

dMLLM-TTS: Self-Verified and Efficient Test-Time Scaling for Diffusion Multi-Modal Large Language Models

Yi Xin^{1, 2, 3*}, Siqi Luo^{3, 4*}, Tianxiang Xu⁵, Qi Qin³, Haoxing Chen², Kaiwen Zhu^{3, 4}, Zhiwei Zhang¹,
 Yangfan He¹, Rongchao Zhang⁵, Jinbin Bai¹, Shuo Cao³, Bin Fu³, Junjun He³, Yihao Liu³,

Yuewen Cao^{3✉}, Xiaohong Liu^{1, 4✉}

¹Shanghai Innovation Institute, ²Nanjing University, ³Shanghai AI Lab,

⁴Shanghai Jiao Tong University, ⁵Peking University



Figure 1. **dMLLM-TTS**: We present the generative effects and performance improvements achieved by applying Test-Time Scaling (TTS) to dMLLMs. Images generated with TTS exhibit higher quality and stronger prompt alignment than those generated without TTS.

Abstract

Diffusion Multi-modal Large Language Models (dMLLMs) have recently emerged as a novel architecture unifying image generation and understanding. However, developing effective and efficient Test-Time Scaling (TTS) methods to unlock their full generative potential remains an underexplored challenge. To address this, we propose dMLLM-TTS, a novel framework operating on two complementary scaling axes: (1) trajectory exploration scaling to enhance the diversity of generated hypotheses, and (2) iterative refine-

ment scaling for stable generation. Conventional TTS approaches typically perform linear search across these two dimensions, incurring substantial computational costs of $\mathcal{O}(NT)$ and requiring an external verifier for best-of- N selection. To overcome these limitations, we propose two innovations. First, we design an efficient hierarchical search algorithm with $\mathcal{O}(N + T)$ complexity that adaptively expands and prunes sampling trajectories. Second, we introduce a self-verified feedback mechanism that leverages the dMLLMs’ intrinsic image understanding capabilities to assess text-image alignment, eliminating the need for exter-

nal verifier. Extensive experiments on the GenEval benchmark across three representative dMLLMs (e.g., Lumina-DiMOO, MMaDA, Muddit) show that our framework substantially improves generation quality while achieving up to 6× greater efficiency than linear search. Codes are available at: <https://github.com/Alpha-VLLM/Lumina-DiMOO>.

1. Introduction

Recent advances in text-to-image (T2I) generation, driven by powerful diffusion models [2, 13, 21, 22] and autoregressive models [3, 33], have achieved significant breakthroughs. This success is largely attributed to the scalability of training resources, as evidenced by consistent performance improvements with increased training compute, model parameters, and dataset sizes. However, this training-time scaling paradigm is encountering diminishing returns, constrained by exorbitant computational costs and a growing scarcity of high-quality data. These challenges have sparked interest in test-time scaling (TTS), a more efficient strategy that enhances the capabilities of pretrained T2I generative models by allocating additional computational resources during inference, which has also yielded highly satisfactory results.

Most recently, Diffusion Multi-Modal Large Language Models (dMLLMs) [25, 32, 36] have emerged, representing a fundamental architectural evolution. dMLLMs feature a unified architecture that natively integrates both image generation and understanding within a discrete diffusion modeling. This development has motivated researchers to investigate test-time scaling strategies specifically for dMLLMs.

In this paper, we introduce dMLLM-TTS, the first test-time scaling framework for dMLLMs that harmonizes the scaling strategy, verification mechanism, and search algorithm into a single, elegant inference process. Drawing inspiration from TTS in diffusion models [18, 26, 30, 40], we conceptualize test-time scaling in dMLLMs along two complementary axes: (1) **Trajectory Exploration Scaling**, which broadens the hypothesis space by generating N diverse initial samples to increase the likelihood of finding images that match the text prompt, and (2) **Iterative Refinement Scaling**, which deepens the generation process with T refinement steps per trajectory to enhance generative stability and final image quality. Combining these two dimensions can significantly enhance the generative capabilities of dMLLMs.

