

Parallelized Autoregressive Visual Generation

Yuqing Wang¹ Shuhuai Ren³ Zhijie Lin^{2†} Yujin Han¹
 Haoyuan Guo² Zhenheng Yang² Difan Zou¹ Jiashi Feng² Xihui Liu^{1*}
¹University of Hong Kong ²ByteDance ³Peking University

Abstract

Autoregressive models have emerged as a powerful approach for visual generation but suffer from slow inference speed due to their sequential token-by-token prediction process. In this paper, we propose a simple yet effective approach for parallelized autoregressive visual generation that improves generation efficiency while preserving the advantages of autoregressive modeling. Our key insight is that parallel generation depends on visual token dependencies—tokens with weak dependencies can be generated in parallel, while strongly dependent adjacent tokens are difficult to generate together, as their independent sampling may lead to inconsistencies. Based on this observation, we develop a parallel generation strategy that generates distant tokens with weak dependencies in parallel while maintaining sequential generation for strongly dependent local tokens. Our approach can be seamlessly integrated into standard autoregressive models without modifying the architecture or tokenizer. Experiments on ImageNet and UCF-101 demonstrate that our method achieves a $3.6\times$ speedup with comparable quality and up to $9.5\times$ speedup with minimal quality degradation across both image and video generation tasks. We hope this work will inspire future research in efficient visual generation and unified autoregressive modeling. Project page: <https://yuqingwang1029.github.io/PAR-project>.

1. Introduction

Autoregressive modeling has achieved remarkable success in language modeling [2, 29, 30, 45, 46], inspiring its application to visual generation [5, 6, 9, 19, 26, 27, 31, 32, 34, 42, 44, 48, 49, 58]. These models show great potential for visual tasks due to their strong scalability and unified modeling capabilities [11, 43, 54]. Current autoregressive visual generation approaches typically rely on a sequential token-by-token generation paradigm: visual data is first encoded into token sequences using an autoencoder [9, 50, 60],




Figure 1. **Comparison of different parallel generation strategies.** Both strategies generate initial tokens [1,2,3,4] sequentially then generate multiple tokens in parallel per step, following the order [5a-5d] to [6a-6d] to [7a-7d], etc. (a) Our approach generates weakly dependent tokens across non-local regions in parallel, preserving coherent patterns and local details. (b) The naive method generates strongly dependent tokens within local regions simultaneously, while independent sampling for strongly correlated tokens can cause inconsistent generation and disrupted patterns, such as distorted tiger faces and fragmented zebra stripes.

then an autoregressive transformer [51] is trained to predict these tokens following a raster scan order [6]. However, this strictly sequential generation process leads to a slow generation speed, severely limiting its practical applications [23, 54]. In this work, we aim to develop an efficient autoregressive visual generation approach that improves generation speed while maintaining the generation quality.

An intuitive way to improve generation efficiency is to predict multiple tokens in parallel at each step. In language modeling, methods like speculative decoding [4, 20, 22] and Jacobi decoding [17, 38] achieve parallel generation through auxiliary draft models or iterative refinement. In the visual domain, approaches like MaskGIT [3] employ non-autoregressive paradigms with masked modeling strategies, while VAR [44] achieves faster speed through next-scale prediction that requires specially designed multi-

†Project lead. *Corresponding author.



Figure 2. **Visualization comparison of our parallel generation and traditional autoregressive generation (LlamaGen [42]).** Our approach (PAR) achieves $3.6\text{-}9.5\times$ speedup over LlamaGen with comparable quality, reducing the generation time from 12.41s to 3.46s (PAR-4 $\times$) and 1.31s (PAR-16 $\times$) per image. Time measurements are conducted with a batch size of 1 on a single A100 GPU.

scale tokenizers and longer token sequences. However, the introduction of additional models and specialized architectures increases model complexity, and may limit the flexibility of autoregressive models as a unified solution across different modalities.

In this work, we ask: *can we achieve parallel visual generation while maintaining the simplicity and flexibility of standard autoregressive models?* We find that parallel generation is closely tied to token dependencies—tokens with strong dependencies need sequential generation, while weakly dependent tokens can be generated in parallel. In autoregressive models, each token is generated through sampling (e.g., top-k) to maintain diversity. Parallel generation requires independent sampling of multiple tokens simultaneously, but the joint distribution of highly dependent tokens cannot be factorized for independent sampling, leading to inconsistent predictions, as demonstrated by the distorted local patterns in Fig. 1 (b). For visual data, such dependencies are naturally correlated with spatial distances—while locally adjacent tokens exhibit strong dependencies, spatially distant tokens often have weak correlations. This motivates us to reconsider how to organize tokens for generation: by identifying spatially distant tokens with weak correlations, we can group them for simultaneous prediction. Such non-local grouping allows us to maintain sequential generation for strongly dependent local tokens while enabling parallel generation across differ-

ent spatial regions. Moreover, we observe that initial tokens in each local region play a crucial role in establishing the global structure - generating them in parallel could lead to conflicting structures across regions, such as repeated parts in different regions without global coordination (see middle row of Fig. 5). Therefore, the initial tokens in each local region should be generated sequentially to establish the global visual structure.

