

Improved Masked Image Generation with Knowledge-Augmented Token Representations

Guotao Liang^{1,2}, Baoquan Zhang^{1*}, Zhiyuan Wen², Zihao Han¹, Yunming Ye^{1,2}

¹Harbin Institute of Technology, Shenzhen;

²Peng Cheng Laboratory;

lianggt@pcl.ac.cn, {baoquanzhang, yeyunming}@hit.edu.cn, zihaoan@stu.hit.edu.cn, wenzhiyuan2012@gmail.com

Abstract

Masked image generation (MIG) has demonstrated remarkable efficiency and high-fidelity images by enabling parallel token prediction. Existing methods typically rely solely on the model itself to learn semantic dependencies among visual token sequences. However, directly learning such semantic dependencies from data is challenging because the individual tokens lack clear semantic meanings, and these sequences are usually long. To address this limitation, we propose a novel Knowledge-Augmented Masked Image Generation framework, named KA-MIG, which introduces explicit knowledge of token-level semantic dependencies (i.e., extracted from the training data) as priors to learn richer representations for improving performance. In particular, we explore and identify three types of advantageous token knowledge graphs, including two positive and one negative graphs (i.e., the co-occurrence graph, the semantic similarity graph, and the position-token incompatibility graph). Based on three prior knowledge graphs, we design a graph-aware encoder to learn token and position-aware representations. After that, a lightweight fusion mechanism is introduced to integrate these enriched representations into the existing MIG methods. Resorting to such prior knowledge, our method effectively enhances the model’s ability to capture semantic dependencies, leading to improved generation quality. Experimental results demonstrate that our method improves upon existing MIG for class-conditional image generation on ImageNet.

Code — <https://github.com/GuotaoLiang/KA-MIG>

Introduction

Class-conditional images are a fundamental task in generative modeling, which aims to synthesize realistic images conditioned on given semantic class labels. Traditional methods are mainly based on generative adversarial networks (GANs) (Brock, Donahue, and Simonyan 2018; Sauer, Schwarz, and Geiger 2022; Kang et al. 2023) which have achieved impressive visual quality and class-conditional controllability. In recent years, inspired by next-token prediction in natural language processing (Vaswani et al. 2017), the autoregressive generation paradigm based on discrete tokens has emerged as a promising alternative

*Corresponding Authors.

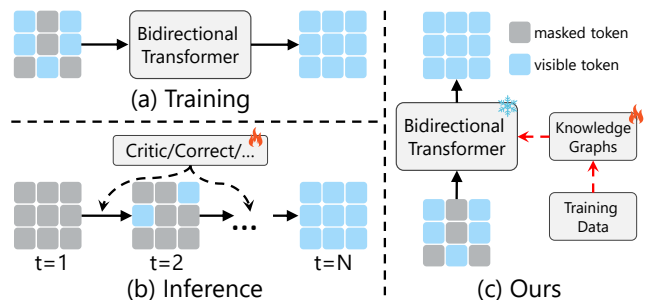


Figure 1: Motivation and framework overview. (a) The training pipeline of Masked Image Generation (MIG). (b) The inference process in MIG, where existing works primarily focus on improving sampling quality. (c) Our proposed framework incorporates external prior knowledge graphs to learn richer representations for enhancing MIG.

(Razavi, Van den Oord, and Vinyals 2019; Esser, Rombach, and Ommer 2021; Sun et al. 2024). However, these methods often suffer from slow sampling and degraded image quality due to the long visual token sequences.

Recent advances in masked image generation (MIG) frameworks, such as MaskGIT (Chang et al. 2022), formulate image synthesis as a parallel decoding process that iteratively predicts masked tokens. This approach achieves a favorable trade-off between sampling speed and quality, and has become a representative paradigm in token-based image generation. However, it still underperforms recent well-developed diffusion-based models (Rombach et al. 2022). To address this, many novel methods have been proposed to improve the generation performance, including introducing criticism or correction strategies (Lezama et al. 2022a,b), or proposing self-guidance sampling method (Hur et al. 2024), or devising a more uniform sampling strategy (Besnier et al. 2025). While these approaches have made notable progress, they primarily focus on refining the decoding strategy, leaving the intrinsic representation and modeling capacity of MaskGIT itself largely underexplored.

Existing methods mainly rely on transformer itself to learn visual token dependencies. However, directly learning such semantic dependencies from training data is challenging for two main reasons. First, individual tokens lack clear semantic meaning, making it difficult for the model

to interpret or relate them effectively. Second, visual token sequences are typically long (e.g., 256 tokens per image), which hinders the capture of complex relations and leads to training inefficiencies and unstable generation performance.