Furthermore, following the experience of TTS in diffusion models, a linear search is typically performed along these two dimensions and an external Vision-Language Model (VLM) [9, 12, 30] is employed to verify and score the final candidates, selecting the one that best matches the prompt. However, this approach has a sampling complexity of $\mathcal{O}(NT)$, and deploying an external VLM incurs addi-

tional overhead. This leads us to two fundamental questions: (1) *Can a dMLLM verify its own generated images, thereby eliminating the need for an external model?* and (2) *Can we design an efficient search algorithm that adaptively allocates compute to the most promising generative trajectories?*

This paper answers both questions affirmatively. We introduce a Self-Verified Feedback (SVF) mechanism that repurposes the dMLLM’s intrinsic multimodal understanding to assess text–image alignment. Specifically, we frame this as a question-answering task, querying the model on whether a given text prompt accurately describes the generated image. The probability of a “yes” response is then used as the alignment score. Guided by this feedback, we design a Hierarchical Trajectory Search (HTS) algorithm. HTS implements a coarse-to-fine strategy that first broadly explores diverse structural hypotheses, then prunes low-potential trajectories, and finally refines the surviving high-potential trajectories for enhanced detail. The SVF score is instrumental in both pruning unpromising trajectories and selecting the best sample from the final candidates. This intelligent allocation reduces the search complexity to a near-linear $\mathcal{O}(N + T)$.

We conduct extensive experiments using the GenEval [7] benchmark across three representative dMLLMs, (e.g., Lumina-DiMOO [32], MMaDA [36], Muddit [25]), consistently observing significant performance improvements. Furthermore, our method surpasses the linear search TTS baseline, offering higher efficiency and delivering superior final outputs. These results underscore the potential of dMLLM-TTS to enhance image generation quality without the need for additional large-scale training.

Our main contributions are summarized as follows:

- **Pioneering TTS Framework.** We establish the first test-time scaling framework for dMLLMs that integrates scaling strategy, verification, and search algorithm.
- **Novel Verifier.** We introduce a self-verification mechanism that leverages the model’s inherent image understanding to internally evaluate generative outcomes, eliminating external verifiers.
- **Efficient Search Algorithm.** We present the Hierarchical Trajectory Search, which optimizes efficiency, achieving $\mathcal{O}(N + T)$ complexity, outperforming conventional linear search baseline with $\mathcal{O}(NT)$ complexity.
- **Superior Performance.** The proposed TTS framework elevates dMLLMs to match state-of-the-art generation models, significantly boosting image quality.

2. Related Work

Diffusion Multi-Modal Large Language Models (dMLLMs). Diffusion modeling has recently progressed from continuous visual generation to discrete multi-modal generation. Early discrete diffusion methods [1, 17] established




Figure 2. **Visualization of the image generation process in dMLLMs.** The first row shows the input latent masks at each step, and the second row depicts the corresponding outputs. Sampling begins with fully masked tokens (gray) and gradually fills the discrete multimodal token space with increasingly confident predictions (blue).

the foundation for token-level denoising in discrete spaces. Building on these formulations, diffusion large language models (dLLMs) [8, 19, 38] have demonstrated that language generation can be reimagined as an iterative parallel denoising process instead of following the traditional autoregressive approach. Extending this paradigm to multi-modal domains, LLaDA-V [37] introduced visual instruction tuning, showcasing the potential of diffusion-based language models to seamlessly connect textual and visual modalities. Subsequent research further explores this direction toward unified diffusion multi-modal large language models (dMLLMs) [25, 32, 36], which can *perform both multi-modal generation and understanding*.

Test-Time Scaling for Image Generation. Prior to the emergence of dMLLMs, diffusion models [2, 6, 22, 29, 30, 39] and autoregressive models [3, 16, 27, 33, 34] dominated as the primary paradigms for image generation. Recent studies on these architectures have offered valuable insights into the principles of TTS. For diffusion models, scaling typically expands the continuous noise space [18, 26, 30] or increases the number of sampling steps [18, 26] to improve quality and diversity, with further progress driven by reflection-based optimization [15, 40]. Autoregressive models, on the other hand, allow for a distinct type of scaling by directly operating over discrete token sequences [4, 33]. Previous TTS approaches commonly rely on external verifiers [9, 11, 35] to evaluate generated images, yet such designs often require additional VLM deployments. This motivates exploring TTS within dMLLMs, which possess *unified multimodal understanding and generation capabilities that naturally support in-loop self-verification*.

3. Methodology

3.1. Preliminary

Image Generation Process of dMLLMs. dMLLMs generate images via a parallel, diffusion-based process that it-

eratively denoises a fully masked sequence into the target output, as shown in Figure 2. Let $\mathcal{Y}$ be the token vocabulary and $[\text{Mask}] \in \mathcal{Y}$ the special mask token. Given a text prompt C , the dMLLMs generate an image token sequence $Z = (z_1, z_2, \dots, z_M)$ (a total of M tokens) through T discrete denoising steps, indexed by $t = 1$ up to T . The initial state Z_0 is a fully masked sequence:

$$Z_0 = (\underbrace{[\text{Mask}], [\text{Mask}], \dots, [\text{Mask}], [\text{Mask}]}_M). \quad (1)$$

At each step t , dMLLMs predicts the tokens for all the masked locations:

$$Z_t = (z_1^t, \dots, z_{M-1}^t, z_M^t). \quad (2)$$

Subsequently, tokens with high confidence is the final choice and low confidence are re-masked and used as input for prediction at step $t + 1$:

$$Z_t = ([\text{Mask}], z_2^t, \dots, [\text{Mask}], z_M^t). \quad (3)$$

To decode a generated token sequence Z into an image, a specific discrete tokenizer [20, 31] is essential.

