# DiT-Air: Revisiting the Efficiency of Diffusion Model Architecture Design in Text to Image Generation

Chen Chen, Rui Qian, Wenze Hu, Tsu-Jui Fu, Jialing Tong, Xinze Wang,
Lezhi Li, Bowen Zhang, Alex Schwing, Wei Liu, Yinfei Yang
Apple Inc.

{chen_chen999}@apple.com

## Abstract

*In this work, we empirically study Diffusion Transformers (DiTs) for text-to-image generation, focusing on architectural choices, text-conditioning strategies, and training protocols. We evaluate a range of DiT-based architectures—including PixArt-style and MMDiT variants—and compare them with a standard DiT variant which directly processes concatenated text and noise inputs. Surprisingly, our findings reveal that the performance of standard DiT is comparable with those specialized models, while demonstrating superior parameter-efficiency, especially when scaled up. Leveraging the layer-wise parameter sharing strategy, we achieve a further reduction of 66% in model size compared to an MMDiT architecture, with minimal performance impact. Building on an in-depth analysis of critical components such as text encoders and Variational Auto-Encoders (VAEs), we introduce DiT-Air and DiT-Air-Lite. With supervised and reward fine-tuning, DiT-Air achieves state-of-the-art performance on GenEval and T2I CompBench, while DiT-Air-Lite remains highly competitive, surpassing most existing models despite its compact size.*

## 1. Introduction

The field of text-to-image synthesis has witnessed remarkable progress, primarily attributable to the wide adoption of diffusion-based models [14, 20]. Diffusion Transformers (DiTs) [33] have emerged as a prominent architectural paradigm, combining the iterative denoising process inherent to diffusion models with the representational efficacy of transformer networks. While existing paradigm variants like PixArt-style [3–5] and MMDiT from Stable Diffusion 3 (SD3) [14] have demonstrated strong performance, key aspects of DiTs—such as the choice of architectural components, text-conditioning mechanisms, and training strategies—have not been exhaustively explored yet [6, 17].
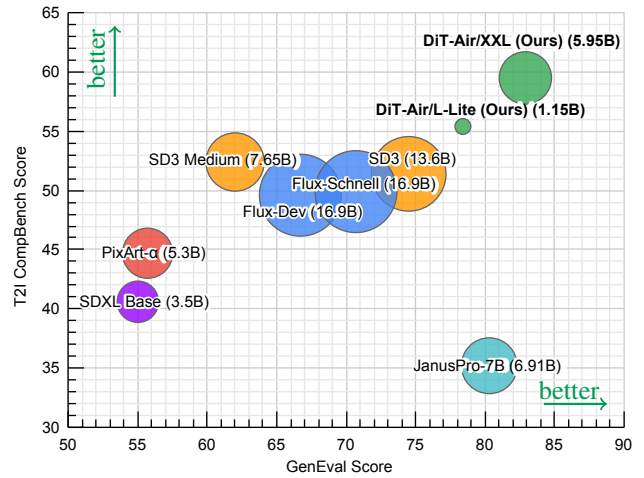


Figure 1. Comparison of text-to-image generation methods on two metrics, GenEval and T2I CompBench (higher is better for both). Despite significantly smaller model size, our proposed DiT-Air achieves state-of-the-art results. Note that, for our model, we report the full model size including text encoder and VAE. A detailed parameter breakdown is provided in Appendix G.

In this work, we conduct a comprehensive investigation into DiT design choices for text-to-image synthesis. Beginning with a comparative analysis of the vanilla DiT [33], PixArt-$\alpha$ [4], and MMDiT [14], we develop a streamlined architecture. This architecture utilizes concatenated text and noise inputs (following MMDiT) and shared AdaLN parameters (following PixArt-$\alpha$), eliminating modality-specific projections. This simplification yields substantial parameter savings (66% compared to MMDiT and 25% compared to PixArt-$\alpha$), while preserving or enhancing performance. Notably, the resulting architecture, named DiT-Air, closely resembles the original DiT, allowing us to leverage existing transformer optimizations. Inspired by parameter-sharing strategies in NLP models such as ALBERT [27], we adopt both full block-sharing and

1

Figure 2. Sample images from our proposed DiT-Air, each with the text prompt below it. See Appendix H for more examples.

attention-only sharing schemes to further push the parameter efficiency of DiT models. Our ablation studies demonstrate that attention-only sharing provides a compelling trade-off, achieving significant parameter reduction with minimal loss in text alignment and generative fidelity.