Motivated by the above limitation, we propose a novel *Knowledge-Augmented Masked Image Generation* framework, called KA-MIG. The novelty lies in leveraging prior semantic knowledge graphs (i.e., token-level relations) to guide the learning of richer token representations for improving generation performance. However, as mentioned above, each token lacks interpretable semantics in the dataset, making it difficult to obtain such knowledge graphs. To handle this, we adopt a data-driven approach to uncover latent structural regularities from large-scale training data. Its advantage lies in its efficiency and compatibility, as it requires neither external annotations nor manually designed rules. Follow this, we construct three types of advantageous token knowledge graphs, consisting of two positive priors and one negative prior: 1) co-occurrence graph, which captures frequent co-appearances of token pairs and reflects latent spatial-semantic correlations; 2) semantic similarity graph, which identifies tokens with analogous meanings in the context of image synthesis; 3) position-token incompatibility graph, which captures negative prior, identifying tokens that are incompatible with certain spatial locations conditioned on class labels. The incompatible means that tokens that have never appeared at that position. For example, in the “airplane” class, tokens representing ground or grass textures seldom occur in the upper-middle image region, as that space is predominantly occupied by the airplane body or sky. Building on these graphs, a lightweight fusion framework is introduced to integrate the enriched prior knowledge into the existing MIG methods. More specifically, for positive priors, we inject the prior features into each layer of the transformer in an additive manner. Different from positive priors, negative prior knowledge is incorporated using a subtractive manner. This simple yet effective fusion scheme guides the model towards capturing richer semantic dependencies and spatial constraints without introducing significant computational complexity.

The contributions of this work are summarized as follows:

- Unlike existing works that focus on improving sampling strategies, we propose a novel knowledge-augmented masked image generation framework, named KA-MIG, which leverages prior semantic knowledge graphs to enrich internal token representations and improve generation performance.
- We systematically construct three types of advantageous token knowledge graphs, consisting of two positive priors and one negative prior, i.e., co-occurrence graph, semantic similarity graph, and position-token incompatibility graph. Resorting to such prior knowledge, we introduce a lightweight additive–subtractive fusion mechanism to effectively integrate both positive and negative knowledge into the existing framework.
- Extensive experiments show that KA-MIG significantly enhances the generation performance of existing MIG models, demonstrating its effectiveness and versatility.

Related Works

Image generation, the task of synthesizing realistic images from noise or structured inputs, has long been a core challenge in computer vision. Generative adversarial networks (GANs) initially dominated image generation research (Goodfellow et al. 2014; Brock, Donahue, and Simonyan 2018). Despite their success, GANs suffer from training instability, mode collapse, and difficulty in modeling diverse data distributions (Sauer, Schwarz, and Geiger 2022; Kang et al. 2023; Miyato et al. 2018). Motivated by this limitation, *diffusion-based models* (Dhariwal and Nichol 2021; Rombach et al. 2022; Han et al. 2025) and *token-based models* (Esser, Rombach, and Ommer 2021) have recently emerged as powerful alternatives. *Diffusion-based models* formulate image generation as progressive denoising process, starting from pure Gaussian noise and learning to reconstruct images through a series of iterative refinements.

Instead of directly generating pixel values, *token-based models* first rely on vector quantization techniques (Van Den Oord, Vinyals et al. 2017) to compress continuous images into sequences of discrete tokens, and then learn to model and generate these token sequences. Two primary prediction paradigms exist: autoregressive and non-autoregressive approaches. Autoregressive models generate tokens sequentially, i.e., predicting next token based on previously generated tokens (Esser, Rombach, and Ommer 2021; Guotao et al. 2024; Liang et al. 2025).

Non-autoregressive models, i.e., masked image generation, generate multiple tokens per step, enabling parallel decoding. A representative work is MaskGIT (Chang et al. 2022), which predicts the entire token map in parallel and iteratively refines it using confidence-based masking. This results in significantly faster inference with competitive generation quality. However, it still underperforms recent well-developed diffusion-based models. Unfortunately, efforts to improve its performance remain relatively limited. Existing works focus on improving the inference sampling strategy. For example, training a critic or correctors post-hoc to correct tokens with intermediate decoding errors (Lezama et al. 2022a,b). Some works (Ni et al. 2024a,b) attempt to directly optimize training and sampling hyper-parameters to obtain better performance. Inspired by the self-guidance approach in diffusion models, SG-MGM (Hur et al. 2024) adapts the self-guidance approach to Masked Image Generation (MIG) to address the multi-modality problem (Gu et al. 2017). A recent work (Besnier et al. 2025) argues that the confidence scheduling strategy of the MaskGIT method impacts the generated images’ diversity and quality, and proposes a new sampling strategy based on Halton low-discrepancy sequence (Halton and Smith 1964), which can reveal new tokens more uniformly.

Unlike prior methods that mainly refine the decoding strategy, our approach focuses on improving model’s internal representation capacity by modeling richer semantic dependencies among tokens. We identify three types of advantageous token knowledge graphs and propose a lightweight additive–subtractive fusion mechanism to integrate prior knowledge into the existing framework effectively.