Based on these insights, we propose a simple yet effective approach for parallel generation in autoregressive visual models. Our key idea is to identify and group weakly dependent visual tokens for simultaneous prediction while maintaining sequential generation for strongly dependent ones. To achieve this, we first divide the image into local regions and generate their initial tokens sequentially to establish global context, then perform parallel generation by identifying and grouping tokens at corresponding positions across spatially distant regions. The process is illustrated in Fig. 1 (a). Our approach can be seamlessly implemented within standard autoregressive transformers through a re-ordering mechanism, with a few learnable token embeddings to facilitate the transition between sequential and parallel generation modes. By ensuring each prediction step has access to all previously generated tokens across regions, we maintain the autoregressive property and preserve global context modeling capabilities. With non-local parallel generation, our approach significantly reduces the number of

inference steps and thereby accelerates generation, while maintaining comparable visual quality through careful token dependency handling.

We verify the effectiveness of our approach on both image and video generation tasks using ImageNet [7] and UCF-101 [39] datasets. For image generation, our method achieves around $3.9\times$ fewer generation steps and $3.6\times$ actual inference-time speedup with comparable generation quality. With more aggressive parallelization, we achieve around $11.3\times$ reduction in steps and $9.5\times$ actual speedup with minimal quality drop (within 0.7 FID for image and 10 FVD for video). The qualitative comparison of generation results between our method and the baseline is shown in Fig. 2. The experiments demonstrate the effectiveness of our approach across different visual domains and its compatibility with various tokenizers like VQGAN [9] and MAGVIT-v2 [60].

In summary, we propose a simple yet effective parallelized autoregressive visual generation approach that carefully handles token dependencies. Our key idea is to identify and group weakly dependent tokens for simultaneous prediction while maintaining sequential generation for strongly dependent ones. Our approach can be seamlessly integrated into standard autoregressive models without architectural modifications. Through extensive experiments with different visual domains and tokenization methods, we demonstrate considerable speedup while preserving generation quality, making autoregressive visual generation more practically usable for real-world applications.

2. Related Work

Autoregressive Visual Generation. Autoregressive modeling has been explored in visual generation for years, from early pixel-based approaches [34, 48, 49] to current token-based methods. Modern approaches typically follow a two-stage paradigm: first compressing visual data into compact token sequences through discrete tokenizers [9, 50, 60], then training a transformer [51] to predict these tokens autoregressively in raster scan order [19, 31, 58]. This paradigm has been successfully extended to video generation [16, 55, 56], where tokens from different frames are predicted sequentially. However, the strictly sequential generation process leads to slow inference speed that scales with sequence length.

Parallel Prediction in Sequential Generation. Various approaches have been proposed to accelerate sequential generation. In language modeling, speculative decoding [4, 20, 22] employs a draft model to generate candidate tokens for main model verification, while Jacobi decoding [17, 38] enables parallel generation through iterative refinement. In visual generation, MaskGIT [3] adopts a non-autoregressive approach with BERT-like masked modeling strategies, taking a different modeling paradigm from

traditional autoregressive generation. VAR [44] proposes next-scale prediction that progressively generates tokens at increasing resolutions, though requiring specialized multi-level tokenizers and longer token sequences. In contrast, our approach enables efficient parallel generation while preserving the autoregressive property and model simplicity, readily applicable to various visual tasks without specialized architectures or additional models.

3. Method

In this section, we present our approach for parallelized visual autoregressive generation. We first discuss the relationship between token dependencies and parallel generation in Sec. 3.1. Based on these insights, we propose our parallel generation approach in Sec. 3.2. Finally, we present the model architecture and implementation details that realize this process within autoregressive transformers in Sec. 3.3.

3.1. Token Dependencies and Parallel Generation

Standard autoregressive models adopt token-by-token sequential generation, which significantly limits generation efficiency. To improve efficiency, we explore the possibility of generating multiple tokens in parallel. However, a critical question arises: *which tokens can be generated in parallel without compromising generation quality?* In this section, we analyze the relationship between token dependencies and parallel generation through pilot studies, providing guidance for designing parallelized autoregressive visual generation models.