3.2. How to Scale at Test Time in dMLLMs

During training, prior text-to-image (T2I) diffusion models learn a denoising vector field that defines a progressive refinement process [10, 24]. At inference time, this mechanism enables scaling through stochastic trajectory sampling and adjustable refinement sampling steps [18]. dMLLMs inherit this property within a discrete latent token space, where generation proceeds by *iteratively refining token sequences initialized from stochastic states*. Thus, test-time scaling in dMLLMs can be viewed as a dual-dimensional process of trajectory exploration and iterative refinement:

- **Trajectory Exploration Scaling.** This dimension enhances the diversity of generated hypotheses. Instead of relying on a single deterministic initialization, the model




Figure 3. **Overview of dMLLM-TTS framework.** (a) dMLLM-TTS scales compute along two axes: trajectory exploration and iterative refinement, guided by Self-Verified Feedback for text–image alignment evaluation. (b) Hierarchical Trajectory Search (HTS) performs coarse-to-fine generation by starting with broad exploration, pruning low-potential trajectories, and refining high-potential trajectories.

samples N distinct stochastic latent token sequences Z_1 as its starting states (see $t = 1$ in Figure 2):

$$Z_1 \sim p_{\text{init}}, \quad (4)$$

where p_{init} denotes the initialization prior. A larger N broadens the hypothesis space, thereby increasing the likelihood of generating images that are semantically aligned with the given text prompt.

- **Iterative Refinement Scaling.** This dimension controls the computational depth per trajectory. Given an initial state Z_1 , prompt C , the model performs $T - 1$ refinement steps:

$$Z_{t+1} = \mathcal{G}_\theta(Z_t, C, t), \quad t = 1, \dots, T - 1. \quad (5)$$

Increasing refinement steps T enhances the granularity of the denoising process, yielding more stable generation trajectories and obtaining higher quality sampled images.

While scaling the number of trajectories or refinement steps yields marked improvements in image quality and text-image alignment, it comes at a substantial computational cost. However, we observe that the potential of a sampling trajectory can be preliminarily assessed during its intermediate process, as illustrated in Figure 2 (later-stage). Thus, the key to effective test-time scaling lies not in merely increasing computation, but in *adaptively allocating it to the most promising trajectories*.

We formalize this process as an *adaptive trajectory search problem*, jointly governed by a generator $\mathcal{G}_\theta$, a verifier $\mathcal{V}$, and a search function f :

$$\text{TTS} = \langle \mathcal{G}_\theta, \mathcal{V}, f \rangle, \quad (6)$$

Here, $\mathcal{V} : \mathcal{Z} \times \mathcal{C} \rightarrow \mathbb{R}$ measures semantic alignment between generated images and the given text prompts, and f adaptively redistributes inference computation under $\mathcal{V}$'s guidance. This establish a principled foundation for scalable test-time scaling in dMLLMs.

3.3. Self-Verified Feedback

Previous test-time scaling methods often depend on external verifiers, such as CLIP [23], VILA-Judge [30], and GPT-4o [12], to assess candidate samples. However, these methods suffer from inherent drawbacks: (i) they require additional VLM deployments or commercial API calls, increasing resource consumption; and (ii) they incur substantial computational overhead due to repeated decode (converting tokens to image) and encode (transforming images to embeddings for VLMs) operations.

For dMLLMs, the model itself is the ideal verifier $\mathcal{V}$, given its inherent image generation and understanding capabilities. This lead us to introduce a **Self-Verified Feedback (SVF)** mechanism, which reuses the model to evaluate generated images via its internal understanding function, enabling efficient, in-loop assessment without external decoding or scoring models, as shown in Figure 3(a) and following template.

Formally, for a given text prompt C and an intermediate generated image token sequence Z_t , the model $\mathcal{G}_\theta$ provides binary responses (“yes” for good quality, “no” for low quality) and leverages the logit probability of “yes” as the text-image alignment score:

$$\Phi_{\text{SVF}} = \text{logit}_{\text{yes}}(\mathcal{G}_\theta(Z_t, C)). \quad (7)$$

The SVF score thus serves as a semantic measure, allowing the model to rank and select trajectories based on internally consistent criteria.

Instruction Template

<Generated Image> Is this image shows {text prompt}? Please answer “Yes” or “No” directly without explanation.

