregressive models in vision and other domains.
 509

REFERENCES

510 George Casella, Christian P Robert, and Martin T Wells. Generalized accept-reject sampling
 511 schemes. *Lecture notes-monograph series*, pp. 342–347, 2004.
 512 Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and
 513 John Jumper. Accelerating large language model decoding with speculative sampling.
 514 *arXiv:2302.01318*, 2023.
 515 Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever.
 516 Generative pretraining from pixels. In *International Conference on Machine Learning*, pp. 1691–
 517 1703, 2020.
 518 Xi Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. Pixelsnail: An improved au-
 519 toregressive generative model. In *International Conference on Machine Learning*, pp. 864–872,
 520 2018.
 521 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
 522 archical image database. In *2009 IEEE conference on Computer Vision and Pattern Recognition*,
 523 pp. 248–255. Ieee, 2009.
 524 Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image
 525 synthesis. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recog-
 526 nition*, pp. 12873–12883, 2021.
 527 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 528 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in
 529 Neural Information Processing Systems*, 30, 2017.
 530 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 531 Neural Information Processing Systems*, 33:6840–6851, 2020.
 532 Doohyuk Jang, Sihwan Park, June Yong Yang, Yeonsung Jung, Jihun Yun, Souvik Kundu, Sung-
 533 Yub Kim, and Eunho Yang. Lantern: Accelerating visual autoregressive models with relaxed
 534 speculative decoding. *arXiv:2410.03355*, 2024.

540 Siqi Kou, Lanxiang Hu, Zhezhi He, Zhijie Deng, and Hao Zhang. Cllms: Consistency large language
 541 models. *arXiv:2403.00835*, 2024.

542

543 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
 544 decoding. In *International Conference on Machine Learning*, pp. 19274–19286, 2023.

545

546 Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image
 547 generation without vector quantization. *Advances in Neural Information Processing Systems*, 37:
 548 56424–56445, 2025.

549

550 Yizhuo Li, Yuying Ge, Yixiao Ge, Ping Luo, and Ying Shan. Dicode: Diffusion-compressed deep
 551 tokens for autoregressive video generation with language models. *arXiv:2412.04446*, 2024a.

552

553 Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. Eagle: Speculative sampling requires
 554 rethinking feature uncertainty. *arXiv:2401.15077*, 2024b.

555

556 Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. Eagle-2: Faster inference of language
 557 models with dynamic draft trees. *arXiv:2406.16858*, 2024c.

558

559 Dongyang Liu, Shitian Zhao, Le Zhuo, Weifeng Lin, Yu Qiao, Hongsheng Li, and Peng Gao.
 560 Lumina-mgpt: Illuminate flexible photorealistic text-to-image generation with multimodal gener-
 561 ative pretraining. *arXiv:2408.02657*, 2024.

562

563 Xiaoxuan Liu, Lanxiang Hu, Peter Bailis, Alvin Cheung, Zhijie Deng, Ion Stoica, and Hao Zhang.
 564 Online speculative decoding. *arXiv:2310.07177*, 2023.

565

566 Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantiza-
 567 tion: Vq-vae made simple. *arXiv:2309.15505*, 2023.

568

569 Xupeng Miao, Gabriele Oliaro, Zhihao Zhang, Xinhao Cheng, Zeyu Wang, Zhengxin Zhang,
 570 Rae Ying Yee Wong, Alan Zhu, Lijie Yang, Xiaoxiang Shi, et al. Specinfer: Accelerating
 571 generative large language model serving with tree-based speculative inference and verification.
 572 *arXiv:2305.09781*, 2023.

573

574 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.
 575 In *International Conference on Machine Learning*, pp. 8162–8171, 2021.

576

577 Sihwan Park, Doohyuk Jang, Sungyub Kim, Souvik Kundu, and Eunho Yang. Lantern++: En-
 578 hanced relaxed speculative decoding with static tree drafting for visual auto-regressive models.
 579 *arXiv:2502.06352*, 2025.

580

581 Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and
 582 Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pp. 4055–
 583 4064, 2018.

584

585 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.
 586 Improved techniques for training gans. *Advances in Neural Information Processing Systems*, 29,
 587 2016.

588

589 Andrea Santilli, Silvio Severino, Emilian Postolache, Valentino Maiorca, Michele Mancusi, Ric-
 590 cardo Marin, and Emanuele Rodolà. Accelerating transformer inference for translation via paral-
 591 lel decoding. *arXiv:2305.10427*, 2023.

592

593 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models.
 594 *arXiv:2010.02502*, 2020.

595

596 Ren Sucheng, Yu Qihang, He Ju, Shen Xiaohui, Yuille Alan, and Chen Liang-Chieh. Beyond next-
 597 token: Next-x prediction for autoregressive visual generation. *arXiv preprint arXiv:2502.20388*,
 598 2025.

599

600 Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan.
 601 Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv:2406.06525*,
 602 2024.

594 Haotian Tang, Yecheng Wu, Shang Yang, Enze Xie, Junsong Chen, Junyu Chen, Zhuoyang Zhang,
 595 Han Cai, Yao Lu, and Song Han. Hart: Efficient visual generation with hybrid autoregressive
 596 transformer. *arXiv:2410.10812*, 2024.

597

598 Yao Teng, Han Shi, Xian Liu, Xuefei Ning, Guohao Dai, Yu Wang, Zhenguo Li, and Xihui Liu. Ac-
 599 celerating auto-regressive text-to-image generation with training-free speculative jacobi decoding.
 600 *arXiv:2410.01699*, 2024.

601 Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling:
 602 Scalable image generation via next-scale prediction. *Advances in Neural Information Processing
 603 Systems*, 37:84839–84865, 2025.

604 Michael Tschannen, Cian Eastwood, and Fabian Mentzer. Givt: Generative infinite-vocabulary
 605 transformers. In *European Conference on Computer Vision*, pp. 292–309. Springer, 2025.

606

607 Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Con-
 608 ditional image generation with pixelcnn decoders. *Advances in Neural Information Processing
 609 Systems*, 29, 2016.

610 Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks.
 611 In *International Conference on Machine Learning*, pp. 1747–1756, 2016.

612

613 Size Wu, Wenwei Zhang, Lumin Xu, Sheng Jin, Zhonghua Wu, Qingyi Tao, Wentao Liu, Wei Li,
 614 and Chen Change Loy. Harmonizing visual representations for unified multimodal understanding
 615 and generation. 2025. URL <https://arxiv.org/abs/2503.21979>.

616 Yilun Xu, Gabriele Corso, Tommi Jaakkola, Arash Vahdat, and Karsten Kreis. Disco-diff: Enhanc-
 617 ing continuous diffusion models with discrete latents. In *International Conference on Machine
 618 Learning*, 2024.

619

620 Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G
 621 Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video
 622 transformer. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recog-
 623 nition*, pp. 10459–10469, 2023a.

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

627

628 Weilin Zhao, Yuxiang Huang, Xu Han, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. Ouroboros:
 629 Speculative decoding with large model enhanced drafting. *arXiv:2402.13720*, 2024a.

630

631 Yao Zhao, Zhitian Xie, Chen Liang, Chenyi Zhuang, and Jinjie Gu. Lookahead: An inference accel-
 632 eration framework for large language model with lossless generation accuracy. In *Proceedings of
 633 the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6344–6355,
 634 2024b.

635

636 Yongchao Zhou, Kaifeng Lyu, Ankit Singh Rawat, Aditya Krishna Menon, Afshin Rostamizadeh,
 637 Sanjiv Kumar, Jean-François Kagy, and Rishabh Agarwal. Distillspec: Improving speculative
 638 decoding via knowledge distillation. *arXiv:2310.08461*, 2023.