Pilot Study. In language modeling, researchers have attempted to group adjacent tokens for multi-token prediction [4, 17, 20, 22, 38, 40, 52]. However, our pilot study reveals that directly predicting adjacent tokens leads to significant quality degradation in visual generation (see Fig. 1(b) and Tab. 4 (d)). In autoregressive generation, each token is generated through sampling strategies (e.g., top-k) to maintain diversity. When generating multiple tokens in parallel, these tokens need to be sampled independently. However, for adjacent visual tokens with strong dependencies, their joint distribution cannot be factorized into independent distributions, as each token is heavily influenced by its neighbors. The impact of such independent sampling is clearly demonstrated in the figure, where generating adjacent tokens in parallel leads to inconsistent local structures like distorted tiger faces and fragmented zebra stripes, as tokens are sampled without considering their neighbors' decisions.

Design Principles. These observations suggest that parallel generation should focus on weakly correlated tokens to minimize the impact of independent sampling. For visual tokens, dependencies naturally decrease with spatial distance - tokens from distant regions typically have weaker correlations than adjacent ones. This motivates us to per-




Figure 3. **Illustration of our non-local parallel generation process.** Stage 1: sequential generation of initial tokens (1-4) for each region (separated by dotted lines) to establish global structure. Stage 2: parallel generation at aligned positions across different regions (e.g., 5a-5d), then moving to next aligned positions (6a-6d, 7a-7d, etc.) for parallel generation. Same numbers indicate tokens generated in the same step, and letter suffix (a,b,c,d) denotes different regions .

form parallel generation across distant regions rather than within local neighborhoods. However, we find that not all distant tokens can be generated in parallel. The initial tokens of each regions are particularly crucial as they jointly determine the global image structure. Parallel generation of these initial tokens, despite their spatial distances, could lead to conflicting global decisions, resulting in issues like repeated patterns or incoherent patches across regions (see the middle row in Fig. 5).

Based on these insights, we propose three key design principles for parallelized autoregressive generation: 1) generate initial tokens for each region sequentially to establish proper global structure; 2) maintain sequential generation within local regions where dependencies are strong; and 3) enable parallel generation across regions where dependencies are weak through proper token organization.

3.2. Non-Local Parallel Generation

Based on the above principles, we propose our approach that enables parallel token prediction while maintaining autoregressive properties. The process is illustrated in Fig. 3.

Cross-region Token Grouping. Let $\{v_i\}_{i=1}^{H \times W}$ denote a sequence of visual tokens arranged in a $H \times W$ grid. We first partition the token grid into $M \times M$ regions. Each region contains $k := \left(\frac{H}{M} \times \frac{W}{M}\right)$ tokens. We then group tokens at corresponding positions across different regions. Let $v_j^{(r)}$ denote the token at position j in region r , where $r \in \{1, \dots, M^2\}$ and $j \in \{1, \dots, k\}$. We then organize these tokens into groups based on their corresponding positions across regions:

$$\left\{ [v_1^{(1)}, \dots, v_1^{(M^2)}], [v_2^{(1)}, \dots, v_2^{(M^2)}], \dots, [v_k^{(1)}, \dots, v_k^{(M^2)}] \right\}. \quad (1)$$

This organization groups together tokens at the same relative position across different regions, facilitating our parallel generation process.

Stage 1: Sequential Generation of Initial Tokens of Each Region. We first generate one initial token for each region sequentially (marked as “1-4” in Fig. 3) to establish the global context. As shown in Fig. 3 (1), we start with the top-left region and generate the initial token for each region

by sampling from the conditional probability distribution:

$$v_1^{(i)} \sim \mathbb{P}(v_1^{(i)} | v_1^{(<i)}), \quad i \in \{1, \dots, M^2\}, \quad (2)$$

where $v_1^{(i)}$ denotes the initial token of the i -th region. Since the number of regions (M^2) is small and fixed, this sequential generation introduces minimal overhead while providing crucial global context for subsequent parallel generation.

Stage 2: Parallel Generation of Cross-region Tokens.

After initializing tokens for all regions, we proceed with parallel generation of the remaining tokens. As illustrated in Fig.3 (2), at each step, we identify the next position j within each region following a raster scan order and simultaneously predict tokens at this position across all regions (e.g., tokens 5a-5d are generated in parallel). The parallel generation at each step can be formulated as:

$$\{v_j^{(r)}\}_{r=1}^{M^2} \sim \mathbb{P}(\{v_j^{(r)}\}_{r=1}^{M^2} | v_{<j}), \quad (3)$$

where $\{v_j^{(r)}\}_{r=1}^{M^2}$ represents the set of tokens at position j across all regions to be generated in parallel, and $v_{<j}$ includes both initial tokens and tokens from previous parallel steps. For example, with $M = 2$ on a 24×24 token grid, after generating 4 initial tokens sequentially, we predict $M^2 = 4$ tokens in parallel at each subsequent step, reducing the total number of generation steps from 576 to 147 (i.e., $4 + \frac{576-4}{4}$). While enabling parallel prediction, our approach maintains the autoregressive property as each prediction is still conditioned on all previous tokens. The key difference is that tokens at corresponding positions across regions, which exhibit weak dependencies, are now generated simultaneously instead of sequentially.