3.4. How to Efficiently Search the Optimal Sample

Conventional search strategies, such as the **Linear Trajectory Search (LTS)** baseline [30, 33], allocate equal inference compute across all trajectories and incur $\mathcal{O}(NT)$ complexity without leveraging intermediate sampling feedback, as shown in Figure 3(a). While effective, this brute-force approach is highly inefficient for dMLLMs’ *coarse-to-fine* generative hierarchy, where early stages establish the global structure, while later stages refine semantic details. Thus, it is redundant to invest compute resources in trajectories that are already flawed in the early stages.

To address this limitation and achieve a better balance between image quality and test-time computational efficiency, we propose **Hierarchical Trajectory Search (HTS)** strategy with $\mathcal{O}(N+T)$ complexity, which adaptively redistributes computation guided by self-verified feedback (SVF) and further performs *local neighborhood exploration in the vicinity of high-potential trajectories to guide subsequent refinement*. Specifically, as shown in Figure 3(b), the HTS can be divided into three smoothly evolving stage:

$$\text{HTS} \Rightarrow \begin{cases} \text{Initial Stochastic Exploration,} & t \leq T_s, \\ \text{Hierarchical Thinning,} & T_s < t \leq T_r, \\ \text{Final Refinement,} & T_r < t \leq T, \end{cases} \quad (8)$$

where T_s and T_r denote the transition steps between the three stages.

Initial Stochastic Exploration. We begin by sampling N stochastic trajectories from the initialization prior:

$$Z_1^{(i)} \sim p_{\text{init}}, \quad i = 1, \dots, N. \quad (9)$$

Each trajectory is then denoised over T_s steps to generate a coarse structural hypothesis:

$$Z_{T_s}^{(i)} = \mathcal{G}_\theta^{T_s}(Z_1^{(i)}, C). \quad (10)$$

where $\mathcal{G}_\theta^{T_s}$ denotes the T_s steps diffusion propagation operator. At this initial, high-noise stage, token representations remain highly stochastic, resulting in blurry generated images (see Figure 2). Consequently, SVF-based evaluation is

less meaningful at this stage. Focus should instead be directed towards broad structural exploration and developing a diverse array of initial hypotheses for further refinement.

Progressive Hierarchical Thinning. We define a progressive *trajectory-width decay schedule* that controls how the search space gradually narrows throughout inference:

$$W_t = \max(\lfloor Nd^{-(t-T_s)} \rfloor, K), \quad d > 1, \quad (11)$$

where K is the minimal retained set, and d is a decay coefficient determining how aggressively the pool contracts. This formulation ensures that the number of active trajectories decreases exponentially until convergence. At each refinement step t , computation is adaptively redistributed toward promising trajectories under the guidance of SVF:

1. **Scoring.** Each trajectory in the current pool $\{Z_t^{(i)}\}_{i=1}^{W_t}$ is assigned a SVF-based score Φ_{SVF} .
2. **Selection.** The top- K trajectories with the highest scores form the surviving set B_t .
3. **Branching.** Each survivor $Z_t^{(j)} \in B_t$ then generates b_t stochastic continuations from a local kernel q :

$$b_t = \left\lfloor \frac{W_{t+1}}{K} \right\rfloor, \quad Z_{t+1}^{(j,k)} \sim q(Z | Z_t^{(j)}), \quad k = 1, \dots, b_t. \quad (12)$$

As t increases, both W_t and b_t decrease geometrically, incrementally focusing computation on potential trajectories with high SVF scores. When $W_t = K$ (i.e., $b_t = 1$), branching stops, and the surviving trajectories transition into the final refinement stage.

Final Refinement. The adaptive thinning process contracts the trajectory pool to a minimum width of K at step T_r , signifying the transition from trajectory exploration to refinement. Subsequently, the K surviving trajectories are refined independently until the final diffusion step T :

$$Z_T^{(j)} = \mathcal{G}_\theta^{T-T_r}(Z_{T_r}^{(j)}, C), \quad j = 1, \dots, K. \quad (13)$$

At this stage, all computational resources are redirected to enhancing the image’s details. By concentrating its inference depth on the text-image aligned trajectories identified by SVF, the model transforms early stochastic exploration into precise, text-aligned outputs. This adaptive contraction strategy dynamically selects the most promising paths for detailed refinement, effectively merging the benefits of broad exploration and deep refinement within a single, compute-efficient inference process.


Complexity Analysis. The total forward cost of HTS can be expressed as:

$$C_{\text{HTS}} = \mathcal{O}(NT_s + \frac{N-dK}{d-1} + K(T-T_r)), \quad (14)$$

where the three terms respectively correspond to (i) early stochastic exploration over N trajectories for T_s steps, (ii)

Table 1. **Quantitative performance comparison on GenEval** across various d-MLLMs and test-time scaling settings.