In addition to architectural innovations, DiT-Air benefits from a thorough analysis of text-conditioning strategies and variational autoencoders (VAEs). Specifically, we evaluate three primary types of text encoders: CLIP, large language models (LLMs), and the T5 model. Our study includes comprehensive ablations of causal versus bidirectional CLIP, layer selection strategies for both CLIP and LLMs, and a final performance comparison of all three encoders. We also introduce a refined variational autoencoder (VAE) [24] that better preserves fine-grained visual details, further boosting image quality, especially complex visual features.

Finally, with a progressive training approach, DiT-Air achieves new state-of-the-art GenEval [15] and T2I Comp-Bench [22] scores of 82.9 and 59.5, respectively. As shown in Figure 1, our model delivers superior performance with outstanding parameter efficiency compared to leading models such as SD3, FLUX, and JanusPro. Example generation results are provided in Figure 2.

Our key contributions are as follows: $(i)$ We systematically study the design choices of a range of DiT-based architectures including PixArt-style and MMDiT variants. $(ii)$ We introduce DiT-Air and DiT-Air-Lite, a novel DiT model family that simply extends a standard DiT by directly processing concatenated text and noise inputs. $(iii)$ We demonstrate parameter efficiency, achieving a 66% model size reduction with minimal performance impact compared to the state-of-the-art MMDiT. $(iv)$ We establish a new state-of-the-art performance on GenEval [15] and T2I Comp-Bench [22].

## 2. Related Works

### 2.1. Text-to-Image Diffusion Models

Diffusion models [11, 20, 41] have achieved remarkable success in text-to-image generation [37, 39]. These mod-

els generate images by iteratively denoising random noise, guided by semantic text embeddings extracted from pre-trained text encoders, such as CLIP [36]. Latent diffusion methods [34, 38] further enhanced training and inference efficiency by operating in a latent space defined by a pre-trained variational autoencoder (VAE) [24], reducing computational costs without sacrificing quality. Recently, flow matching objectives [31, 32, 43] have been introduced to connect source and target image distributions through simplified paths, offering further gains in image quality [14]. Our approach builds upon this paradigm by leveraging pre-trained text encoders for conditioning and optimizing a conditional flow matching objective in latent space.

### 2.2. Diffusion Transformers and Text Conditioning

Diffusion Transformers (DiTs) [33] were initially proposed to extend the advantages of transformer architectures to class-conditional image generation within diffusion models. Built upon vanilla vision transformers [13], widely utilized in image understanding tasks, DiT explored various approaches for incorporating noise level (time) and class label conditions. The study compared zero-initialized adaptive layer normalization (AdaLN), cross-attention, and in-context conditioning, demonstrating that AdaLN was the most effective strategy through comprehensive experiments.

Subsequent works extended DiTs to text-to-image generation, establishing them as a popular choice for high-quality open-source models [4, 14, 26]. Among them, PixArt models use cross attention between fixed text embeddings and image features after each of the self attention layers to inject text conditions into transformer models. MMDiT expands the transformer to a dual-stream design, with separate query, key, value, and output (QKVO) projections and MLPs for text embeddings and image features. Text and image features interact via scaled dot product attention applied on concatenated feature sequences. Compared to vanilla DiTs with a similar number of layers and feature dimensions, both PixArt and MMDiT expand the model size and use the extra parameters to convert semantically rich text embeddings into the visual space.

To justify these extra complications for text conditioning, the SD3 paper [14] reports that MMDiT outperforms several DiT variants. However, this study was conducted on the relatively small CC12M [2] dataset and at smaller model scales, without fully accounting for the additional parameters introduced by MMDiT's dual-stream architecture. Moreover, a scaling law study by Liang et al. [29] compared cross-attention with in-context conditioning but excluded AdaLN and focused only on models with 1M to 1B parameters, limiting its applicability to larger-scale settings. Through large-scale experiments, we demonstrate that our DiT-Air, which adheres closely to the simpler vanilla DiT architecture, can achieve comparable or superior performance to these more complex models, particularly at scale.

**Parameter Sharing in Transformers.** Parameter sharing has emerged as an effective approach to enhance efficiency in transformer-based architectures. For instance, ALBERT [27] demonstrates that sharing parameters across layers in a BERT-like model can substantially reduce model size while retaining competitive performance on NLP tasks. The parameter sharing approach in our proposed diffusion model architecture (Section 3.3.1) is inspired by this work, and offers new avenues for balancing generative quality and model compactness in DiT architectures.