3.3. Model Architecture Details

We illustrate our parallel generation framework using a standard autoregressive transformer for class-conditioned image generation.

Framework Implementation. As shown in Fig. 4 (a), our model architecture consists of an autoregressive transformer that processes the input sequence and generates visual tokens. The input sequence begins with a class token




Figure 4. **Overview of our parallel autoregressive generation framework.** (a) **Model implementation.** The model first generates initial tokens sequentially [1,2,3,4], then uses learnable tokens [M1,M2,M3] to help transition into parallel prediction mode. (b) **Comparison of visible context between our parallel prediction approach (left) and traditional single-token prediction (right).** The colored cells indicate available context during generation. In traditional AR, when predicting token 6d, the model can access all previous tokens including 6a – 6c. Without full attention, our parallel approach would limit each token (e.g., 6b) to only see tokens up to the same position in the previous group (e.g., up to 5b). We enable group-wise full attention to allow access to the entire previous group.

(C) followed by visual tokens to be generated. To achieve n -token parallel prediction, we design a special sequence structure with three distinct parts: 1) initial sequential tokens [1,2,...,n] that are generated one at a time, 2) a transition part with $n - 1$ learnable tokens [M1,M2,M3] that helps the model enter parallel prediction mode, and 3) subsequent token groups that are predicted n tokens at a time (e.g., [5a, 5b, 5c, 5d], [6a, 6b, 6c, 6d]). For predicting each group, the model takes all previous tokens as input while maintaining a fixed offset of n tokens between input and target sequences. The learnable tokens share the same dimension as regular tokens for seamless integration. To maintain spatial relationships under our reordered sequence, we employ 2D Rotary Position Embedding (RoPE) [41], which preserves each token’s original spatial position information regardless of its sequence position. The above designs enable parallel prediction while preserving the standard autoregressive transformer architecture.

Group-wise Bi-directional Attention with Global Autoregression. Our framework combines sequential generation of initial tokens with parallel generation of subsequent token groups. As illustrated in Fig.4 (b), in traditional autoregressive models, when predicting token 6d, the model can access all previous tokens including 6a – 6c. However, naive parallel generation with causal masking would restrict each token (e.g., 6b) to only see tokens up to the same position in the previous group (e.g., up to 5b), limiting the available context. To address this limitation while maintaining parallelism, we enable bi-directional attention within each prediction group while preserving causal attention between groups. This allows each token in the current group to access the entire previous group as context (e.g., all tokens [5a – 5d] are visible when predicting any token in [6a – 6d]). This design enriches the local context for parallel prediction while maintaining the global autoregressive property, ensuring compatibility with standard optimizations like KV-cache.

Model	Params	Layers	Hidden	Heads
PAR-L	343M	24	1024	16
PAR-XL	775M	36	1280	20
PAR-XXL	1.4B	48	1536	24
PAR-3B	3.1B	24	3200	32

Table 1. **Model sizes and architecture configurations of PAR.** The configurations are following previous works [25, 30, 42, 45].

Extension to Video Generation. Our parallel generation framework can be naturally extended to video generation. The tokenization process reduces both spatial and temporal dimensions, resulting in tokens arranged in a $T \times H \times W$ grid, where each latent frame aggregates information from multiple input frames. We treat these temporally compressed tokens similarly to image tokens and apply our parallel generation strategy along the spatial dimensions, with the only modification being the use of 3D position embeddings. While we also explored parallel generation along the temporal dimension, we found it less effective than spatial parallelization. This is because temporal dependencies exhibit stronger sequential characteristics that are fundamental to video coherence, making them less suitable for parallel prediction compared to spatial relationships. The exploration of effective temporal parallel strategies remains as future work.

4. Experiments

4.1. Experimental Setup

Image Generation. For fair comparison with existing token-by-token autoregressive visual generation methods, we adopt similar settings as [42], using a VQGAN tokenizer [9] with a 16,384 codebook size and $16 \times$ downsampling ratio. Models are trained on ImageNet-1K [7] for 300 epochs, with 384×384 images tokenized into 24×24 sequences. We evaluate on the ImageNet validation set at 256×256 resolution using FID [12] as the primary met-



Figure 5. **Qualitative comparison of parallel generation strategies.** **Top:** Our method with sequential initial tokens followed by parallel distant token prediction produces high-quality and coherent images. **Middle:** Direct parallel prediction without sequential initial tokens leads to inconsistent global structures. **Bottom:** Parallel prediction of adjacent tokens results in distorted local patterns and broken details.