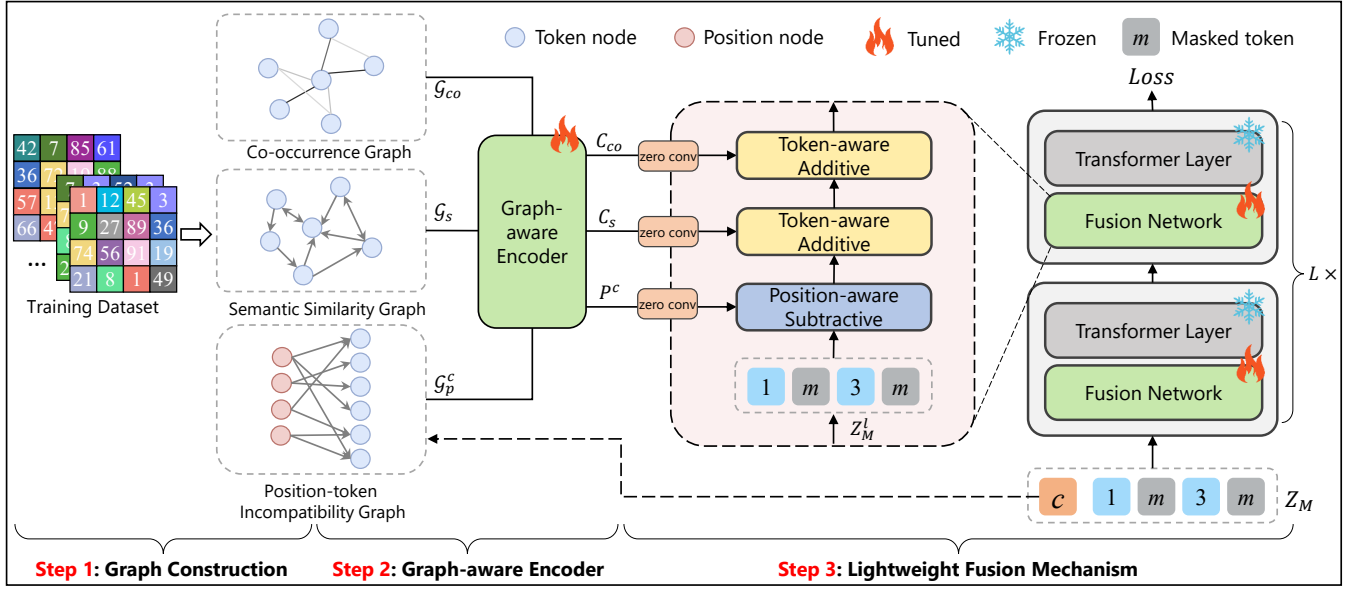


Figure 2: The illustration of our proposed KA-MIG framework. We first construct three types of prior knowledge graphs from the training dataset, *i.e.*, \mathcal{G}_{co} , \mathcal{G}_s , and \mathcal{G}_p^c . These graphs are then fed into a graph-aware encoder to learn richer token- and position-aware representations, *i.e.*, C_{co} , C_s , and P^c . Finally, the enriched prior representations are integrated into the existing MIG framework through a lightweight fusion network.

Preliminaries: MIG

Mask image generation (MIG) relies on a pre-trained VQ-VAE model (Van Den Oord, Vinyals et al. 2017) that compresses continuous images into a discrete token sequence. The VQ-VAE consists of an encoder E_{vq} , a quantizer Q with a learnable codebook \mathcal{C} , and a decoder D_{vq} . Given an image $x \in \mathbb{R}^{H \times W \times 3}$, the encoder and quantizer compress the image into a discrete token sequence: $Z = Q(E_{vq}(x))$, $Z = [z_i]_{i=1}^N$, where z_i is a specific entry in codebook, N is the length of the sequence. The decoder is used to reconstruct the original image based on Z . MIG aims to learn to generate reasonable discrete token sequences.

Training strategy. Give an external condition c , MIG first samples a binary mask matrix $M = [m_i]_{i=1}^N$. This token z_i is replaced by the special token [MASK] if $m_i = 1$, otherwise z_i will keep unchanged when $m_i = 0$. MIG, a multi-layer bidirectional transformer $p_\theta(\cdot)$, is trained to restore the masked tokens to clean tokens based on unmasked tokens. The training objective is to minimize the negative log-likelihood of the masked tokens:

$$\mathcal{L}_{\text{MIG}} = - \mathbb{E}_{Z \in \mathcal{D}} \left[\sum_{\forall i \in [1, N], m_i = 1} \log p_\theta(z_i | Z_{\overline{M}}, c) \right]. \quad (1)$$

Generation strategy. MIG starts from a fully masked token sequence and iteratively predicts all masked positions in parallel. At each step, a subset of tokens with the highest confidence scores are selected and kept, gradually completing the sequence over multiple iterations.

Proposed Method: KA-MIG

Existing works attempt to refine sampling strategies to improve generation performance, while ignoring a fundamental limitation that the model itself struggles to effectively capture semantic dependencies among visual tokens because individual tokens lack clear semantic meanings. To address this, we propose KA-MIG, a novel framework that introduces explicit knowledge of token-level semantic dependencies (*i.e.*, extracted from the training data) as priors to guide the learning of richer token representations. As illustrated in Figure 2, our KA-MIG consists of three key steps, *i.e.*, graph construction, graph-aware encoder, and lightweight fusion mechanism. Here, we elaborate on each step in order:

Step 1: Graph Construction. Due to the absence of explicit semantic meanings assigned to individual tokens in the dataset, this poses a challenge for us to construct prior knowledge graphs. To address this, we seek to uncover latent structural regularities in a data-driven manner. Specifically, we construct three types of advantageous graphs, consisting of two positive priors and one negative prior:

Co-occurrence Graph \mathcal{G}_{co} indicates that two tokens are likely to co-occur in a local range, which reflects the latent spatial-semantic correlations. Such correlations have been shown to be effective in graph learning (Nguyen et al. 2022) and recommendation systems (Han et al. 2022). Specifically, we construct a weighted, undirected token-to-token graph \mathcal{G}_{co} . Each edge \mathcal{E}_{ij}^{co} between token i and token j records how often the two tokens co-occur in a first-order local neighborhood (*i.e.*, vertical, horizontal, and diagonal directions) within the training data. To reduce noise, we prune the graph by removing low-frequency edges.