# 3. Architecture Design

## 3.1. Latent Diffusion Framework

A typical text-to-image generation pipeline based on latent diffusion encodes the target image $\mathbf{x}$ into a latent representation $\mathbf{z}_0 \in \mathbb{R}^{h \times w \times c}$ via a variational autoencoder (VAE), *i.e.*, $\mathbf{z}_0 = E_{\text{VAE}}(\mathbf{x})$. Meanwhile, the text prompt $p$ is processed by one or more pretrained text encoders (*e.g.*, CLIP, T5), and the resulting embeddings are projected to produce token embeddings $\mathbf{c} \in \mathbb{R}^{l_{\text{text}} \times d}$.

During training (see Figure 3), a forward diffusion process adds noise to $\mathbf{z}_0$, producing a noisy latent $\mathbf{z}_t$ for a randomly sampled $t \in (0, 1)$. The model $f_\theta$ then learns to reverse this process by predicting a target quantity—such as noise $\epsilon$ in denoising models, velocity $\mathbf{v}$ in v-prediction models, or vector field $\mathbf{F}$ in flow-matching models—conditioned on the current latent $\mathbf{z}_t$, the text embedding $\mathbf{c}$, and the timestep $t$:

$$\hat{y}_t = f_\theta(\mathbf{z}_t, \mathbf{c}, t).$$

The training objective minimizes the discrepancy between $\hat{y}_t$ and the true target, effectively denoising $\mathbf{z}_t$. At inference, $f_\theta$ iteratively transforms a random sample $\mathbf{z}_T$ back to a clean latent $\hat{\mathbf{z}}_0$, which is then decoded by the VAE to produce the final image.
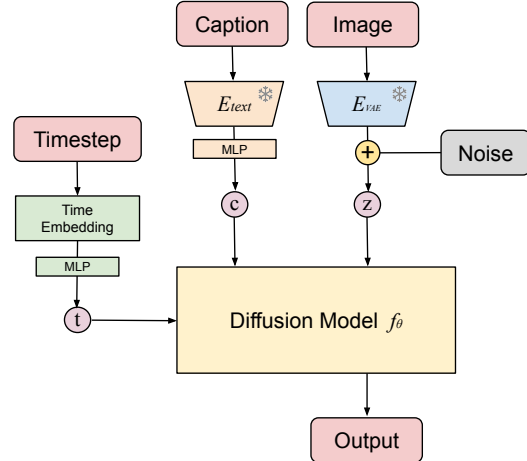


Figure 3. **Overview of Latent Diffusion Training.** During training, $\mathbf{x}$ is encoded into a latent $\mathbf{z}_0$ via a VAE, and the text prompt $p$ is mapped to embeddings $\mathbf{c}$. A forward diffusion adds noise to $\mathbf{z}_0$, and the model learns to reverse this process by predicting the noise (or similar target) at each timestep.

## 3.2. Text-conditioned Diffusion Transformers

In this section, we examine the backbone diffusion model $f_\theta$, specifically focusing on two widely-used Diffusion Transformer (DiT) variants: PixArt-$\alpha$ and MMDiT.

**PixArt-$\alpha$.** PixArt-$\alpha$ (Figure 4a) follows a two-step attention process: (1) self-attention on patchified visual tokens, and (2) cross-attention to fixed text embeddings $\mathbf{c}$. These text embeddings remain the same across all layers, providing a global conditioning signal for generation.

**MMDiT.** MMDiT (Figure 4b) adopts a dual-stream approach in each transformer block: text and visual tokens have separate query, key, value, and output (QKVO) projections. After computing self-attention over the concatenated sequence, tokens from each modality are processed by separate MLPs, preserving modality-specific transformations. This design supports rich multimodal interactions but substantially increases the parameter count.