Type	Model	#Para.	FID↓	IS↑	Precision↑	Recall↑	Steps	Time(s)↓
GAN	BigGAN [1]	112M	6.95	224.5	0.89	0.38	1	—
	GigaGAN [15]	569M	3.45	225.5	0.84	0.61	1	—
	StyleGan-XL [35]	166M	2.30	265.1	0.78	0.53	1	0.08
Diffusion	ADM [8]	554M	10.94	101.0	0.69	0.63	250	44.68
	CDM [13]	—	4.88	158.7	—	—	8100	—
	LDM-4 [33]	400M	3.60	247.7	—	—	250	—
	DiT-XL/2 [28]	675M	2.27	278.2	0.83	0.57	250	11.97
Mask	MaskGIT [3]	227M	6.18	182.1	0.80	0.51	8	0.13
VAR	VAR-d30 [44]	2B	1.97	334.7	0.81	0.61	10	0.27
MAR	MAR [21]	943M	1.55	303.7	0.81	0.62	64	28.24
AR	VQGAN [9]	227M	18.65	80.4	0.78	0.26	256	5.05
	VQGAN [9]	1.4B	15.78	74.3	—	—	256	5.05
	VQGAN-re [9]	1.4B	5.20	280.3	—	—	256	6.38
	ViT-VQGAN [57]	1.7B	4.17	175.1	—	—	1024	>6.38
	ViT-VQGAN-re [57]	1.7B	3.04	227.4	—	—	1024	>6.38
	RQTran. [19]	3.8B	7.55	134.0	—	—	256	5.58
	RQTran.-re [19]	3.8B	3.80	323.7	—	—	256	5.58
AR	LlamaGen-L [42]	343M	3.07	256.1	0.83	0.52	576	12.58
	LlamaGen-XL [42]	775M	2.62	244.1	0.80	0.57	576	18.66
	LlamaGen-XXL [42]	1.4B	2.34	253.9	0.80	0.59	576	24.91
	LlamaGen-3B [42]	3.1B	2.18	263.3	0.81	0.58	576	12.41
AR	PAR-L-4×	343M	3.76	218.9	0.84	0.50	147	3.38
	PAR-XL-4×	775M	2.61	259.2	0.82	0.56	147	4.94
	PAR-XXL-4×	1.4B	2.35	263.2	0.82	0.57	147	6.84
	PAR-3B-4×	3.1B	2.29	255.5	0.82	0.58	147	3.46
	PAR-XXL-16×	1.4B	3.02	270.6	0.81	0.56	51	2.28
	PAR-3B-16×	3.1B	2.88	262.5	0.82	0.56	51	1.31

Table 2. **Class-conditional image generation on ImageNet 256×256 benchmark.** “↓” or “↑” indicate lower or higher values are better. “-re” means using rejection sampling. PAR-4× and PAR-16× means generating 4 and 16 tokens per step in parallel, respectively.

ric, complemented by IS and Precision/Recall [18]. We experiment with model sizes from 343M to 3.1B parameters (Tab.1), reporting both generation steps and latency time.

Video Generation. We evaluate on the UCF-101 [39] dataset using MAGVIT-v2 tokenizer [60] reproduced by [32]. Each 17-frame video (128×128 resolution) is compressed by $8 \times$ spatially and $4 \times$ temporally into a $5 \times 16 \times 16$ token sequence (1280 tokens per video). For fair comparison, we implement both next-token prediction and our parallel generation approach using the same architecture. The position of video codes is encoded via 3D positional embeddings. Our reproduced MAGVIT-v2 tokenizer uses a 64K visual vocabulary instead of the original 262K to facilitate model training. We use Fréchet Video Distance (FVD) [47] to evaluate generation quality.

Detailed training configurations for both video and image generation are provided in the supplementary material.

4.2. Main Results

4.2.1. Image Generation

Tab. 2 presents comprehensive comparisons of class-conditional image generation with various state-of-the-art methods, including GAN [1, 15, 35] (one-shot generation), Diffusion [8, 13, 28, 33] (iterative denoising), Mask [3] (mask token prediction), VAR [44] (next-scale prediction), MAR [21] (continuous mask token prediction), and AR [9, 42, 57] (autoregressive generation). Our PAR achieves competitive performance while maintaining faster inference speed than most state-of-the-art models. Specifically, when comparing with representative models from different categories, our method shows advantages. Compared with the mask-based method MaskGIT [3], our method achieves substantially better generation quality (FID 2.29 vs. 6.18) despite requiring more steps. For VAR [44], while it achieves slightly better FID (1.97 vs. 2.29), our method maintains a simpler framework with fewer tokens per image and preserves the pure autoregressive nature, making it more flexible for multi-modal integration.