Figure 3: Visualization of semantic similarity token via reconstruction. Peak Signal-to-noise Ratio (PSNR) (Fardo et al. 2016) measures the pixel-level similarity between two images. The red box indicates the location of token (1013). We compare three versions of the same image: (a) the original reconstruction, (b) replacing token (1013) with its most similar token (463) based on our similarity graph \mathcal{G}_s , and (c) replacing it with the least similar token (149). The perceptual fidelity of (b) supports the semantic closeness, while (c) introduces noticeable degradation.

Semantic Similarity Graph \mathcal{G}_s indicates that two tokens express similar semantics in the context of image synthesis, which is akin to the definition of synonyms in natural language. Its advantage is that it encourages the model to capture interchangeable token patterns and improve robustness. However, constructing this graph is very challenging due to the absence of explicit semantic meanings for individual tokens. To capture such patterns, we hypothesize that tokens with similar positional distributions across a wide range of images are likely to convey similar semantics. Based on this, for each token, we construct a position distribution vector of length N , where each entry represents the frequency of the token appearing at a specific spatial position in the training images. We then compute the Jensen-Shannon (JS) divergence between position distribution vectors to quantify token similarity. For each token, we retain the top two most similar tokens to form a directed token-to-token graph \mathcal{G}_s . To validate our approach, as illustrated in Figure 3, we replace token (1013) with its most similar token (463) and its least similar token (149), and input the modified token maps into the decoder for image reconstruction. The results show that replacing with token (463) yields no perceptible difference from the original image, whereas token (149) significantly degrades image quality.

Position-token Incompatibility Graph \mathcal{G}_p serves as a negative prior, by identifying tokens that are incompatible with specific spatial locations under particular class labels. For example, in the “airplane” class, tokens representing ground or grass textures rarely appear in the upper-middle region of the image, as the airplane body or sky typically occupies that area. Incorporating such a negative prior helps the model avoid implausible compositions, resulting in more semantically consistent and spatially coherent image synthesis. Specifically, for each class c , we construct a directed position-to-token graph \mathcal{G}_p^c , where each entry $\mathcal{E}_{ij}^{pc} = 1$ denotes that token j is incompatible with spatial position i . To construct this graph, we scan all training images of class c and record tokens that never appear at each spatial location.

Step 2: Graph-aware Encoder After obtaining the prior

graphs described above, our goal is to learn richer representations from them for injection into the generation process, thereby enhancing image quality. To achieve this, we design two distinct encoders to process positive and negative priors. For the positive priors \mathcal{G}_{co} and \mathcal{G}_s , the goal is to learn global token representations that reflect semantic dependencies among tokens. To this end, we employ two graph convolutional networks $f_{\theta_{co}}$ and f_{θ_s} , to extract global token representations, respectively:

$$\begin{aligned} C_{co} &= f_{\theta_{co}}(\mathcal{G}_{co}, C), \\ C_s &= f_{\theta_s}(\mathcal{G}_s, C), \end{aligned} \quad (2)$$

where C is the codebook embedding.

For the negative prior, the objective is to learn position-specific features that reflect which tokens are unlikely to occur at specific position. Specifically, given class c , for each position p_i^c , we collect the set of incompatible tokens $\mathcal{I}_{i,j}$ from \mathcal{G}_p^c , and compute its representation by averaging the corresponding token embeddings:

$$p_i^c = \frac{1}{|\mathcal{I}_{i,j}|} \sum_{t \in \mathcal{I}_{i,j}} C_t W, \quad (3)$$

where C_t is the embedding of token t in the codebook C , W is a learnable weight matrix. This yields a position embedding $P^c \in \mathbb{R}^{N \times d}$ that encodes negative spatial constraints.

Step 3: Lightweight Fusion Mechanism Considering that existing MIG models have already captured rich knowledge from large-scale training, it is crucial to carefully integrate additional prior information so as not to interfere with the learned representations. Motivated by this, a lightweight fusion mechanism is proposed to solve this, as shown in Figure 2. Specifically, given class c , input masked sequence Z_M , and a pre-trained MIG, a fusion network is inserted before the input of each transformer layer of MIG. For the positive priors C_{co} and C_s , we enhance the unmasked token representations $Z_{\overline{M}}$ by additive operations at l -th layer:

$$Z_M^l = Z_{\overline{M}}^l + f_{pos}^l(C_{co}[Z_{\overline{M}}]) + f_{pos}^l(C_s[Z_{\overline{M}}]), \quad (4)$$

where $C_{co}[Z_{\overline{M}}]$ and $C_s[Z_{\overline{M}}]$ denote the unmask token embedding retrieved from the global prior codebook embeddings, $f_{pos}(\cdot)$ is a zero convolution mapping function (Zhang, Rao, and Agrawala 2023).

In contrast, for the negative prior P^c , we suppress the token features at each position by subtraction:

$$Z_M^l = Z_{\overline{M}}^l - \alpha f_{neg}^l(P^c), \quad (5)$$

where $f_{neg}(\cdot)$ is also implemented as a zero convolution function, α is a hyperparameter.