Table 1 compares the main parameter components for PixArt-$\alpha$ and MMDiT, highlighting that MMDiT consumes significantly more parameters due to:

- Per-layer AdaLN: MMDiT uses independent AdaLN parameters in each layer, while PixArt-$\alpha$ shares them.
- Dual-Stream Design: MMDiT duplicates QKVO and MLPs for text and image tokens, effectively doubling related parameter sets.
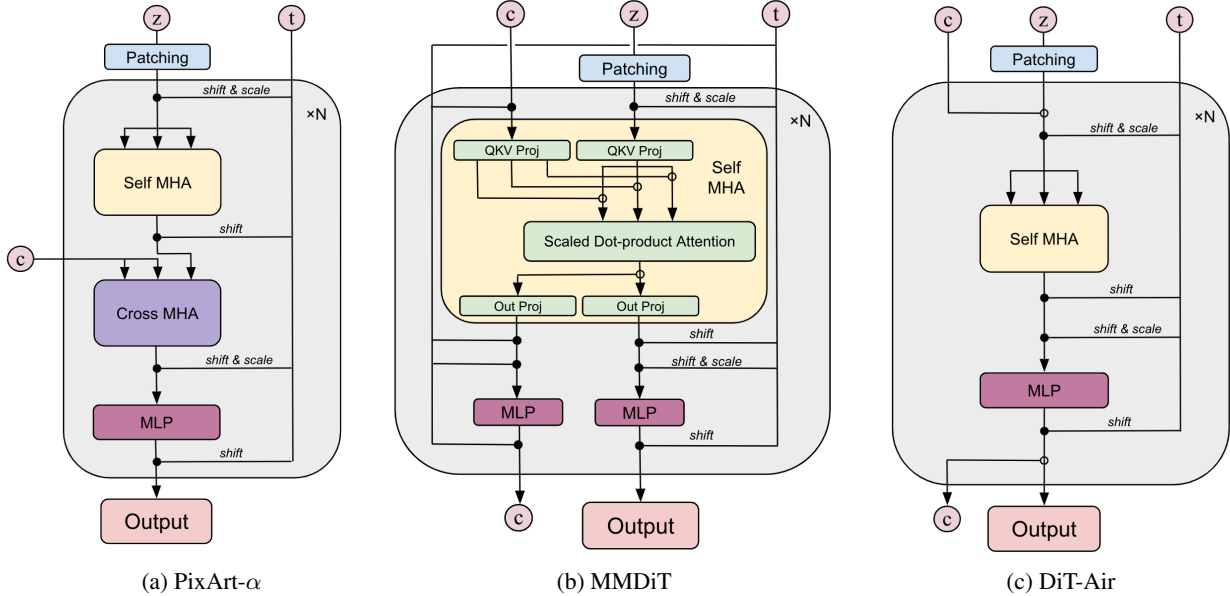
**Figure 4. Comparison of Diffusion Transformer Architectures.** Element-wise operations are denoted by $\bullet$, and sequence-wise operations by $\circ$. The details of inputs $\mathbf{c}$, $\mathbf{z}$, $t$ can be found in Figure 3. PixArt-$\alpha$ relies on sequential self- and cross-attention, whereas MMDiT uses a dual-stream approach with separate parameters for text and image tokens. Our proposed DiT-Air resembles a vanilla DiT that processes concatenated text and noises.

Table 1. Parameter counts for PixArt-$\alpha$, MMDiT, DiT-Air, and DiT-Air-lite (full vs. QKVO). $d$ represents the effective embedding size in transformer blocks (*i.e.*, the model's hidden dimension), and $N$ denotes the number of layers. Our DiT-Air is a compact DiT with shared dual-stream AdaLN, hence saves parameters significantly.

| Component | PixArt-$\alpha$ | MMDiT | DiT-Air | DiT-Air-Lite (full) | DiT-Air-Lite (attention) |
|---|---|---|---|---|---|
| AdaLN | $6d^2$ | $12Nd^2$ | $12d^2$ | $12d^2$ | $12d^2$ |
| Self-MHA | $4Nd^2$ | $8Nd^2$ | $4Nd^2$ | $4d^2$ | $4d^2$ |
| Cross-MHA | $4Nd^2$ | – | – | – | – |
| MLP | $8Nd^2$ | $16Nd^2$ | $8Nd^2$ | $8d^2$ | $8Nd^2$ |
| Total | $(6+16N)d^2$ | $36Nd^2$ | $(12+12N)d^2$ | $24d^2$ | $(16+8N)d^2$ |

### 3.3. DiT-Air: A Compact Text-to-Image DiT

In this section, we introduce **DiT-Air**, a compact DiT architecture tailored for text-conditioned generation. Instead of having the dual-stream approach of MMDiT, DiT-Air employs a unified set of QKVO projections and MLPs for both text and image tokens, offering a streamlined yet effective design (see Figure 4c).