Compared to our baseline model LlamaGen [42], PAR achieves $3.9 \times$ reduction in generation steps (147 vs. 576) and $3.58 \times$ speedup in wall-clock time (3.46s vs. 12.41s) while maintaining comparable quality (FID 2.29 vs. 2.18). With more aggressive parallelization, PAR-3B-16 $\times$ further accelerates generation to 1.31s ($9.5 \times$ speedup) with only 0.7 FID degradation compared to the baseline, demonstrating the effectiveness of our parallel generation strategy in balancing efficiency and quality.

4.2.2. Video Generation

We evaluate our approach on the UCF-101 [39] dataset for class-conditional video generation. Table 3 shows comparisons with various state-of-the-art methods across different categories. Among recent works, MAGVIT-v2 [60]

Type	Method	#Param	FVD↓	Steps	Time(s)
Diffusion	VideoFusion [24]	N/A	173	-	-
	Make-A-Video [36]	N/A	81.3	-	-
	HPDM-L [37]	725M	66.3	-	-
Mask.	MAGVIT [59]	306M	76	-	-
	MAGVIT-v2 [60]	840M	58	-	-
AR	CogVideo [14]	9.4B	626	-	-
	TATS [10]	321M	332	-	-
	OmniTokenizer [53]	650M	191	5120	336.70
	MAGVIT-v2-AR [60]	840M	109	1280	-
AR	PAR-1 $\times$	792M	94.1	1280	43.30
	PAR-4 $\times$	792M	99.5	323	11.27
	PAR-16 $\times$	792M	103.4	95	3.44

Table 3. **Comparison of class-conditional video generation methods on UCF-101 benchmark.** FVD measures generation quality, where lower values ($\downarrow$) indicate better performance. PAR-1 $\times$ represents our token-by-token baseline, while PAR-4 $\times$ and PAR-16 $\times$ indicate our parallel generation variants with different speedup ratios, achieving competitive FVD scores with significantly reduced generation steps and wall-clock time.

achieves strong performance with an FVD of 58 using masked token prediction, while its autoregressive variant MAGVIT-v2-AR obtains an FVD of 109 with 1280 generation steps. Our next-token-prediction baseline (PAR-1 $\times$) achieves a competitive FVD of 94.1, demonstrating the effectiveness of our implementation. More importantly, our parallel generation variants significantly reduce both generation steps and wall-clock time while maintaining comparable quality. Specifically, PAR-4 $\times$ reduces the generation steps from 1280 to 323 with minimal FVD increase (99.5 vs. 94.1), achieving $3.8 \times$ speedup (11.27s vs. 43.30s). Further parallelization with PAR-16 $\times$ achieves $12.6 \times$ speedup (3.44s vs. 43.30s) with 103.4 FVD, while reducing generation steps to 95. Due to space limit, we provide visualization results of video generation in supplementary materials.

4.3. Ablation Study

In this section, we conduct comprehensive ablation studies to investigate the effectiveness of our key design choices on the ImageNet 256 $\times$ 256 validation set (Tab. 4). Unless specified, we use the PAR-XL model with parallel group size $n=4$ as default setting.

Initial sequential token generation. We first evaluate the importance of initial sequential token generation by comparing models with and without this phase in Tab. 4 (a). Results show that initial sequential generation reduces FID from 3.67 to 2.61, with only 3 additional steps (147 vs. 144). We also visualize the comparison in Fig. 5. Without initial sequential generation (middle row), the generated images exhibit inconsistent global structures, such as misaligned dogs with duplicated body parts, as initial tokens are generated without awareness of each other. In contrast, our approach with initial sequential generation (top row) produces more coherent and natural-looking images. The

	FID↓	IS↑	steps↓
w/o	3.67	221.36	144
w	2.61	259.17	147

(a) **Importance of initial sequential token generation.** Sequential generation of initial tokens improves FID by 1.06 with negligible step increase.

n	FID↓	IS↑	steps↓
1	2.34	253.90	576
4	2.35	263.24	147
16	3.02	270.57	51

(b) **Number of parallel predicted tokens (PAR-XXL).** $n=1$ is the token-by-token baseline. $n=4$ reduces steps by $4\times$ with similar FID (2.35 vs. 2.34), while $n=16$ reduces steps by $11.3\times$ at the cost of 0.67 FID.

attn	FID↓	IS↑	steps↓
causal	3.64	228.08	147
full	2.61	259.17	147

(c) **Attention pattern between parallel tokens.** Full attention allows complete context access from previous parallel groups (vs. causal attention’s limited access), bringing 1.03 FID improvement.

order	pattern	FID↓	IS↑	steps↓
raster	one	2.62	244.08	576
distant	one	2.64	262.72	576
raster	multi	5.64	265.46	147
distant	multi	2.61	259.17	147

(d) **Comparison of different scan orders under single-token and multi-token prediction.** Our region-based distant ordering shows similar performance with raster scan in single-token setting, but significantly outperforms in multi-token prediction (2.61 vs. 5.64 FID).