639

640

641

642

643

644

645

646

647

648 APPENDIX
649650 A DETAILED PROOFS
651652 We will provide a more detailed process and proof of continuous speculative decoding.
653654 The output of the whole model is composed of the condition from the autoregressive model and the
655 denoising process via the diffusion model. So it can be written as:

656
$$657 p(x_{N,0:T}|x_1, \dots, x_{N-1}) = p(x_{N,T}) \prod_{t=1}^T p(x_{N,t-1}|x_{N,t}, x_1, \dots, x_{N-1}),$$

658

659 where N denotes the number of autoregressive step, x_1, \dots, x_{N-1} are tokens generated by previous
660 steps, and $t \in [0, T]$ is diffusion timestep. For simplicity, we omit N and $x_{1:N-1}$ and let $Y = x_{0:T}$
661 to obtain $p(Y) = p(x_T) \prod_{t=1}^T p(x_{t-1}|x_t)$.
662663 A.1 DISCUSSION OF THE APPROXIMATED RATIO
664

665 A.1.1 THE APPROXIMATION OF DIFFUSION DISTRIBUTION

666 For a token x , the acceptance criterion is defined as the ratio of its probability under target dis-
667 tribution to the one under draft distribution, that is, $p(x)/q(x)$. In discrete form, the probability
668 can be directly obtained. But in continuous form, the probability is usually obtained via diffusion
669 process (Ho et al., 2020; Song et al., 2020). Specifically, $x = x_0$ is sampled via reverse diffusion
670 (denoising) process $p(x_T) \prod_{t=1}^T p(x_{t-1}|x_t)$. So the probability of the entire denoising process is:
671

672
$$673 p(x_0) = \int_{x_{1:T}} p(x_T) \prod_{t=1}^T p(x_{t-1}|x_T) dx_{1:T}, \quad (14)$$

674

675 However, Equation 14 is analytically intractable because the product of these complex terms rarely
676 yields a simple, closed-form analytical solution to the high-dimensional integral.
677678 As an alternative to modeling the marginal distribution of a token x_0 , we model the joint probability
679 of a single fixed path $p(Y)$, where $Y = [x_0, x_1, \dots, x_T]$, to characterize the probability of each
680 denoising trajectory. The objective of $p(x)$ is to maintain the marginal distribution, while $p(Y)$ aims
681 to maintain the probability of the denoising trajectory itself. Both methods are effective, yet they
682 address distribution maintenance from different perspectives.
683684 However, using $p(Y)$ is impractical for speculative decoding because the acceptance rate is low. To
685 validate this point, we empirically recorded the average value of different kinds of likelihood ratios
686 over 10,000 samples with a draft length of 4, including: (i) the single-path ratio $p(Y)/q(Y)$, (ii)
687 the two-path ratio $p(Y_p)/q(Y_q)$ without denoising trajectory alignment, and (iii) the two-path ratio
688 $p(Y_p)/q(Y_q)$ with denoising trajectory alignment, as shown in Table 6.
689690

Likelihood ratio	Value	Acceptance rate
$p(Y)/q(Y)$	5.33×10^{-23}	0.0%
$p(Y_p)/q(Y_q)$, w/o align	0.067	14%
$p(Y_p)/q(Y_q)$, w/ align	1.86	32%

691 Table 6: Likelihood ratio comparison of different calculation approaches.
692693 As shown in the first row, the path-space likelihood ratio is extremely small, leading to a 0% accep-
694 tance rate. This is because the draft model’s trajectory Y_p inherently diverges from the target model’s
695 expected trajectory. In each denoising step, samples drawn from the draft model’s distribution q are
696 unlikely to fall near μ of the target distribution p , which results in a low single-step ratio p/q . As
697 the multi-step denoising process proceeds, the overall $p(Y)/q(Y)$ becomes extremely small.
698699 The second row compares ratios derived from independent trajectories (Y_p and Y_q), while the final x_0
700 is generated by the draft model. This x_0 , without alignment, is highly unlikely to fall near the target
701

702 model's target distribution. However, the probability values for the other steps in these trajectories
 703 are derived from the model's own path, and thus maintain a relatively reasonable value.
 704

705 The third row incorporates denoising trajectory alignment. This improvement is twofold: first, our
 706 manuscript demonstrates that the expected distance decreases; and second, our analysis shows that
 707 the correlation of p and q becomes 1 (common issue 2). Consequently, the samples generated by q
 708 have a high probability under p , resulting in an increased p/q .
 709

710 For the above reasons, our work adopts a practical approximation for $p(x_0|x_T)/q(x_0|x_T)$: the ratio
 711 of the joint probabilities $p(Y_p)/q(Y_q)$, where both Y_p and Y_q share the same x_0 . This ratio serves as
 712 a surrogate for the intractable marginal ratio. By utilizing denoising trajectory alignment, we ensure
 713 that Y_p and Y_q are tightly coupled (as discussed in common issue 2), making the likelihood ratio a
 714 valid surrogate and ensuring its numerical stability in practical applications.
 715

714 A.1.2 HOW ACCURATELY $p(Y_p)/q(Y_q)$ APPROXIMATES $p(Y)/q(Y)$

716 We define the joint probability of a single denoising trajectory as $p(Y)$, where Y is the sequence
 717 $[x_0, x_1, \dots, x_T]$. The path space ratio $R = p(Y)/q(Y)$, where $Y = Y_q$, should be an unbiased
 718 estimator of $p(Y)$ to ensure that the sampling method preserves the target distribution, that is:
 719

$$720 \mathbb{E}_{Y_q \sim q}[R] = \int q(Y_q) \frac{p(Y_q)}{q(Y_q)} dY_q = \int p(Y_q) dY_q = 1.$$

721 However, it is very inefficient. As shown in common issue 1, $p(Y)/q(Y)$ is extremely small, because
 722 the trajectory sampled from the draft model may fall into the region where the target model assigns
 723 negligible probability mass.
 724

725 To improve the efficiency, we propose a new estimator, $\tilde{R} := p(Y_p)/q(Y_q)$. \tilde{R} is a biased estimator;
 726 its expectation under q is not equal to 1. We analyze the difference between the expectation of \tilde{R}
 727 and 1, i.e., its **bias**:
 728

$$729 B = \mathbb{E}_q[\tilde{R}] - 1 = \mathbb{E}_q \left[\frac{p(Y_p)}{q(Y_q)} \right] - \mathbb{E}_q \left[\frac{p(Y_q)}{q(Y_q)} \right] = \mathbb{E}_q \left[\frac{p(Y_p) - p(Y_q)}{q(Y_q)} \right].$$

730 We perform a first-order Taylor expansion of $p(Y_p)$ at Y_q . Equivalently:
 731

$$732 p(Y_p) = p(Y_q + (Y_p - Y_q)) \approx p(Y_q) + \nabla p(Y_q)^T (Y_p - Y_q).$$

733 Substitute into B :

$$734 B \approx \mathbb{E}_q \left[\frac{p(Y_q) + \nabla p(Y_q)^T (Y_p - Y_q) - p(Y_q)}{q(Y_q)} \right] = \mathbb{E}_q \left[\frac{\nabla p(Y_q)^T (Y_p - Y_q)}{q(Y_q)} \right].$$

735 The magnitude of B is:
 736

$$737 |B| \approx \left| \mathbb{E}_q \left[\frac{\nabla p(Y_q)^T (Y_p - Y_q)}{q(Y_q)} \right] \right| \leq \mathbb{E}_q \left[\frac{\|\nabla p(Y_q)^T\|}{|q(Y_q)|} \|Y_p - Y_q\| \right]. \quad (15)$$

738 Therefore, the bound of the bias B is proportional to $\mathbb{E}[\|Y_p - Y_q\|]$, yielding an explicit mean-square
 739 error (MSE) bound.
 740

741 The reduction of expected distance is discussed in Appendix A.2. The expected distance without
 742 denoising trajectory alignment is:
 743

$$744 \mathbb{E} [\|x_{t-1}^q - x_{t-1}^p\|^2] = \|\mu_t^q - \mu_t^p\|^2 + \text{tr}[\Sigma_t^p + \Sigma_t^q]. \quad (16)$$

745 The expected distance with denoising trajectory alignment is:
 746

$$747 \mathbb{E}_{align} [\|x_{t-1}^q - x_{t-1}^p\|^2] = \|\mu_t^q - \mu_t^p\|^2 + \text{tr}[\Sigma_t^p + \Sigma_t^q - 2\sqrt{\Sigma_t^p \Sigma_t^q}]. \quad (17)$$