Training and Generation

Following the above process, we adopt Equation 1 as the training objective to optimize our model. Since the prior knowledge is learned independently of the MIG framework, the corresponding representations can be precomputed and stored. During inference, only lightweight additive and subtractive operations are performed, thereby incurring minimal computational overhead. More detailed analysis can be found in Section Ablation Study. This makes our method both practical for real-world deployment and fully compatible with existing MIG frameworks.




Figure 4: Visualizations for sampling images on ImageNet-256 using selected classes (270: White wolf, and 96: Toucan).

Model	Type	ImageNet 256×256					
		#Params	Steps	FID↓	IS↑	Prec↑	Rec↑
BigGAN-deep (Brock et al. 2018)	GANs	112M	1	6.95	224.5	0.89	0.38
GigaGAN (Kang et al. 2023)	GANs	569M	1	3.45	225.5	0.84	0.61
ADM (Dhariwal and Nichol 2021)	Diff.	554M	250	10.94	101.0	0.69	0.63
CDM (Ho et al. 2022)	Diff.	—	8100	4.88	158.7	—	—
LDM-4 (Rombach et al. 2022)	Diff.	400M	250	3.60	247.7	—	—
VQVAE-2 (Razavi et al., 2019)	AR	13.5B	5120	31.11	45	0.36	0.57
VQGAN (Esser et al. 2021)	AR	227M	256	18.65	80.4	0.78	0.26
LlamaGen-XI (Sun et al. 2024)	AR	775M	576	2.62	244.1	0.80	0.57
LlamaGen-XXI (Sun et al. 2024)	AR	1.4B	576	2.34	253.9	0.80	0.59
VAR-d20 (Tian et al. 2024)	AR	600M	10	2.57	302.6	0.83	0.56
VAR-d30 (Tian et al. 2024)	AR	2.0B	10	1.92	323.1	0.82	0.59
Token-Critic (Lezama et al. 2022a)	MIG	422M	36	4.69	174.5	0.76	0.53
DPC-light (Lezama et al. 2022b)	MIG	—	66	4.8	249.0	0.80	0.50
DPC-full (Lezama et al. 2022b)	MIG	—	180	4.45	244.8	0.78	0.52
MaskGIT-SAG (Hur et al. 2024)	MIG	—	12	3.35	259.7	0.81	0.52
MaskGIT-Halton (Besnier et al. 2025)	MIG	705M	32	3.74	279.5	0.81	0.60
MaskGIT (Chang et al. 2022)	MIG	227M	8	6.18	182.1	0.80	0.52
MaskGIT-KA (Ours)	MIG	245M	8	5.69	170.2	0.81	0.50
AutoNAT (Ni et al. 2024a)	MIG	194M	8	2.68	278.8	—	—
AutoNAT-KA (Ours)	MIG	211M	8	2.45	274.1	0.82	0.56
TiTok-b64 (Yu et al. 2024)	MIG	177M	8	2.48	214.7	—	—
TiTok-b64-KA (Ours)	MIG	194M	8	2.40	217.0	0.78	0.60
TiTok-s128 (Yu et al. 2024)	MIG	177M	64	1.97	281.8	—	—
TiTok-s128-KA (Ours)	MIG	194M	64	1.90	271.9	0.78	0.61

Table 1: Class-conditional image generation on ImageNet-256. Our method is marked with gray. We calculate FID-50K following (Dhariwal and Nichol 2021). The best results are highlighted in bold. Diff: diffusion, AR: autoregressive.

Experiments

Experimental Settings

Benchmarks and Datasets. As our method is model-agnostic, allowing it to be applied to different MIG architectures. Then, we choose MaskGIT (Chang et al. 2022), AutoNAT (Ni et al. 2024a), and recently state-of-the-art work TiTok (Yu et al. 2024) as our backbone networks. We evaluate the effectiveness of our method in class conditional generation using the ImageNet benchmark (Deng et al. 2009). Due to AutoNAT and TiTok not having publicly available checkpoints on 512×512 , we only evaluated them at 256×256 .

Evaluation Metrics. Following previous works (Dhariwal and Nichol 2021; Chang et al. 2022; Ni et al. 2024a; Yu et al. 2024; Hur et al. 2024), we adopt the standard Fréchet Inception Distance (FID) (Heusel and Ramsauer 2017), In-

ception Score (IS) (Salimans et al. 2016), Precision, and Recall (Kynkäänniemi et al. 2019) to evaluate the trade-off between sample fidelity and diversity.

Implementation Details. Our training and generation configurations follow the backbone model with minor adjustments to batch sizes, learning rates, and weight regularization to obtain better performance. All knowledge graphs are normalized. The $f_{\theta_{c.o}}$ and f_{θ_s} are 3-layer GCNs with the ReLU activation function. We freeze the backbone network and only fine-tune the classification layer and the additional parameters introduced by our method. More detailed experimental settings can be found in the supplementary material.