Params	FID↓	IS↑	steps
343M	3.76	218.92	147
775M	2.61	259.17	147
1.4B	2.35	263.24	147
3.1B	2.29	255.46	147

(e) **Scaling of model size ($4\times$ parallel).** Generation quality steadily improves with more parameters, from 343M (FID 3.76) to 3.1B (FID 2.29).

Table 4. **Ablation studies on image generation model designs.**

results illustrate the importance of initial sequential token generation for establishing proper global structure.

Number of parallel predicted tokens. The number of tokens predicted in parallel (n) controls the trade-off between efficiency and quality. As shown in Tab. 4(b), with $n = 4$ ($M = 2$), our approach reduces generation steps from 576 to 147 while maintaining comparable quality (FID 2.35 vs. 2.34). Further increasing to $n = 16$ ($M = 4$) achieves more aggressive parallelization with only 51 steps, at the cost of slight quality degradation (FID increase of 0.67). This is consistent with our analysis that tokens from distant regions have weaker dependencies and can be generated in parallel. As shown in Fig. 2, both PAR- $4\times$ and PAR- $16\times$ preserve visual fidelity while achieving significant speedup (3.46s and

1.31s vs. 12.41s).

Impact of attention pattern. To enable effective parallel prediction while preserving rich context modeling, we study different attention patterns between parallel predicted tokens. With $n = 4$ parallel tokens, enabling full attention within groups reduces FID from 3.64 to 2.61 compared to causal attention, as it allows each token to access complete context from previous groups. This supports our design of combining bi-directional attention within groups with autoregressive attention between groups.

Impact of token ordering and prediction pattern. We compare raster scan and our distant ordering under different prediction settings. As shown in Tab. 4(d), while both achieve comparable quality in single-token prediction (FID 2.62 vs. 2.64), their performance differs significantly with multi-token prediction - raster scan degrades severely (FID 5.64) while our distant ordering maintains quality (FID 2.61). This indicates that the choice of parallel predicted tokens is critical. When using raster scan, adjacent tokens with strong dependencies are forced to generate simultaneously, leading to distorted local patterns as shown in Fig. 5 (bottom row). In contrast, our region-based distant ordering groups weakly correlated tokens for parallel prediction, preserving both local details and global coherence (top row).

Model scaling analysis. In Tab. 4(e), we study how our parallel prediction approach scales with model size. With $n = 4$ parallel tokens, increasing model size from 343M to 3.1B parameters steadily improves generation quality (FID decreases from 3.76 to 2.29). Comparing with sequential generation baseline (LlamaGen) in Tab. 2, while smaller models show a noticeable quality gap (343M: FID 3.76 vs. 3.07), larger models achieve comparable performance (775M: 2.61 vs. 2.62; 1.4B: 2.35 vs. 2.34) while reducing generation steps from 576 to 147. This demonstrates that increased model capacity helps mitigate the quality trade-off from parallel prediction, suggesting stronger capability in modeling joint distribution of parallel tokens.

5. Conclusion

We propose PAR, a novel autoregressive visual generation approach that enables efficient parallel generation while preserving the advantages of autoregressive modeling. Our key finding is that the feasibility of parallel generation depends on token dependencies - tokens with weak dependencies can be generated in parallel while strongly dependent tokens lead to inconsistent results. Based on this insight, our PAR organizes tokens based on their dependency strengths rather than spatial proximity. The effectiveness of our approach across different visual domains validates this strategy for efficient autoregressive visual generation. We hope our work can inspire future research on visual generation and other sequence prediction tasks.

Acknowledgements

This work is partially supported by the National Nature Science Foundation of China (No. 62402406). The authors would like to thank Peize Sun, Yi Jiang, Xian Liu and Yao Teng for helpful discussions on autoregressive models.