756 The two distances satisfies:

$$758 \mathbb{E} [||x_{t-1}^q - x_{t-1}^p||^2] \geq \mathbb{E}_{align} [||x_{t-1}^q - x_{t-1}^p||^2] \quad (18)$$

760 The distance over the entire trajectory satisfies:

$$762 \mathbb{E} [||Y_p - Y_q||] \leq \sqrt{\sum_{t=1}^T \mathbb{E} ||x_t^q - x_t^p||^2}. \quad (19)$$

765 Substituting into B yields the final first-order error bound:

$$767 |B| \leq \mathbb{E}_q \left[\frac{||\nabla p(Y_q)^T||}{q(Y_q)} ||Y_p - Y_q|| \right] \leq \mathbb{E}_q \left[\frac{||\nabla p(Y_q)^T||}{q(Y_q)} \sqrt{\sum_{t=1}^T ||\mu_t^q - \mu_t^p||^2 + \text{tr}[\Sigma_t^p + \Sigma_t^q]} \right] \quad (20)$$

$$772 |B_{align}| \leq \mathbb{E}_q \left[\frac{||\nabla p(Y_q)^T||}{q(Y_q)} \sqrt{\sum_{t=1}^T ||\mu_t^q - \mu_t^p||^2 + \text{tr}[\Sigma_t^p + \Sigma_t^q] - 2\sqrt{\Sigma_t^p \Sigma_t^q}} \right] \quad (21)$$

775 The error bound shows that the squared drift difference $||\mu_t^q - \mu_t^p||^2$ and the covariance term $\text{tr}(\Sigma_t^p + \Sigma_t^q)$ or $\text{tr}(\Sigma_t^p + \Sigma_t^q - 2\sqrt{\Sigma_t^p \Sigma_t^q})$ together determine the expected bias incurred by approximating $p(Y)/q(Y)$.

779 Since the use of denoising trajectory alignment produces the cross term $-2\sqrt{\Sigma_t^p \Sigma_t^q}$, the bias $|B_{align}|$ is typically smaller than $|B|$, thereby supporting a more accurate approximation of $p(Y)/q(Y)$.

782 A.1.3 THE EXTENT TO WHICH $p(Y_p)/q(Y_q)$ CAN IMPROVE THE RATIO

783 Let:

$$785 \log R = \log p(Y_q) - \log q(Y_q), \quad (22)$$

$$786 \log \tilde{R} = \log p(Y_p) - \log q(Y_q). \quad (23)$$

788 The expected difference between the two quantities can be expressed as:

$$790 \Delta l := \mathbb{E}[\log \tilde{R}] - \mathbb{E}[\log R] = \mathbb{E}[\log p(Y_p)] - \mathbb{E}[\log p(Y_q)].$$

792 For each term, we have:

$$793 \mathbb{E}[\log p(Y_p)] = \mathbb{E}_{Y \sim p}[\log p(Y)] = -H(p), \quad (24)$$

$$795 \mathbb{E}[\log p(Y_q)] = -H(q) - D_{KL}(q||p). \quad (25)$$

796 Substituting into Δl yields:

$$798 \Delta l = [-H(p)] - [-H(q) - D_{KL}(q||p)] \\ 799 = D_{KL}(q||p) + [H(q) - H(p)]. \quad (26)$$

801 Empirically, the larger target model is expected to be more capable and to produce more confident
802 (lower-entropy) predictive distributions than the smaller draft model. Therefore, we assume the draft
803 model's entropy $H(q)$ is larger than the target model's entropy $H(p)$. We have:

$$805 \Delta l = D_{KL}(q||p) + [H(q) - H(p)] \geq 0. \quad (27)$$

806 It indicates that in the log domain, \tilde{R} exceeds R by Δl . This implies that our method increases the
807 expected log-ratio, $\mathbb{E}[\log \tilde{R}] - \mathbb{E}[\log R] = \Delta l$, which subsequently leads to a higher expectation
808 of the ratio, $\mathbb{E}[\tilde{R}]$. This explains the empirically observed higher likelihood ratios and increased
809 acceptance rates.

810 A.2 DENOISING TRAJECTORY ALIGNMENT
811812 We obtain $x = x_0$ through the denoising process:

813
814
$$p(x_{0:T}) = p(x_T) \prod_{t=1}^T p(x_{t-1}|x_t), \quad (28)$$

815

816 with the conditioned probability distributions as Gaussian approximated by a neural network θ :

817
818
$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)). \quad (29)$$

819 Therefore, $p_\theta(x_{t-1}|x_t)$ can be computed using the PDF of the Gaussian distribution. The computation
820 and corresponding notation of $q_\theta(x_{t-1}|x_t)$ are the same.821 Empirically, x_{t-1} is obtained by sampling the Gaussian distribution on the right-hand side by **repara-**
822 **parameterization**. That is, we first sample $\varepsilon_t \sim \mathcal{N}(0, I)$, and then we obtain the result by scale and
823 shift $x_{t-1} = \sqrt{\Sigma_\theta(x_t, t)} \cdot \varepsilon_t + \mu_\theta(x_t, t)$. To this end, we can compute $p(x)$ and $q(x)$ to obtain the
824 ratio $p(x)/q(x)$.825 However, as described in Sec. 3, directly computing the $p(x)$ and $q(x)$ is algebraically correct but
826 may lead to a low acceptance rate due to a distinct denoising trajectory. Thus, we employ the same
827 ε_t in $p(x)$ and $q(x)$ to align their trajectory as closely as possible without affecting the denoising
828 procedure and the results.830 **Proof of Theorem 1** Denoising trajectory alignment enhances consistency by reducing the ex-
831 pected inter-sample distance throughout the denoising process. Suppose $x_{t-1}^p = \sqrt{\Sigma_t^p} \cdot \varepsilon_t^p + \mu_t^p$ and
832 $x_{t-1}^q = \sqrt{\Sigma_t^q} \cdot \varepsilon_t^q + \mu_t^q$.833 **Without alignment** ($\varepsilon_t^p \neq \varepsilon_t^q$), let:

834
835
$$\begin{aligned} X &= x_{t-1}^q - x_{t-1}^p \\ &= (\mu_t^q - \mu_t^p) + \left(\sqrt{\Sigma_t^p} \cdot \varepsilon_t^p - \sqrt{\Sigma_t^q} \cdot \varepsilon_t^q \right) \\ &= \mu + Y, \end{aligned} \quad (30)$$

836

837 where $\mu = \mu_t^q - \mu_t^p$ and $Y = \sqrt{\Sigma_t^p} \cdot \varepsilon_t^p - \sqrt{\Sigma_t^q} \cdot \varepsilon_t^q$. The $\|X\|^2$ is given by:

838
839
$$\|X\|^2 = X^T X = (\mu + Y)^T (\mu + Y) = \mu^T \mu + \mu^T Y + Y^T \mu + Y^T Y. \quad (31)$$

840

841 Therefore:

842
843
$$\mathbb{E} [\|X\|^2] = \mathbb{E} [\mu^T \mu] + \mathbb{E} [\mu^T Y] + \mathbb{E} [Y^T \mu] + \mathbb{E} [Y^T Y]. \quad (32)$$

844

845 First of all, $\mathbb{E} [\mu^T \mu] = \mu^T \mu$.

846 Since:

847
848
$$\mathbb{E} [Y] = \mathbb{E} \left[\sqrt{\Sigma_t^p} \cdot \varepsilon_t^p - \sqrt{\Sigma_t^q} \cdot \varepsilon_t^q \right] = 0 \quad (\text{for } \mathbb{E} [\varepsilon_t^p] = \mathbb{E} [\varepsilon_t^q] = 0), \quad (33)$$

849

850 we have $\mathbb{E} [\mu^T Y] = \mathbb{E} [Y^T \mu] = 0$.851 For $\mathbb{E} [Y^T Y]$:

852
853
$$\begin{aligned} Y^T Y &= (\sqrt{\Sigma_t^p} \cdot \varepsilon_t^p - \sqrt{\Sigma_t^q} \cdot \varepsilon_t^q)^T (\sqrt{\Sigma_t^p} \cdot \varepsilon_t^p - \sqrt{\Sigma_t^q} \cdot \varepsilon_t^q) \\ &= (\varepsilon_t^p)^T (\sqrt{\Sigma_t^p})^T \sqrt{\Sigma_t^p} \varepsilon_t^p + (\varepsilon_t^q)^T (\sqrt{\Sigma_t^q})^T \sqrt{\Sigma_t^q} \varepsilon_t^q \\ &\quad - (\varepsilon_t^p)^T (\sqrt{\Sigma_t^p})^T \sqrt{\Sigma_t^q} \varepsilon_t^q - (\varepsilon_t^q)^T (\sqrt{\Sigma_t^q})^T \sqrt{\Sigma_t^p} \varepsilon_t^p. \end{aligned} \quad (34)$$

854

855 Note that $(\sqrt{\Sigma_t^p})^T \sqrt{\Sigma_t^p} = \Sigma_t^p$ and $(\sqrt{\Sigma_t^q})^T \sqrt{\Sigma_t^q} = \Sigma_t^q$, we have:

856
857
$$Y^T Y = (\varepsilon_t^p)^T \Sigma_t^p \varepsilon_t^p + (\varepsilon_t^q)^T \Sigma_t^q \varepsilon_t^q - (\varepsilon_t^p)^T C \varepsilon_t^q - (\varepsilon_t^q)^T C^T \varepsilon_t^p \quad (C = \sqrt{\Sigma_t^p} \sqrt{\Sigma_t^q}). \quad (35)$$

858

864 Since $\mathbb{E} [\varepsilon^T M \varepsilon] = \text{tr}[M]$ where $\varepsilon \sim \mathcal{N}(0, I_n)$ and M is a matrix, and ε_t^p and ε_t^q are independent
 865 (expectation of the cross term is 0), we have:

$$866 \mathbb{E} [Y^T Y] = \text{tr}[\Sigma_t^p] + \text{tr}[\Sigma_t^q] + 0 + 0 = \text{tr}[\Sigma_t^p + \Sigma_t^q]. \quad (36)$$

868 Put these results to Equation 32 makes:

$$869 \mathbb{E}_{\varepsilon_t^p \neq \varepsilon_t^q} \left[\|x_{t-1}^q - x_{t-1}^p\|^2 \right] = \|\mu_t^q - \mu_t^p\|^2 + \text{tr}[\Sigma_t^q + \Sigma_t^p]. \quad (37)$$

872 **With alignment** ($\varepsilon_t = \varepsilon_t^p = \varepsilon_t^q$), we have:

$$873 \begin{aligned} X &= (\mu_t^q - \mu_t^p) + \left(\sqrt{\Sigma_t^p} - \sqrt{\Sigma_t^q} \right) \varepsilon_t \\ 874 &= \mu + Y \varepsilon_t, \end{aligned} \quad (38)$$

877 where $\mu = \mu_t^q - \mu_t^p$ and $Y = \sqrt{\Sigma_t^p} - \sqrt{\Sigma_t^q}$. We have:

$$878 \begin{aligned} 879 \|X\|^2 &= (\mu + Y \varepsilon_t)^T (\mu + Y \varepsilon_t) \\ 880 &= \mu^T \mu + \mu^T Y \varepsilon_t + \varepsilon_t^T Y^T \mu + \varepsilon_t^T Y^T Y \varepsilon_t. \end{aligned} \quad (39)$$

881 Since $\mu^T Y \varepsilon_t$ and $\varepsilon_t^T Y^T \mu$ are scalars, we have:

$$883 \mu^T Y \varepsilon_t + \varepsilon_t^T Y^T \mu = 2\mu^T Y \varepsilon_t. \quad (40)$$

884 So:

$$886 \|X\|^2 = \|\mu\|^2 + 2\mu^T Y \varepsilon_t + \varepsilon_t^T Y^T Y \varepsilon_t. \quad (41)$$

888 Therefore:

$$889 \mathbb{E} [\|X\|^2] = \mathbb{E} [\|\mu\|^2] + 2\mathbb{E} [\mu^T Y \varepsilon_t] + \mathbb{E} [\varepsilon_t^T Y^T Y \varepsilon_t]. \quad (42)$$

891 Note that $\mathbb{E} [\|\mu\|^2] = \|\mu\|^2$, and $2\mathbb{E} [\mu^T Y \varepsilon_t] = 2\mu^T Y \mathbb{E} [\varepsilon_t] = 0$.

892 Also, $\mathbb{E} [\varepsilon_t^T Y^T Y \varepsilon_t] = \text{tr}[Y^T Y]$, we have:

$$894 \begin{aligned} \text{tr}[Y^T Y] &= \text{tr}[(\sqrt{\Sigma_t^p} - \sqrt{\Sigma_t^q})^T (\sqrt{\Sigma_t^p} - \sqrt{\Sigma_t^q})] \\ 895 &= \text{tr}[\Sigma_t^p] + \text{tr}[\Sigma_t^q] - 2\text{tr}[\sqrt{\Sigma_t^p \Sigma_t^q}] \\ 896 &= \text{tr}[\Sigma_t^p + \Sigma_t^q] - 2\text{tr}[\sqrt{\Sigma_t^p \Sigma_t^q}]. \end{aligned} \quad (43)$$

900 Put these results in Equation 39 makes:

$$901 \mathbb{E}_{\varepsilon_t^p = \varepsilon_t^q} \left[\|x_{t-1}^q - x_{t-1}^p\|^2 \right] = \|\mu_t^q - \mu_t^p\|^2 + \text{tr}[\Sigma_t^p + \Sigma_t^q] - 2\text{tr}[\sqrt{\Sigma_t^p \Sigma_t^q}]. \quad (44)$$

904 Subtracting Equation 44 from Equation 37 yields:

$$905 \begin{aligned} \Delta \mathbb{E} &= \mathbb{E}_{\varepsilon_t^p \neq \varepsilon_t^q} \left[\|x_{t-1}^q - x_{t-1}^p\|^2 \right] - \mathbb{E}_{\varepsilon_t^p = \varepsilon_t^q} \left[\|x_{t-1}^q - x_{t-1}^p\|^2 \right] \\ 906 &= \text{tr}[\Sigma_t^q + \Sigma_t^p] - \text{tr}[\Sigma_t^p + \Sigma_t^q] + 2\text{tr}[\sqrt{\Sigma_t^p \Sigma_t^q}] \\ 907 &= 2 \cdot \text{tr} \left[\sqrt{\Sigma_t^q \Sigma_t^p} \right] \geq 0. \end{aligned} \quad (45)$$

911 Furthermore, alignment also simplifies the computation of $\frac{p(x)}{q(x)}$. Note that in Gaussian distribution,
 912 we have:

$$914 \begin{aligned} p(x) &= \frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma|}} \exp \left\{ \frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu) \right\} \\ 915 &= \frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma|}} \exp \{ \varepsilon_t^T \varepsilon_t \}. \end{aligned} \quad (46)$$

918 Since we have the same ϵ_t of both $p(x)$ and $q(x)$, the exponential term can be eliminated to obtain:
919

$$\begin{aligned} 920 \frac{p(x)}{q(x)} &= \frac{\frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma_p|}} \exp\left\{\frac{1}{2}(x - \mu_p)^T \Sigma_p^{-1} (x - \mu_p)\right\}}{921 \frac{1}{(\sqrt{2\pi})^n \sqrt{|\Sigma_q|}} \exp\left\{\frac{1}{2}(x - \mu_q)^T \Sigma_q^{-1} (x - \mu_q)\right\}} \\ 922 &= \frac{\sqrt{|\Sigma_q|}}{\sqrt{|\Sigma_p|}}. \end{aligned} \quad (47)$$