Discussion of Results

ImageNet-256 In Table 1, we present a comparison of our method against the state-of-the-art generative models and

	Step	FID↓	CLIP-Score↑
MaskGen (Kim et al. 2025)	16	22.27	25.58
Ours	16	21.01	26.10

Table 2: Text-to-image generation on COCO-30K.

the backbone models on ImageNet-256. The comparison models involve GAN-based, diffusion-based models, recently proposed state-of-the-art autoregressive models, and models with improved MIG sampling strategies. From the results, we can obtain several key conclusions: 1) Compared with the backbone models, our method consistently enhances generation performance, demonstrating both its effectiveness and versatility. This improvement is attributed to the enriched internal representations brought by incorporating prior semantic knowledge graphs, which help the model better capture inter-token dependencies; 2) Compared with MaskGIT-KA and AutoNAT-KA, the performance improvement of TiTok-KA is relatively marginal. This result is expected and reasonable. On the one hand, TiTok uses a significantly shorter token sequence (64 or 128 tokens) than MaskGIT and AutoNAT (256 tokens). Longer sequences entail more complex token interactions. Then, integrating such prior knowledge is helpful to improve the modeling ability in this aspect. This explains the more substantial improvements observed on MaskGIT and AutoNAT, supporting the validity of our core motivation. On the other hand, despite TiTok’s compact design, our method still yields consistent performance gains, further demonstrating the generalizability and robustness of the proposed approach. 3) Compared to GAN, Diff, and AR generative models, our TiTok-s128-KA achieves the best FID score, demonstrating the superiority of the proposed approach. Moreover, the AutoNAT model achieves more competitive performance after integrating our method. For example, in FID, AutoNAT-KA achieves a score of 2.45, outperforming LlamaGen-XL (2.62), VAR-d20 (2.57), and TiTok-b64 (2.48). These results further validate the effectiveness and general applicability of our method; 4) Our method introduces only approximately 20 million additional parameters, demonstrating its lightweight design and efficiency. Finally, we also provide a qualitative comparison of image generation in Figure 4.

Text-to-image generation on MS-COCO. We further evaluate the effectiveness of our method in the text-to-image generation scenario, using the MS-COCO (Lin and Maire 2014) benchmark. We select MaskGen (Kim et al. 2025) as our backbone. Since the \mathcal{G}_p is customized to the class condition, we only incorporate \mathcal{G}_{co} and \mathcal{G}_s into MaskGen. The results are summarized in Table 2, which demonstrates the effectiveness and versatility of our method.

Ablation Study

Are all three prior knowledge graphs both effective? To investigate this, we conduct extensive ablation studies to evaluate the individual and combined impact of the three graphs. We use AutoNAT as the backbone model, and the results are presented in Table 3. The results lead to several

		FID↓	IS↑
(i)	Baseline (AutoNAT)	2.68	278.8
(ii)	+ \mathcal{G}_s	2.49	279.6
(iii)	+ \mathcal{G}_p	2.51	285.6
(iv)	+ \mathcal{G}_{co}	2.51	282.1
(v)	+ $\mathcal{G}_s + \mathcal{G}_p$	2.46	279.9
(vi)	+ $\mathcal{G}_{co} + \mathcal{G}_p$	2.46	280.7
(vii)	+ $\mathcal{G}_{co} + \mathcal{G}_s$	2.48	277.4
(viii)	+ $\mathcal{G}_{co} + \mathcal{G}_s + \mathcal{G}_p$	2.45	274.1

Table 3: Ablation study of three graphs on ImageNet-256. “+ \mathcal{G}_{co} ” denotes introducing co-occurrence graph.

key observations: First, each graph plays a crucial role in improving performance, which is reasonable because each captures semantic dependencies among tokens from different perspectives, thereby validating the rationale behind our design. Second, by comparing (ii) to (iv), we observe that \mathcal{G}_s leads to the most significant performance improvement. This can be attributed to its ability to guide the model toward learning interchangeable token patterns, which helps enhance the robustness and diversity of the generated results. Third, compared to the use of individual graphs in (ii) to (iv), combining any two graphs (*i.e.*, (v) to (vii)) leads to further performance gains. This indicates that the prior graphs are highly complementary, contributing to richer and more diverse representations. Finally, comparing the results of (viii) with those of (i) to (vii) shows that incorporating all three graphs yields the best overall performance, further demonstrating the effectiveness of our proposed approach.

Efficiency Analysis. As described above, our method is independent of the MIG framework, which means that graph features can be precomputed to reduce computational overhead. To demonstrate this, we conduct a detailed analysis in Table 4 to evaluate the impact of different storage strategies on model parameters and TFLOPs for the three graphs. We compute the TFLOPs of the entire generation process (excluding the VQ decoder) using an NVIDIA RTX 4090 GPU. “Computed Online” denotes that the prior graph features are computed dynamically during inference, whereas “Precomputed” indicates that the features are calculated offline prior to inference and cached for reuse. Observed from the results, compared to online computation, precomputing the \mathcal{G}_{co} and \mathcal{G}_s introduces minimal additional parameters (only +0.79M each) and virtually no extra TFLOPs, making them highly suitable for pre-inference caching. In contrast, \mathcal{G}_p introduces substantial storage overhead (+196M) due to its class-conditional nature—each class must store its own graph. To mitigate this, we compute \mathcal{G}_p online during inference, which incurs only a minor computational cost (+0.06 TFLOPs) while significantly reducing memory usage. This hybrid configuration (marked with ✓ in the table) demonstrates that our approach is not only effective but also practical for real-world deployment, offering a favorable trade-off between performance and efficiency.