References

- [1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 6, 7
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Sastry, et al. Language models are few-shot learners. In *NeurIPS*, 2020. 1
- [3] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked generative image transformer. In *CVPR*, 2022. 1, 3, 6, 7
- [4] Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*, 2023. 1, 3
- [5] Liang Chen, Zekun Wang, Shuhuai Ren, Lei Li, Haozhe Zhao, Yunshui Li, Zefan Cai, Hongcheng Guo, Lei Zhang, Yizhe Xiong, et al. Next token prediction towards multimodal intelligence: A comprehensive survey. *arXiv preprint arXiv:2412.18619*, 2024. 1
- [6] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *ICML*, pages 1691–1703, 2020. 1
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 3, 5
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, pages 8780–8794, 2021. 6, 7
- [9] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pages 12873–12883, 2021. 1, 3, 5, 6, 7
- [10] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In *ECCV*, pages 102–118, 2022. 7
- [11] Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*, 2020. 1
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017. 5
- [13] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research*, 23(1):2249–2281, 2022. 6, 7
- [14] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *ICLR*, 2022. 7
- [15] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *CVPR*, pages 10124–10134, 2023. 6, 7
- [16] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig Adam, Hassan Akbari, Yair Alon, Vignesh Birodkar, et al. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023. 3
- [17] Siqi Kou, Lanxiang Hu, Zhezhi He, Zhijie Deng, and Hao Zhang. Cllms: Consistency large language models. *arXiv preprint arXiv:2403.00835*, 2024. 1, 3
- [18] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *NeurIPS*, 32, 2019. 7
- [19] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *CVPR*, 2022. 1, 3, 6
- [20] Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *ICML*, 2023. 1, 3
- [21] Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image generation without vector quantization. *arXiv preprint arXiv:2406.11838*, 2024. 6, 7
- [22] Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. Eagle: Speculative sampling requires rethinking feature uncertainty. *arXiv preprint arXiv:2401.15077*, 2024. 1, 3
- [23] Dongyang Liu, Shitian Zhao, Le Zhuo, Weifeng Lin, Yu Qiao, Hongsheng Li, and Peng Gao. Lumina-mgpt: Illuminate flexible photorealistic text-to-image generation with multimodal generative pretraining. *arXiv preprint arXiv:2408.02657*, 2024. 1
- [24] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. *arXiv preprint arXiv:2303.08320*, 2023. 7
- [25] OpenLM-Research. Openllama 3b. https://huggingface.co/openlm-research/open_llama_3b, 2023. 5
- [26] Ziqi Pang, Tianyuan Zhang, Fujun Luan, Yunze Man, Hao Tan, Kai Zhang, William T. Freeman, and Yu-Xiong Wang. Randar: Decoder-only autoregressive visual generation in random orders. *arXiv preprint arXiv:2412.01827*, 2024. 1
- [27] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, pages 4055–4064, 2018. 1
- [28] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, pages 4195–4205, 2023. 6, 7
- [29] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 1

- [30] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. 1, 5
- [31] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, pages 8821–8831, 2021. 1, 3
- [32] Shuhuai Ren, Shuming Ma, Xu Sun, and Furu Wei. Next block prediction: Video generation via semi-autoregressive modeling. *arXiv preprint arXiv:2502.07737*, 2025. 1, 7
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 6, 7
- [34] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint arXiv:1701.05517*, 2017. 1, 3
- [35] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–10, 2022. 6, 7
- [36] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. In *ICLR*, 2022. 7
- [37] Ivan Skorokhodov, Willi Menapace, Aliaksandr Siarohin, and Sergey Tulyakov. Hierarchical patch diffusion models for high-resolution video generation. In *CVPR*, 2024. 7
- [38] Yang Song, Chenlin Meng, Renjie Liao, and Stefano Ermon. Accelerating feedforward computation via parallel nonlinear equation solving. In *ICML*, 2021. 1, 3
- [39] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 3, 7
- [40] Mitchell Stern, Noam Shazeer, and Jakob Uszkoreit. Block-wise parallel decoding for deep autoregressive models. *Advances in Neural Information Processing Systems*, 31, 2018. 3
- [41] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 5
- [42] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. 1, 2, 5, 6, 7
- [43] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 1
- [44] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024. 1, 3, 6, 7
- [45] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1, 5
- [46] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 1
- [47] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 7
- [48] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *NeurIPS*, 2016. 1, 3
- [49] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016. 1, 3
- [50] Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. 1, 3
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1, 3
- [52] Chunqi Wang, Ji Zhang, and Haiqing Chen. Semi-autoregressive neural machine translation. *arXiv preprint arXiv:1808.08583*, 2018. 3
- [53] Junke Wang, Yi Jiang, Zehuan Yuan, Binyue Peng, Zuxuan Wu, and Yu-Gang Jiang. Omnitokenizer: A joint image-video tokenizer for visual generation. *arXiv preprint arXiv:2406.09399*, 2024. 7
- [54] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. 1
- [55] Yuqing Wang, Tianwei Xiong, Daquan Zhou, Zhijie Lin, Yang Zhao, Bingyi Kang, Jiashi Feng, and Xihui Liu. Loong: Generating minute-level long videos with autoregressive language models. *arXiv preprint arXiv:2410.02757*, 2024. 3
- [56] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 3
- [57] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021. 6, 7
- [58] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022. 1, 3
- [59] Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video transformer. In *CVPR*, 2023. 7

- [60] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats diffusion—tokenizer is key to visual generation. In *ICLR*, 2024. [1](#), [3](#), [7](#)