926 Therefore, for all timesteps t :
927

$$\begin{aligned} 928 \frac{p(x)}{q(x)} &= \frac{p(x_T) \prod_{t=2}^T p(x_{t-1}|x_t)}{q(x_T) \prod_{t=2}^T q(x_{t-1}|x_t)} \cdot \frac{p(x_0|x_1)}{q(x_0|x_1)} \\ 929 &= \frac{\prod_{t=2}^T \sqrt{|\Sigma_{q,t}|}}{\prod_{t=2}^T \sqrt{|\Sigma_{p,t}|}} \cdot \frac{p(x_0|x_1)}{q(x_0|x_1)} \\ 930 &= \Sigma \cdot \frac{p(x_0|x_1)}{q(x_0|x_1)}, \end{aligned} \quad (48)$$

936 where Σ is the cumulative product of $\frac{\sqrt{|\Sigma_{q,t}|}}{\sqrt{|\Sigma_{p,t}|}}$ along the denoising intermediate steps. Also, since
937 $x_0 \sim q(x)$ is verified by the target model, $p(x_0|x_1)$ is not obtained by denoising. It is obtained by
938 substituting x_0 into $p(x_0|x_1)$. We keep the two terms since they should be computed separately.
939

941 A.3 ACCEPTANCE-REJECTION SAMPLING

943 After rejection, we should resample a new output from:

$$945 p'(Y) = \frac{\max(0, p(Y) - q(Y))}{\int_{x'} \max(0, p(x') - q(x')) dx'}. \quad (49)$$

947 But Z is hard to obtain. This integral $Z = \int_{x'} \max(0, p(x') - q(x')) dx'$ is difficult to compute and
948 may introduce computation errors if we employ an approximation. This integral also does not have
949 an analytical form.
950

951 On the other hand, sampling from proposal distribution $p(Y)$ requires the diffusion loss module to
952 forward for another time, since the entire distribution is determined by all the denoising steps. But
953 in practice, extra model inference introduces heavy overhead and extra latency, and may reduce the
954 speed of speculative decoding, which is harmful for the whole algorithm.
955

956 **Proof of Corollary 1** The introduction of acceptance-rejection sampling can eliminate Z by $M = 1/Z$. The denoising trajectory alignment can simplify the expression and avoid repetitive diffusion
957 model inference. The result is given by:
958

$$\begin{aligned} 959 \alpha_s &= \frac{\max(0, p(Y) - q(Y))/Z}{p(Y)/Z} \\ 960 &= \frac{\max(0, p(Y) - q(Y))}{p(Y)} \\ 961 &= \frac{\max(0, p(x_T)p_\theta(x_0|x_1^p) \prod_{t=2}^T p_\theta(x_{t-1}^p|x_t^p) - q(x_T)q_\theta(x_0|x_1^q) \prod_{t=2}^T q_\theta(x_{t-1}^q|x_t^q))}{p(x_T)p_\theta(x_0|x_1^p) \prod_{t=2}^T p_\theta(x_{t-1}^p|x_t^p)} \\ 962 &= \frac{\max(0, \Sigma \cdot p_\theta(x_0|x_1^p) - q_\theta(x_0|x_1^q))}{\Sigma \cdot p_\theta(x_0|x_1^p)} \end{aligned} \quad (50)$$

969 Afterward, we can obtain the computable results by eliminating the intermediate denoising term
970 denoted as Σ . The final expression can be derived easily. The modified distribution can be sampled
971 using this approach.
972

B ALGORITHM

Algorithm 1 shows this procedure of continuous speculative decoding algorithm with our implementation of denoising trajectory alignment and acceptance-rejection sampling.

Algorithm 1 ContinuousSpeculativeDecodingStep

Inputs: $M_p, M_q, prefix$.

▷ Sample γ guesses $x_{1,\dots,\gamma}$ from M_q autoregressively

for $i = 1$ to γ **do**

- $q_i(Y_q) \leftarrow M_q(prefix + [x_1, \dots, x_{i-1}])$
- $x_i \sim q_i(Y_q)$

end for

▷ Run M_p in parallel, keep the ϵ_t the same in M_q

$p_1(Y_p), \dots, p_{\gamma+1}(Y_p) \leftarrow$

$$M_p(prefix), \dots, M_p(prefix + [x_1, \dots, x_\gamma])$$

$$\Sigma \leftarrow \frac{\prod_{t=2}^T \sqrt{|\Sigma_{q,t}|}}{\prod_{t=2}^T \sqrt{|\Sigma_{p,t}|}}$$

▷ Determine the number of accepted guesses n .

$r_1 \sim U(0, 1), \dots, r_\gamma \sim U(0, 1)$

$$\frac{p_i(Y_p)}{q_i(Y_q)} \leftarrow \Sigma \cdot \frac{p_i(x|x_1^p)}{q_i(x|x_1^q)}$$

$n \leftarrow \min(\{i - 1 \mid 1 \leq i \leq \gamma, r_i > \frac{p_i(x)}{q_i(x)}\} \cup \{\gamma\})$

▷ Sample the modified distribution via

▷ acceptance-rejection sampling.

if $n < \gamma$ **then**

repeat

- $x_t \leftarrow p_n(x|x_1^p)$
- $\alpha_s \leftarrow \frac{\max(0, \Sigma \cdot p_n(x_t|x_1^p) - q_n(x_t|x_1^q))}{\Sigma \cdot p_n(x_t|x_1^p)}$
- $\epsilon \sim U(0, 1)$

until $\epsilon \leq \alpha_s$

end if

▷ Return one token from M_p , and n tokens from M_q

return $prefix + [x_1, \dots, x_n, x_t]$

C LIMITATIONS

C.1 WALL-TIME IMPROVEMENT

As described in Leviathan et al. (2023), the expected walltime improvement is assumed to be:

$$\frac{1 - \alpha^{\gamma+1}}{(1 - \alpha)(\gamma c + 1)}, \quad (51)$$

where α is the acceptance rate of draft tokens, γ is the draft length, and c is the inference time ratio between the draft and target models. However, the existing draft model and the target model do not differ significantly in scale. For example, the inference time ratio c of MAR-B (208M) over MAR-H (943M) is 0.38 (bs=128), which is **far more larger** than the number 0.05 or close to 0 mentioned in [Leviathan et al. \(2023\)](#). Increasing the batch size would reduce c , which is why our method shows better results on large batch size.

We anticipate that our algorithm will achieve more significant runtime improvements with larger target models, like 7B, 13B, as well as smaller draft models, like 97M, 125M. This direction warrants further investigation in future research.

1026 **D IMPLEMENTATION DETAILS**
1027

1028 We have conducted extensive experiments with open-sourced continuous visual autoregressive
 1029 model MAR (Li et al., 2025) and xAR (Sucheng et al., 2025) on ImageNet (Deng et al., 2009)
 1030 256×256 generation, and unified model Harmon (Wu et al., 2025) on text-to-image generation.
 1031 The draft model is chosen from MAR-B (208M), xAR-B (172M) and Harmon-0.5B. The target
 1032 model is chosen from MAR-H (943M), xAR-H (1.1B) and Harmon-1.5B, respectively. We use of-
 1033 ficial pretrained checkpoints for all models. Since original xAR model is set to predict next cell of
 1034 the image, the whole image would be generated in 4 steps, we let xAR to autoregressively predict
 1035 next token at each position, as described as $k = 1$ setting (Sucheng et al., 2025). So the autore-
 1036 gressive step of all the involved draft model is set to 1. However, default MAR models have shown
 1037 significant results for bidirectional attention in MAR. When target models verify the draft tokens,
 1038 each output token can be regarded as the last since they can see every previous token. For MAR and
 1039 xAR, their draft and target models utilize their respective class tokens $[\text{cls}]$, which are not shared
 1040 during the speculative decoding process. Their diffusion loss is not shared either. The batch size
 1041 ranges in $\{1, 8, 128, 256\}$. The FID and IS are computed on 50k generated images, averaged on
 1042 ten runs of evaluations. For Harmon models, batch size ranges in $\{1, 8, 16, 32\}$, since larger batch
 1043 size leads to cuda-out-of-memory. The generation resolution includes both 256 and 512. The draft
 1044 and target model use their own text embedding respectively. The generation speed is measured on a
 1045 single NVIDIA A100 GPU.