How is the token embedding space learned from our semantic similarity graph? To examine whether our method effectively captures semantic similarity relationships, we visualize the token embeddings learned from \mathcal{G}_s using t-SNE,

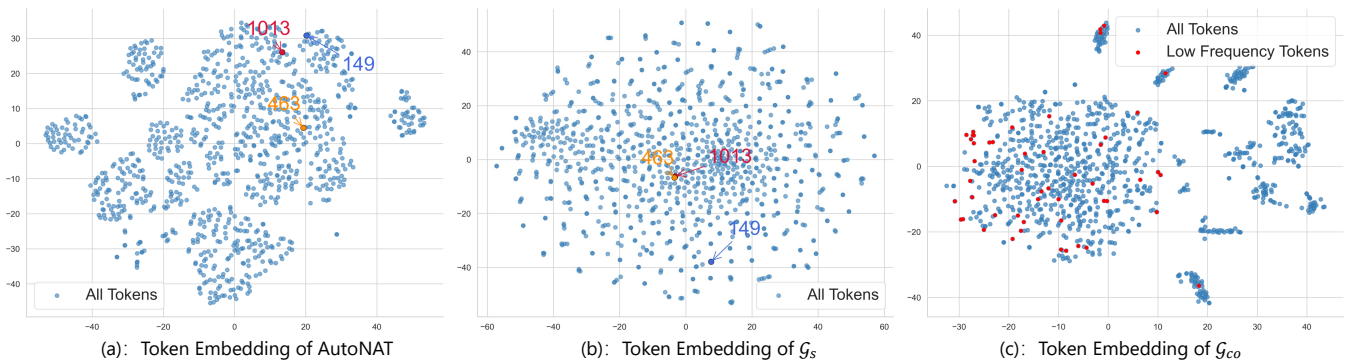


Figure 5: Visualization of token embeddings from (a) AutoNAT, (b) semantic similarity graph, and (c) co-occurrence graph.



Figure 6: Visualizations for text-to-image.

as shown in Figure 5 (b). The results show that many nodes form small clusters of two or three tokens, which aligns with our graph construction strategy, where each node connects to its top-2 most similar neighbors. More interesting, we observe that token (1013) is very similar to token (463) but is far from token (149), which is consistent with the example visualization in Figure 3. These observations demonstrate that our graph effectively captures fine-grained semantic relationships between tokens, further validating the rationale and effectiveness of introducing the \mathcal{G}_s .

How is the token embedding space learned from our co-occurrence graph? Following the above setting, we visualize the token embedding learned from \mathcal{G}_{co} in Figure 5 (c). The visualization shows that tokens form multiple coherent clusters, indicating that tokens with strong co-occurrence relationships are mapped close to each other. This confirms that the graph effectively encodes meaningful semantic locality. Furthermore, low-frequency tokens (in red) are not isolated but embedded among semantically similar neighbors, demonstrating that rare tokens enable the learning of meaningful semantics through contextual information.

Is \mathcal{G}_p effective in reducing positional bias? We conduct a token prediction experiment to demonstrate that \mathcal{G}_p effectively improves the ranking of ground-truth tokens. We randomly sample 10,000 images, apply random masking, and input them into the model to predict the corresponding ground-truth tokens. The results are presented in Ta-

Storage Type	Computed Online			Precomputed		
Prior Graph	\mathcal{G}_{co}	\mathcal{G}_s	\mathcal{G}_p	\mathcal{G}_{co}	\mathcal{G}_s	\mathcal{G}_p
Extra Params	+16M	+16M	+15M	+0.79M	+0.79M	+196M
Extra TFLOPs	+0.09	+0.09	+0.06	~ 0	~ 0	~ 0
Selected			✓	✓	✓	

Table 4: Comparison of different storage and computation strategies for the three proposed prior graphs. Our selected configuration (marked with ✓) achieves a favorable trade-off by precomputing lightweight graphs (\mathcal{G}_{co} and \mathcal{G}_s) to reduce runtime cost, while computing the heavier graph (\mathcal{G}_p) online to avoid excessive memory consumption.

	Mask Ratio	NDCG@100↑
AutoNAT	0.3	0.0282
+ \mathcal{G}_p	0.3	0.0284
AutoNAT	0.5	0.0288
+ \mathcal{G}_p	0.5	0.0289

Table 5: Comparison of NDCG@100 scores between the baseline model (AutoNAT) and our method with \mathcal{G}_p under different mask ratios. The results show that \mathcal{G}_p improves the ranking of ground-truth tokens.

ble 5, which shows that \mathcal{G}_p improves the ranking of ground-truth tokens. This aligns with the design motivation of \mathcal{G}_p : by identifying and suppressing tokens that are incompatible with specific spatial positions under a given class label, the model is better guided toward ranking true tokens higher.

Conclusion

In this work, we presented a novel Knowledge-Augmented Masked Image Generation framework that leverages explicitly constructed token-level semantic prior graphs to learn richer representations for improving performance. In particular, we uncover three types of advantageous token knowledge graphs, consisting of two positive priors and one negative prior: co-occurrence, semantic similarity, and position-token incompatibility graph. We also design a lightweight additive-subtractive fusion mechanism to integrate these priors into the existing framework effectively. Extensive experiments demonstrate the effectiveness and versatility of our approach across multiple backbone models.

Acknowledgements

This work was supported by the Guangdong Basic and Applied Basic Research Foundation under Grant No. 2025A1515011674, the Shenzhen Peacock Program under Grant No. ZX20230597, NSFC under Grant No. 62502120, the National Nature Science Foundation of China No. 62272130, Shenzhen Key Laboratory of Internet Information Collaboration, Harbin Institute of Technology (Shenzhen), Shenzhen, 518055, China, KCXFZ20240903093006009, and SYSPG20241211173609009. It was also supported by the Major Key Project of Peng Cheng Laboratory.