	#Steps	#Trajectory	Search	Single Obj.	Two Obj.	Counting	Colors	Position	Attribute	Overall
<i>Text-to-Image Models</i>										
FLUX.1-dev [13]	-	-	-	0.99	0.81	0.75	0.80	0.21	0.48	0.67
Janus-Pro [3]	-	-	-	0.99	0.89	0.59	0.90	0.79	0.66	0.80
BAGEL [5]	-	-	-	0.99	0.94	0.81	0.88	0.64	0.63	0.82
GPT-4o [12]	-	-	-	0.99	0.92	0.85	0.92	0.75	0.61	0.84
Qwen-Image [28]	-	-	-	0.99	0.92	0.89	0.88	0.76	0.77	0.87
<i>Test-Time Scaling</i>										
Lumina-DiMOO [32]	T = 8	N = 1	-	0.96	0.85	0.70	0.85	0.72	0.62	0.78
+ Test-Time Scaling	T = 16	N = 16 (K = 4)	LTS	1.0 ^{+4.2%}	0.93 ^{+9.4%}	0.81 ^{+15.7%}	0.90 ^{+5.9%}	0.83 ^{+15.3%}	0.76 ^{+22.6%}	0.87 ^{+11.5%}
			HTS	1.0 ^{+4.2%}	0.94 ^{+10.6%}	0.84 ^{+20.0%}	0.92 ^{+8.2%}	0.86 ^{+18.9%}	0.79 ^{+27.4%}	0.89 ^{+14.1%}
	T = 32	N = 32 (K = 8)	LTS	1.0 ^{+4.2%}	0.97 ^{+14.1%}	0.89 ^{+27.1%}	0.93 ^{+9.4%}	0.87 ^{+20.8%}	0.80 ^{+29.0%}	0.91 ^{+15.4%}
			HTS	1.0 ^{+4.2%}	0.97 ^{+14.1%}	0.91 ^{+30.0%}	0.94 ^{+10.6%}	0.89 ^{+23.6%}	0.82 ^{+32.3%}	0.92 ^{+17.9%}
MMaDA [36]	T = 8	N = 1	-	0.94	0.60	0.32	0.76	0.19	0.22	0.51
+ Test-Time Scaling	T = 16	N = 16 (K = 4)	LTS	0.98 ^{+4.3%}	0.73 ^{+21.7%}	0.47 ^{+46.9%}	0.88 ^{+15.8%}	0.25 ^{+31.6%}	0.31 ^{+42.9%}	0.60 ^{+17.6%}
			HTS	0.99 ^{+5.3%}	0.75 ^{+25.0%}	0.48 ^{+50.0%}	0.90 ^{+18.4%}	0.27 ^{+42.1%}	0.32 ^{+45.5%}	0.62 ^{+21.6%}
	T = 32	N = 32 (K = 8)	LTS	0.99 ^{+5.3%}	0.77 ^{+28.3%}	0.50 ^{+56.2%}	0.91 ^{+19.7%}	0.29 ^{+52.6%}	0.36 ^{+63.6%}	0.64 ^{+25.5%}
			HTS	0.99 ^{+5.3%}	0.81 ^{+35.0%}	0.52 ^{+62.5%}	0.92 ^{+21.1%}	0.33 ^{+73.6%}	0.37 ^{+68.2%}	0.66 ^{+29.4%}
Muddit [25]	T = 8	N = 1	-	0.95	0.53	0.48	0.80	0.13	0.29	0.53
+ Test-Time Scaling	T = 16	N = 16 (K = 4)	LTS	0.98 ^{+3.2%}	0.68 ^{+28.3%}	0.59 ^{+22.9%}	0.86 ^{+7.5%}	0.22 ^{+69.2%}	0.39 ^{+34.5%}	0.62 ^{+17.0%}
			HTS	0.99 ^{+4.2%}	0.71 ^{+33.9%}	0.59 ^{+22.9%}	0.88 ^{+10.0%}	0.24 ^{+84.6%}	0.41 ^{+41.4%}	0.64 ^{+20.8%}
	T = 32	N = 32 (K = 8)	LTS	0.99 ^{+4.2%}	0.73 ^{+37.3%}	0.60 ^{+25.0%}	0.89 ^{+11.2%}	0.24 ^{+84.6%}	0.42 ^{+44.8%}	0.65 ^{+22.6%}
			HTS	0.99 ^{+4.2%}	0.76 ^{+43.4%}	0.62 ^{+29.2%}	0.88 ^{+10.0%}	0.25 ^{+92.3%}	0.44 ^{+51.7%}	0.67 ^{+26.4%}

Figure 4. **Qualitative improvement ratio in TTS performance across various text prompt complexities** examined through diverse dMLLMs on GenEval benchmark dimensions. TTS markedly enhances performance across all measured dimensions.

hierarchical thinning with geometric decay factor $d > 1$, and (iii) final refinement over the K surviving trajectories.