1046 **E BORDER IMPACTS**
10471048 **E.1 BORDER FORMS OF OUTPUT DISTRIBUTION**
1049

1050 Various forms of continuous output spaces exist in the visual AR model. For instance,
 1051 GIVT (Tschannen et al., 2025) and DiCoDe (Li et al., 2024a) employ Gaussian Mixture Models
 1052 (GMMs) as output distribution. The PDF of token x in GMMs is expressed as:

$$1054 \quad p(x|\theta) = \sum_{k=1}^K \pi_k \mathcal{N}(x|\mu_k, \Sigma_k), \quad (52)$$

1055 where θ represents the model parameters, π_k denotes the weights of each Gaussian distribution
 1056 indexed by k , and μ_k and Σ_k indicate the mean and covariance of distribution k , respectively. Our
 1057 method is compatible with GMMs by computing $p(x)$, $q(x)$, and $p'(x)$. The modified distribution
 1058 $p'(x)$ can still be computed through acceptance-rejection sampling. In practice, considering that
 1059 current GMM methods cannot achieve competitive performance to MAR and the lack of open-
 1060 source weights with different model sizes, we haven't conducted related experiments.

1061 This applicability highlights a critical insight: our method relies on an explicit expression for the
 1062 output distribution to compute $p(x)$, $q(x)$, and subsequently $p'(x)$. As long as the specific functional
 1063 form of the distribution can be obtained, our algorithm remains universally applicable, regardless of
 1064 its particular form.

1065 **E.2 MORE VARIANTS OF DIFFUSION MODEL**
1066

1067 Other variants of diffusion samplers are still applicable to our method. In this context, we utilize
 1068 DDIM (Song et al., 2020) as an example since it can be derived from DDPM without additional
 1069 training.

1070 Theoretically, DDIM (Song et al., 2020) models a conditioned Gaussian distribution during the
 1071 reverse process, expressed as $p_\theta^{(t)}(x_{t-1}|x_t) = q_\sigma(x_{t-1}|x_t, x_0)$. Its forward (diffusion) process
 1072 is characterized as a non-Markov process. The term q_σ relies on both x_{t-1} and x_0 , where x_0 is
 1073 predicted by the model f_θ . However, this dependency does not alter the form of the output PDF. The
 1074 PDF of DDIM given in (Song et al., 2020) remains consistent with that of DDPM and is computed
 1075 using the expression $p_\theta(x_T) \prod_{t=1}^T p_\theta^{(t)}(x_{t-1}|x_t)$. Notably, our algorithm requires an explicit output
 1076 PDF and does not rely on other properties. In general, the denoising process proceeds through
 1077 sequential sampling from the previous sample. Let the PDF of initial noise be $p(x_T)$, and let the
 1078

1080 PDF for each sampling step be $p(x_{t-1}|x_t, c_{\text{other}})$, where c_{other} represents any auxiliary conditions.
 1081 The final output PDF is then determined as the cumulative product of all intermediate PDFs, namely:
 1082

$$1083 \quad p(x_{0:T}|c_{\text{other}}) = p(x_T) \prod_{t=1}^T p(x_{t-1}|x_t, c_{\text{other}}). \quad (53)$$

1086 This expression is valid across all diffusion models, regardless of specific samplers or diffusion
 1087 processes like DDIM, Rectified Flow, etc.

1088 Practically, the original implementation of MAR is realized through DDPM. Converting DDPM to
 1089 DDIM does not require additional training procedure (Song et al., 2020). We adopt the DDPM to
 1090 the DDIM in the diffusion loss of MAR and present the relevant results in Table 7. We use MAR-B
 1091 as the draft model and MAR-H as the target model. The batch size is 256, and the denoising step
 1092 of DDIM is set to 100. Our method and formulas remain applicable to other variants of diffusion
 1093 models and continue to demonstrate good performance.

1095 draft length γ	4	8	16	32
1096 speed up ratio	1.58	1.92	2.15	2.30
1097 acceptance rate α	0.26	0.24	0.21	0.18

1098 Table 7: Speed up ratio and acceptance rate on MAR with DDIM as diffusion loss under different
 1099 number of draft lengths.
 1100

1102 E.3 BORDER IMPLEMENTATION DOMAINS

1104 Continuous speculative decoding is not confined to autoregressive image generation tasks. It can
 1105 also be applied to other domains, such as autoregressive audio or video generation in continuous
 1106 spaces and broader tasks and scenarios within the continuous domain. Due to the limitations of
 1107 currently available models, we have not been able to identify usable models beyond the domain of
 1108 autoregressive image generation. However, as an algorithm without additional training or perfor-
 1109 mance loss, speculative decoding remains one of the most optimal acceleration methods. We hope
 1110 that continuous speculative decoding can provide more valuable insights and ideas for researchers
 1111 in various fields and stimulate further research on speculative decoding across diverse domains.

1113 F INFORMATION ABOUT USE OF AI ASSISTANTS

1115 In the preparation of this work, we employ AI assistants to assist with refining academic language.
 1116 The AI tools were used solely for improving clarity, grammatical correctness, and syntactic effi-
 1117 ciency—tasks analogous to those performed by a human editor or linter. All conceptual contribu-
 1118 tions, technical claims, and critical analysis remain the authors’ own.

1120 G ADDITIONAL EXPERIMENTS

1123 G.1 EFFECTIVENESS OF DRAFT & VERIFICATION.

1124 Figure 10a demonstrates the comparative generation results using a pure draft model versus the draft
 1125 & verification paradigm. During the draft & verification process, those suboptimal token regions
 1126 in the draft results are systematically identified and substituted with higher-quality tokens during
 1127 verification by the target model. This approach maintains the overall compositional integrity while
 1128 significantly enriching the detail and quality of the generated images.

1130 G.2 ACCEPTANCE-REJECTION SAMPLING

1132 Standard acceptance-rejection sampling favors large total variance $p(x) - q(x)$ for high acceptance
 1133 rate, while speculative decoding prefers small differences for better draft model alignment. There-
 fore, we observe the practical total inference wall-time of target MAR model during the whole

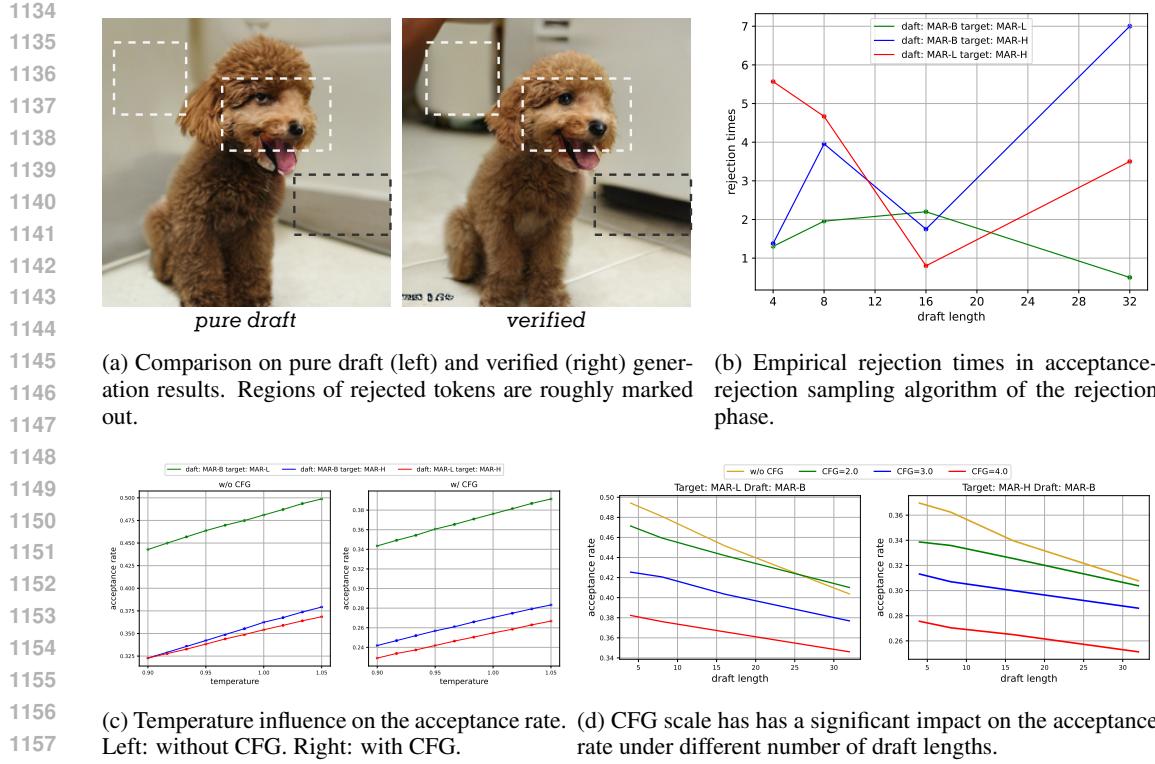


Figure 10: Ablation studies on various experiment factors conducted on MAR models.