References

- Besnier, V.; Chen, M.; Hurych, D.; Valle, E.; and Cord, M. 2025. Halton scheduler for masked generative image transformer. *arXiv preprint arXiv:2503.17076*.
- Brock, A.; Donahue, J.; and Simonyan, K. 2018. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.
- Chang, H.; Zhang, H.; Jiang, L.; Liu, C.; and Freeman, W. T. 2022. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11315–11325.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *CVPR*, 12873–12883.
- Fardo, F. A.; Conforto, V. H.; De Oliveira, F. C.; and Rodrigues, P. S. 2016. A formal evaluation of PSNR as quality measurement parameter for image segmentation algorithms. *arXiv preprint arXiv:1605.07116*.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gu, J.; Bradbury, J.; Xiong, C.; Li, V. O.; and Socher, R. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Guotao, L.; Zhang, B.; Wang, Y.; Ye, Y.; Li, X.; and Chuyao, L. 2024. LG-VQ: Language-Guided Codebook Learning. *Advances in Neural Information Processing Systems*, 37: 139700–139724.
- Halton, J.; and Smith, G. 1964. Radical inverse quasi-random point sequence, algorithm 247. *Commun. ACM*, 7(12): 701.
- Han, Q.; Zhang, C.; Chen, R.; Lai, R.; Song, H.; and Li, L. 2022. Multi-faceted global item relation learning for session-based recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 1705–1715.
- Han, Z.; Zhang, B.; Zhang, L.; Feng, S.; Lin, K.; Liang, G.; Ye, Y.; et al. 2025. AsyncDSB: Schedule-Asynchronous Diffusion Schrödinger Bridge for Image Inpainting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3374–3382.
- Heusel, M.; and Ramsauer, H. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30.
- Ho, J.; Saharia, C.; Chan, W.; Fleet, D. J.; Norouzi, M.; and Salimans, T. 2022. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47): 1–33.
- Hur, J.; Lee, D.; Han, G.; Choi, J.; Jeon, Y.; and Kim, J. 2024. Unlocking the capabilities of masked generative models for image synthesis via self-guidance. *Advances in Neural Information Processing Systems*, 37: 130977–130999.
- Kang, M.; Zhu, J.-Y.; Zhang, R.; Park, J.; Shechtman, E.; Paris, S.; and Park, T. 2023. Scaling up gans for text-to-image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10124–10134.
- Kim, D.; He, J.; Yu, Q.; Yang, C.; Shen, X.; Kwak, S.; and Chen, L.-C. 2025. Democratizing text-to-image masked generative models with compact text-aware one-dimensional tokens. *arXiv preprint arXiv:2501.07730*.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. *Advances in neural information processing systems*, 32.
- Lezama, J.; Chang, H.; Jiang, L.; and Essa, I. 2022a. Improved masked image generation with token-critic. In *European Conference on Computer Vision*, 70–86. Springer.
- Lezama, J.; Salimans, T.; Jiang, L.; Chang, H.; Ho, J.; and Essa, I. 2022b. Discrete predictor-corrector diffusion models for image synthesis. In *The Eleventh International Conference on Learning Representations*.
- Liang, G.; Zhang, B.; Wen, Z.; Zhao, J.; Ye, Y.; Ye, K.; and He, Y. 2025. Towards Improved Text-Aligned Codebook Learning: Multi-Hierarchical Codebook-Text Alignment with Long Text. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 4060–4069.
- Lin, T.-Y.; and Maire, M. 2014. Microsoft coco: Common objects in context. In *ECCV*, 740–755. Springer.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*.
- Nguyen, D. Q.; Tong, V.; Phung, D.; and Nguyen, D. Q. 2022. Node co-occurrence based graph neural networks for knowledge graph link prediction. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 1589–1592.
- Ni, Z.; Wang, Y.; Zhou, R.; Guo, J.; Hu, J.; Liu, Z.; Song, S.; Yao, Y.; and Huang, G. 2024a. Revisiting non-autoregressive transformers for efficient image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7007–7016.

Ni, Z.; Wang, Y.; Zhou, R.; Lu, R.; Guo, J.; Hu, J.; Liu, Z.; Yao, Y.; and Huang, G. 2024b. Adanat: Exploring adaptive policy for token-based image generation. In *European Conference on Computer Vision*, 302–319. Springer.

Razavi, A.; Van den Oord, A.; and Vinyals, O. 2019. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. *Advances in neural information processing systems*, 29.

Sauer, A.; Schwarz, K.; and Geiger, A. 2022. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 conference proceedings*, 1–10.

Sun, P.; Jiang, Y.; Chen, S.; Zhang, S.; Peng, B.; Luo, P.; and Yuan, Z. 2024. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*.

Tian, K.; Jiang, Y.; Yuan, Z.; Peng, B.; and Wang, L. 2024. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems*, 37: 84839–84865.

Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Yu, Q.; Weber, M.; Deng, X.; Shen, X.; Cremers, D.; and Chen, L.-C. 2024. An image is worth 32 tokens for reconstruction and generation. *Advances in Neural Information Processing Systems*, 37: 128940–128966.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.