In practice, we start from a wide exploration with a large number of trajectories N and quickly prune to a compact set ($K \ll N$), while the early warm-up stage ($T_s \ll T$) only covers a small portion of the diffusion process to establish structural diversity. Hence the complexity simplifies to:

$$C_{\text{HTS}} \approx \mathcal{O}(N + T). \quad (15)$$

This near-linear scaling demonstrates that HTS effectively reallocates computation from early wide exploration to late fine refinement, avoiding the quadratic $\mathcal{O}(NT)$ cost of linear trajectory search (LTS). By performing coarse structural exploration under high noise and gradually contracting the trajectory pool to K candidates for fine-grained synthesis,

HTS efficiently balances search width and depth during inference.

4. Experiment

4.1. Experiment Setup

Evaluation Models. We validate our proposed test-time scaling method (*i.e.*, scaling strategy, verifier, and search algorithm) using three popular open-source pre-trained dMLLMs, including Lumina-DiMOO [32], MMaDA [36], and Muddit [25]. These models span a parameter range from 1B to 8B, which allows us to verify the generality of our method across different model scales. Due to the lack of open-source access to other notable dMLLMs, such as Llava-O [14], it is currently not feasible to conduct experiments with them.




Figure 5. **Comparison between linear and hierarchical trajectory search.** The red curve illustrates linear trajectory search, while the blue curve depicts hierarchical trajectory search, with a dashed line indicating predictions based on a geometric series decay approximation. Curve fitting shows that similar subsequent trends tend to converge towards an upper limit.

Evaluation Metrics. To quantify the performance of text-to-image (T2I) generation, we employ the GenEval [7] benchmark, which is specifically designed to assess object-centric generation using compositional prompts with varied object attributes. GenEval, consisting of 553 text prompts, is the most widely used benchmark within the TTS domain. To measure computational cost, we adopt the metric of the number of function evaluations (NFE), consistent with previous TTS studies on generative models [18, 33, 40].

Implementation Details. The baseline performance for each model is established using 8 sampling steps. To ensure a fair comparison, all images are generated at a 512 resolution with CFG=4.0. For image verification, the model performs one step to produce a “Yes” or “No” response. In addition to our proposed Self-Verified Feedback, we conducted further experiments using external verifiers, as detailed in Section 4.3. For the HTS strategy, the primary experimental setting maintains an N:K ratio of 4:1. As for T_s in the HTS, it is empirically set to $T/4$, while T_r is determined by N and K.

4.2. Main Results

TTS Consistently Delivers Stable Performance Gains Across Various d-MLLMs. We applied our dMLLM-TTS method to the Lumina-DiMOO, MMaDA, and Muddit models, yielding significant and consistent performance gains on the GenEval benchmark (Table 1). Specifically, Lumina-DiMOO’s score rose from 0.78 to 0.92 (+17.9%), MMaDA’s from 0.51 to 0.66 (+29.4%), and Muddit’s from 0.53 to 0.67 (+26.4%). This enhancement elevates their capabilities to match or even surpass SOTA text-to-image models. For example, the TTS-enhanced Lumina-DiMOO (0.92) outperforms leading models such as Qwen-Image (0.87) and GPT-4o (0.84). These findings demonstrate that dMLLM-TTS is an effective and widely applicable method for improving the generative quality of d-MLLMs.

TTS Significantly Boosts Performance Across All Dimensions. The ability to handle text prompts of varying complexity across multiple dimensions is a key challenge

of T2I generation. Our dMLLM-TTS consistently boosts performance across six dimensions, as shown in Figure 4. For common prompt sets (e.g., Single Object, Two Object, Colors), dMLLM-TTS provides notable improvements, although the gains are inherently limited by the strong baseline performance. Crucially, for complex prompt sets (e.g., Counting, Position, Attribute), dMLLM-TTS achieves substantial performance gains. This highlights its effectiveness in elevating the dMLLMs’ performance ceiling, especially on challenging compositional generation tasks.

HTS Proves More Efficient and Superior to LTS Baseline. We evaluated our Hierarchical Trajectory Search (HTS) against the Linear Trajectory Search (LTS) baseline in Figure 5. The results reveal clear advantages for our HTS approach in both efficiency and performance. HTS requires substantially less inference compute to reach equivalent scores, achieving $5\times$ faster on Lumina-DiMOO and $6\times$ faster on MMaDA and Muddit. Furthermore, beyond just accelerating the search, HTS consistently converges to a higher overall score than the baseline. These findings confirm that HTS offers a powerful combination of computational savings and superior results.