Draft length γ	4	8	16	32
Target model runtime	64s	56s	53s	51s
Rejection-sampling rutime	0.0378s	0.0189s	0.0095s	0.0047s

Table 8: Overall runtime of target model and acceptance-rejection sampling during the whole speculative decoding process at different draft length γ . Batch size is set to 1.

speculative decoding process at batch size=1, and the wall-time and sampling steps of acceptance-rejection sampling. As show in Table 8, the rejection-sampling actually only accounts for a quite small fraction of the model runtime. Overall the inference speed is improved by speculative decoding. Figure 10b illustrates the relationship between the rejection times and the draft length. Empirically, acceptance-rejection sampling often requires only a few sampling steps. The runtime consumed by this sampling process is negligible compared to the overall model inference time.

G.3 TEMPERATURE

Temperature τ is a crucial hyperparameter during the denoising process in MAR. The temperature setting affects the consistency between the outputs of the draft and target models. Figure 10c illustrates the impact of the temperature τ on the acceptance rate during the generation process. The number of drafts length is set to 8. The temperature influences the PDF of the final output distribution; a lower temperature may result in a sharper distribution, while a higher temperature may lead to a flatter distribution. The ratio $p(x)/q(x)$ can be influenced based on this.

G.4 CLASSIFIER-FREE GUIDANCE

Figure 10d illustrates the relationship between draft length and the acceptance rate under different CFG scales. As the CFG scale increases, there is an overall trend of decreasing acceptance rates. This trend remains consistent mainly across each draft length. This phenomenon may indicate that

1188 the inconsistency between the draft and target models may increase as class guidance strengthens,
 1189 reducing the acceptance rate.
 1190

1191 G.5 COMPARISON WITH MASKED GENERATION

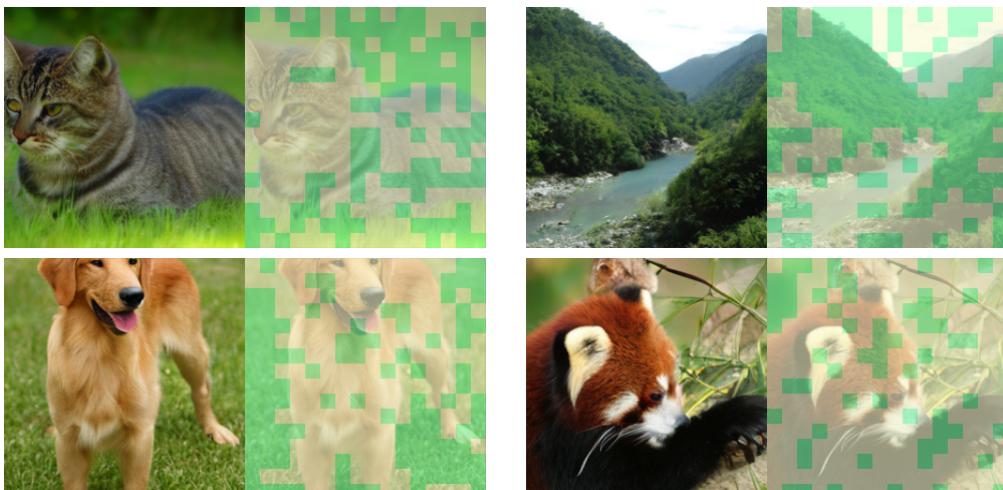
1193 MAR can achieve generation acceleration by generating multiple tokens per step. However, the cost
 1194 of acceleration comes with a noticeable performance loss. Speculative decoding typically provides
 1195 a $2 \times$ speedup, consistent with the conclusion drawn in the application of LLMs (Leviathan et al.,
 1196 2023). The performance is well maintained, as theoretically proved. In contrast, while masked
 1197 generation can achieve a higher acceleration ratio (up to $10 \times$ when the number of masks reaches
 1198 32), it cannot maintain performance, as shown in Table 9. We show the speedup ratio on MAR-H
 1199 with 256 batch size. As the acceleration ratio increases, the model’s performance suffers significant
 1200 degradation.

# mask	2	4	8	16	32	64
speed up	1.99	3.49	5.54	7.85	10.02	11.49
FID	1.56	2.37	3.66	4.99	17.43	59.07

1204 Table 9: Speed-up ratio and FID under different mask generation steps. Larger mask generation step
 1205 can bring the model a better speedup, but it also leads to significant performance degradation.

1208 G.6 VISUALIZATION OF ACCEPTANCE

1210 We visualize the acceptance and rejection region of each position through a 2D heatmap. As shown
 1211 in Figure 11, dark green blocks represent accepted tokens, while light green blocks represent re-
 1212 jected. We observe that tokens representing backgrounds or regions with simpler textures tend to be
 1213 accepted. In contrast, more detailed positions are more likely to be rejected.

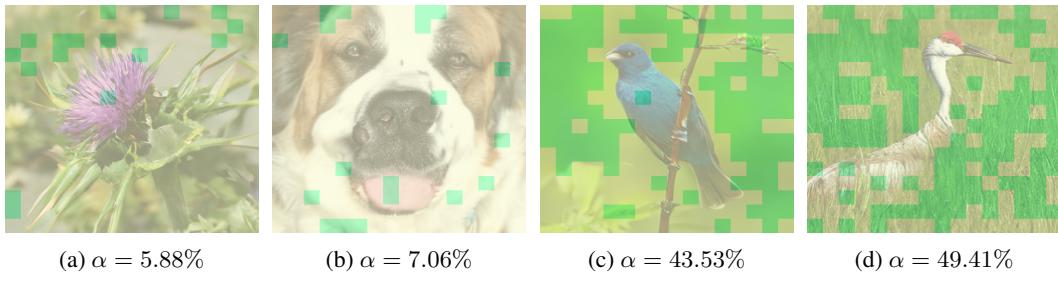


1231 Figure 11: Visualizations of accepted token heatmap. Dark green: accepted. Light green: rejected.
 1232

1234 G.7 FAILURE MODES

1236 We present a visualization of the images associated with the observed failure modes and success
 1237 modes. The resulting visualization is presented in Figure 12. Subfigures (a) and (b) depict the failure
 1238 modes, while subfigures (c) and (d) illustrate the success modes. Our analysis reveals that the failure
 1239 modes are predominantly localized in regions exhibiting high levels of detail and intricate texture
 1240 representation. Subfigures (a) and (b) are characterized by rich details and fine textures, where the
 1241 limited capacity of the draft model results in generated content that is below an acceptable quality
 threshold. In contrast, subfigures (c) and (d) possess comparatively lower complexity in terms of

1242 detail. The two subfigures include substantial background area devoid of pronounced details. The
 1243 acceptance rate is substantially elevated within these less-detailed regions.
 1244



1254 Figure 12: Subfigure (a) and (b): failure modes. Subfigure (c) and (d): success modes. The visual-
 1255 izations reveal that the failure modes are predominantly localized in regions exhibiting high levels
 1256 of detail and intricate texture representation.
 1257

H MORE QUALITATIVE RESULTS

1261 In Figure 13, 14, 15, 16 and 17, we provide more additional images generated under our contin-
 1262 uous speculative decoding by MAR compared with the target model only. While the target model
 1263 has achieved satisfactory quality in generating realistic and high-fidelity images, our continuous
 1264 speculative decoding can show comparable performance, similar generation results, and much faster
 1265 inference speed.
 1266