4.3. Analysis and Discussion

Effect of Trajectory Exploration Scaling. As depicted in Figure 6 (left), our analysis validates the effectiveness of Trajectory Exploration Scaling for improving GenEval scores. With a fixed step of 32, we systematically scaled the exploration trajectory from $N = 1$ to $N = 32$, observing a clear, positive impact across all evaluated models. Notably, the performance gains are inversely correlated with the models’ initial scores. The baseline models, MMaDA and Muddit, achieved the most significant improvements of +20.2% and +16.8%, respectively. Even the top-performing Lumina-DiMOO model benefited from a substantial +8.8% score increase. This evidence strongly suggests that our trajectory exploration scaling method is a robust and valuable technique for enhancing model performance




Figure 6. **Trajectory Exploration Scaling (Left) and Iterative Refinement Scaling (Right)**. Increasing the number of explored trajectories ($N = 1 \rightarrow 32$) refinement steps ($T = 8 \rightarrow 64$) consistently improves performance across all dMLLMs.




Figure 7. **Image Generation Process without (Top) and with (Bottom) dMLLM-TTS**. The baseline models produce unsatisfactory text-to-image results. However, by incorporating our TTS strategies, the generation process is significantly improved.

Table 2. **Comparison results of various verifiers.**

Verifier	Lumina-DiMOO	MMaDA	Muddit
SVF (Ours)	0.92	0.66	0.67
VILA-Judge	0.90 ↓	0.70 ↑	0.70 ↑
GPT-4o	0.95 ↑	0.71 ↑	0.74 ↑

Effect of Iterative Refinement Scaling. Figure 6 (right) demonstrates the effect of iterative refinement scaling across three models. As the number of refinement steps increases from 8 to 64, overall scores consistently improve, underscoring the effectiveness of our iterative refinement scaling strategy. However, indefinite scaling is not practical. The ideal number of refinement steps depends on the complexity of the text prompt. The optimal number of refinement steps typically depends on the complexity of the text prompt. For dMLLMs sampling at a resolution of 512×512 , up to 64 steps typically provides a strong balance between performance gains and computational efficiency.

Comparison with External Verifiers. Our analysis shows that external verifiers, particularly powerful commercial model GPT-4o, outperform our self-verified feedback approach, as reported in Table 2. This gap is primarily attributable to their superior image understanding capabilities.

ties. These results suggest that current dMLLMs still have considerable room to improve in visual comprehension.

Qualitative Examples. To intuitively illustrate the effectiveness of dMLLM-TTS, we present qualitative examples of generated images and corresponding generation processes. As shown in Figure 7, baseline models often produce a bad initial state and follow a wrong generation trajectory. In contrast, our dMLLM-TTS leads to a good initial state and maintain a perfect trajectory. This underscores the significant improvements achieved with dMLLM-TTS, resulting in more precise and faithful image generations that align with the text prompt.

5. Conclusion

In this paper, we present a novel TTS framework specifically designed for dMLLMs. By integrating Self-Verified Feedback for internal evaluation and Hierarchical Trajectory Search for adaptive inference, our framework eliminates reliance on external verifiers and optimizes computational allocation. Experiments show that dMLLM-TTS framework significantly enhances text-image alignment, especially for complex prompts, establishing an efficient and powerful paradigm for scalable inference within dMLLMs.