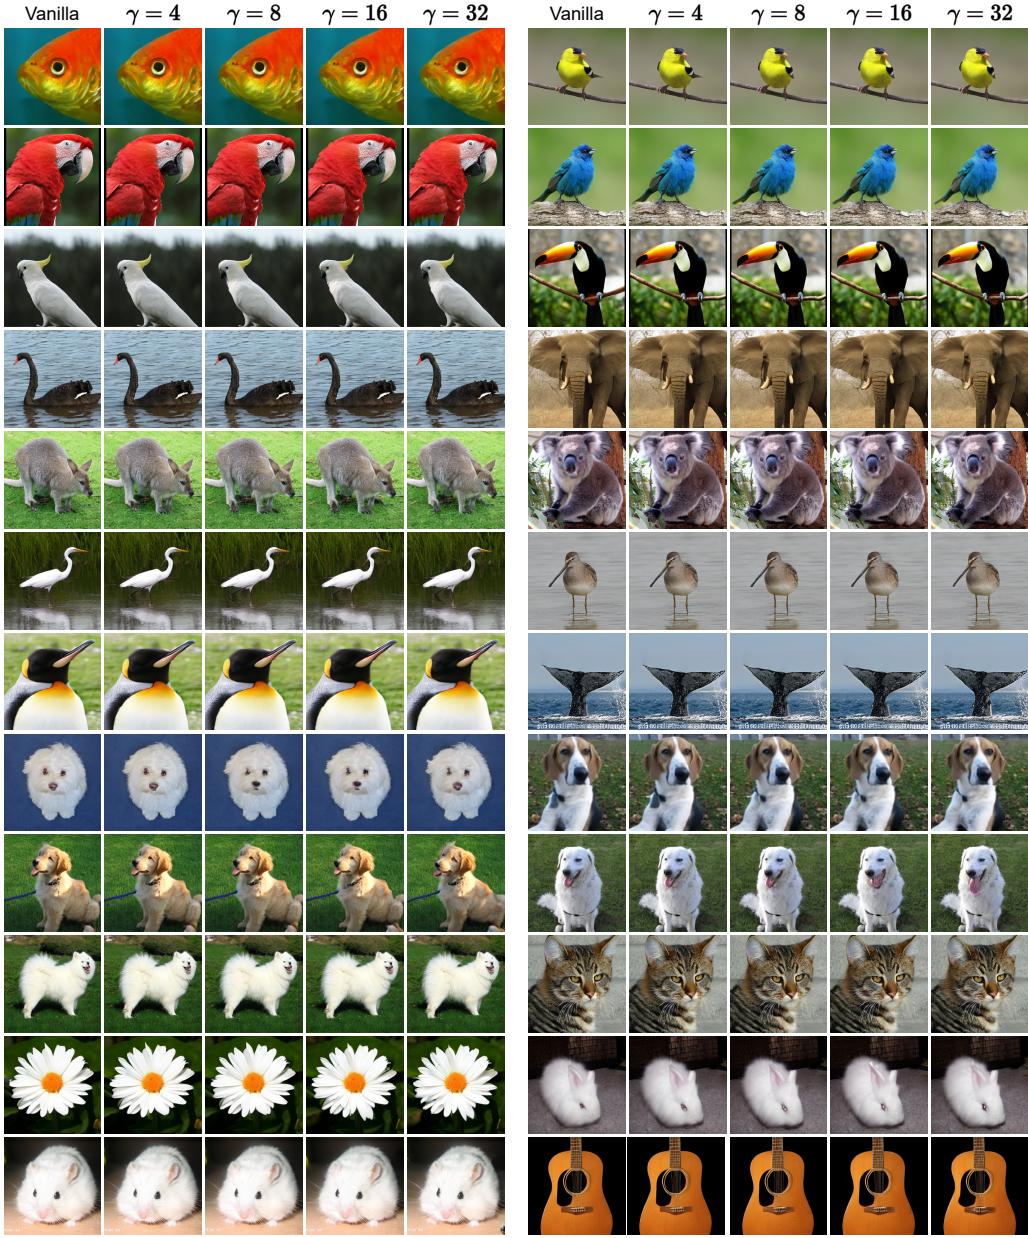
1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295

1296
1297
1298
1299
1300
1301

1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341

1342 Figure 13: Visual quality with increasing draft length γ compared with vanilla target model only
 1343 generation. *Best viewed zoom-in.*

1344
1345
1346
1347
1348
1349



1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

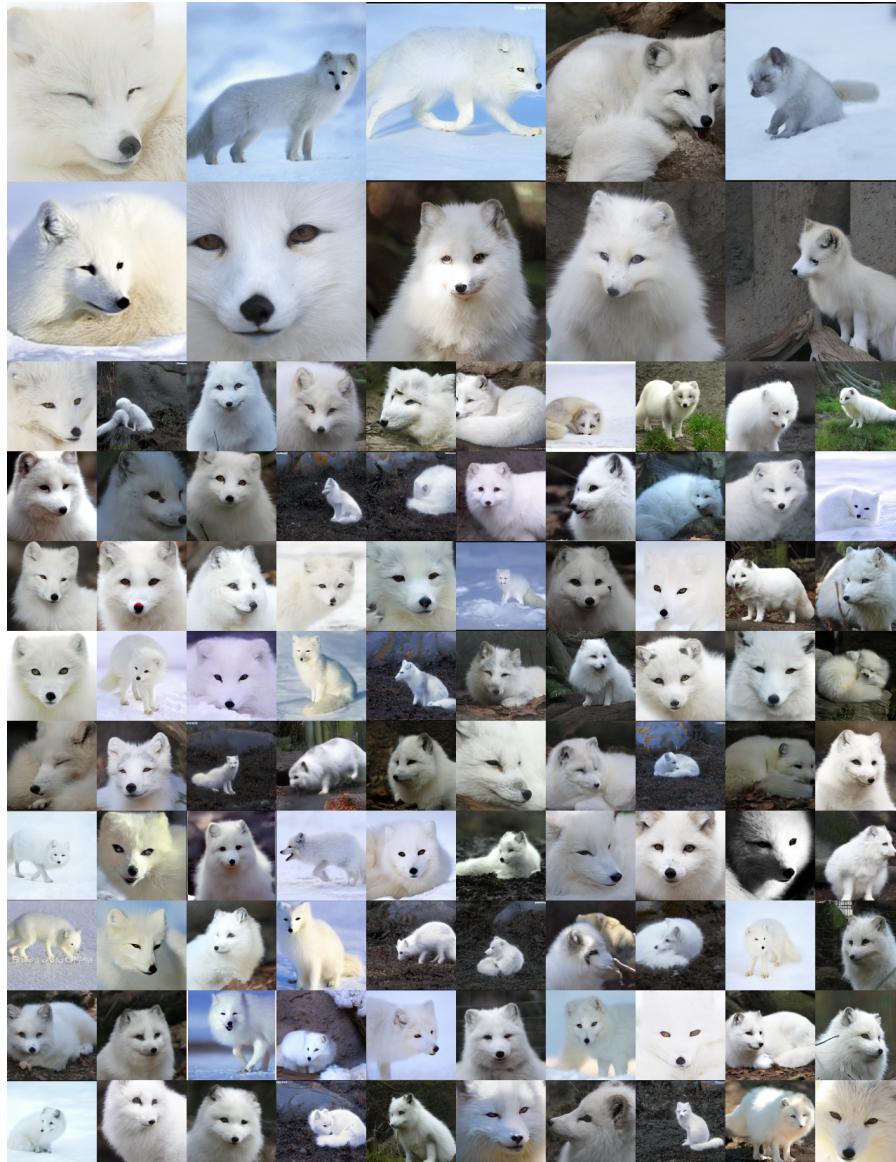


Figure 14: Visualization examples under $\gamma = 4$. Class label: arctic fox (297).

Figure 15: Visualization examples under $\gamma = 8$. Class label: balloon (417).

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

Figure 16: Visualization examples under $\gamma = 16$. Class label: ice cream (928).

Figure 17: Visualization examples under $\gamma = 32$. Class label: volcano (980).