References

- [1] Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [2] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, et al. Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023. 2, 3
- [3] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 2, 3, 6
- [4] Zhekai Chen, Ruihang Chu, Yukang Chen, Shiwei Zhang, Yujie Wei, Yingya Zhang, and Xihui Liu. Tts-var: A test-time scaling framework for visual auto-regressive generation. *arXiv preprint arXiv:2507.18537*, 2025. 3
- [5] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 6
- [6] Peng Gao, Le Zhuo, Dongyang Liu, Ruoyi Du, Xu Luo, Longtian Qiu, Yuhang Zhang, Rongjie Huang, Shijie Geng, Renrui Zhang, et al. Lumina-t2x: Scalable flow-based large diffusion transformer for flexible resolution generation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. 3
- [7] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2024. 2, 7
- [8] Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, et al. Scaling diffusion language models via adaptation from autoregressive models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. 3
- [9] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021. 2, 3
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [11] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 3
- [12] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024. 2, 4, 6
- [13] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 2, 6
- [14] Shufan Li, Jiuxiang Gu, Kangning Liu, Zhe Lin, Zijun Wei, Aditya Grover, and Jason Kuen. Lavid-a-o: Elastic large masked diffusion models for unified multimodal understanding and generation. *arXiv preprint arXiv:2509.19244*, 2025. 6
- [15] Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Arsh Koneru, Yusuke Kato, Kazuki Kozuka, and Aditya Grover. Reflect-dit: Inference-time scaling for text-to-image diffusion transformers via in-context reflection. *arXiv preprint arXiv:2503.12271*, 2025. 3
- [16] 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. 3
- [17] Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024. 2
- [18] Nanye Ma, Shangyuan Tong, Haolin Jia, Hexiang Hu, Yu-Chuan Su, Mingda Zhang, Xuan Yang, Yandong Li, Tommi Jaakkola, Xuhui Jia, et al. Inference-time scaling for diffusion models beyond scaling denoising steps. *arXiv preprint arXiv:2501.09732*, 2025. 2, 3, 7
- [19] Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025. 3
- [20] Suraj Patil, William Berman, Robin Rombach, and Patrick von Platen. amused: An open muse reproduction. *arXiv preprint arXiv:2401.01808*, 2024. 3
- [21] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023. 2
- [22] Qi Qin, Le Zhuo, Yi Xin, Ruoyi Du, Zhen Li, Bin Fu, Yiting Lu, Jiakang Yuan, Xinyue Li, Dongyang Liu, et al. Lumina-image 2.0: A unified and efficient image generative framework. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2025. 2, 3
- [23] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021. 4
- [24] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [25] Qingyu Shi, Jinbin Bai, Zhuoran Zhao, Wenhao Chai, Kaidong Yu, Jianzong Wu, Shuangyong Song, Yunhai Tong,

- Xiangtai Li, Xuelong Li, and Shuicheng Yan. Muddit: Liberating generation beyond text-to-image with a unified discrete diffusion model, 2025. 2, 3, 6
- [26] Raghu Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, et al. A general framework for inference-time scaling and steering of diffusion models. *arXiv preprint arXiv:2501.06848*, 2025. 2, 3
- [27] 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. 3
- [28] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 6
- [29] Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, et al. Sana: Efficient high-resolution image synthesis with linear diffusion transformers. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. 3
- [30] Enze Xie, Junsong Chen, et al. Sana 1.5: Efficient scaling of training-time and inference-time compute in linear diffusion transformer. *Proceedings of the International Conference on Machine Learning (ICML)*, 2025. 2, 3, 4, 5
- [31] Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024. 3
- [32] Yi Xin, Qi Qin, Siqi Luo, Kaiwen Zhu, Juncheng Yan, Yan Tai, Jiayi Lei, Yuewen Cao, Keqi Wang, Yibin Wang, et al. Lumina-dimoo: An omni diffusion large language model for multi-modal generation and understanding. *arXiv preprint arXiv:2510.06308*, 2025. 2, 3, 6
- [33] Yi Xin, Juncheng Yan, Qi Qin, Zhen Li, Dongyang Liu, Shicheng Li, Victor Shearman & Sterling Huang, Yupeng Zhou, Renrui Zhang, Le Zhuo, et al. Lumina-mgpt 2.0: Stand-alone autoregressive image modeling. *arXiv preprint arXiv:2507.17801*, 2025. 2, 3, 5, 7
- [34] Yi Xin, Le Zhuo, Qi Qin, Siqi Luo, Yuewen Cao, Bin Fu, Yangfan He, Hongsheng Li, Guangtao Zhai, Xiaohong Liu, et al. Resurrect mask autoregressive modeling for efficient and scalable image generation. *arXiv preprint arXiv:2507.13032*, 2025. 3
- [35] Jiazhen Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 3
- [36] Ling Yang, Ye Tian, Bowen Li, Xincheng Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada: Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*, 2025. 2, 3, 6
- [37] Zebin You, Shen Nie, Xiaolu Zhang, Jun Hu, Jun Zhou, Zhiwu Lu, Ji-Rong Wen, and Chongxuan Li. LLaDA-V: Large language diffusion models with visual instruction tuning. *arXiv preprint arXiv:2505.16933*, 2025. 3
- [38] Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Hu, Jun Zhou, Jianfei Chen, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. LLaDA 1.5: Variance-reduced preference optimization for large language diffusion models. *arXiv preprint arXiv:2505.19223*, 2025. 3
- [39] Le Zhuo, Ruoyi Du, Han Xiao, Yangguang Li, Dongyang Liu, Rongjie Huang, Wenze Liu, Xiangyang Zhu, Fu-Yun Wang, Zhanyu Ma, et al. Lumina-next: Making lumina-t2x stronger and faster with next-dit. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [40] Le Zhuo, Liangbing Zhao, Sayak Paul, Yue Liao, Renrui Zhang, Yi Xin, et al. From reflection to perfection: Scaling inference-time optimization for text-to-image diffusion models via reflection tuning. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2025. 2, 3, 7