For adaptive layer normalization (AdaLN), DiT-Air combines the dual-stream AdaLN from MMDiT with the parameter-sharing strategy of PixArt-$\alpha$. By sharing AdaLN parameters across all layers, DiT-Air effectively limits parameter growth as model depth increases, achieving a balanced trade-off between efficiency and multimodal capacity. Notably, this approach maintains a constant parameter overhead for AdaLN, irrespective of model depth.

As shown in Table 1, DiT-Air's design reduces the parameter count by approximately $24Nd^2$ compared to MMDiT, while preserving the same computational complexity (FLOPs). Overall, DiT-Air consumes about 66% fewer parameters than MMDiT and 25% fewer than PixArt-$\alpha$ for large $N$.

However, two open questions remain:

- *Does sharing AdaLN across all layers adversely affect image quality?*
- *How does merging text and image streams influence text alignment and fidelity?*

We investigate these aspects in Section 4. Further, we also explore whether parameter savings can be pushed even more aggressively, as described next.

### 3.3.1. DiT-Air-Lite

While DiT-Air strikes a balance between multimodal capacity and parameter efficiency, some applications (*e.g.*, large-scale inference or deployment on resource-constrained hardware) demand even greater reductions in parameter count. Inspired by the shared AdaLN and ALBERT [27], which shares parameters across layers to reduce model size, we propose DiT-Lite to further reduce parameters by sharing Transformer block parameters across layers, either entirely or partially (see Table 1).

Table 2. Performance comparison between MMDiT variants with and without sharing AdaLN. Both models are based on the MMDiT/B sized configuration.

| Model | Params | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|---|
| Per-layer AdaLN | 902M | 0.422 | 14.7 | 32.9 | 20.28 | 69.8 | 5.62 | 51.1 |
| Shared AdaLN | 631M | 0.422 | 15.0 | 32.9 | 20.28 | 69.8 | 5.59 | 51.0 |

**Full Block-Sharing Variant.** In this variant, the *entire* Transformer block (QKVO projections and MLP) is shared across all $N$ layers. This reduces the overall parameter count to $\sim 24d^2$, making the model extremely compact. However, it also limits representational diversity, since every layer applies the exact same transformation, which can degrade performance on complex prompts.

**Attention-Sharing Variant.** A more moderate approach involves *sharing only* the QKVO projections while maintaining a distinct MLP for each layer. As shown in Table 1, this partial-sharing strategy reduces the parameter count by approximately 33% compared to the non-shared version, while generally preserving higher fidelity and better text alignment than the fully shared configuration. Having dedicated MLPs for each layer allows the model to capture depth-specific nuances in the text-to-image mapping.

## 4. Experiments

In this section, we evaluate the impact of architectural variations and design choices using a standardized training and evaluation protocol. Specifically, by fixing other components (*e.g.*, VAE, text encoder), we isolate the effects of attention mechanisms and parameter utilization, allowing for a clear analysis of each architectural change.

### 4.1. Experimental Setup

We outline the dataset, training setup, model scaling strategy, and evaluation metrics used in our experiments.

**Data**: We conduct all ablations on in-house data containing 1.5 billion text-image pairs. Following DALL·E 3 [1], we enrich the dataset using synthetic captions generated by a pretrained captioning model. To balance real and synthetic data, we adopt a 1:9 ratio between original and synthetic captions. Unlike prior work that often relies on smaller subsets, all ablation studies are performed on the full dataset, ensuring a more reliable assessment of real-world performance.

**Training and Inference**: To ensure fair and consistent comparisons across experiments, we standardize key components: all models use a shared in-house variational autoencoder (VAE) and an in-house CLIP-H model as the default text encoder unless stated otherwise. Implementa-

tion details for these components are provided in the Appendix A.

All models are trained using a flow-matching objective, optimized with AdaFactor at a fixed learning rate of 1e-4, and a global batch size of 4096 for 1 million steps. For inference, we employ a Heun SDE solver [23] with 50 sampling steps and a classifier-free guidance scale of 7.5.

For model scaling experiments, we evaluate models at five specifications: S, B, L, XL and XXL, which correspond to 12, 18, 24, 30, and 38 transformer layers, respectively. The hidden dimension $d$ scales proportionally with depth as $d = 64 \times n_{\text{layer}}$, following the scaling strategy of SD3 [14]. Unless otherwise specified, all ablation studies use the B-size model. All experiments are done using the axlearn framework.[1]

**Evaluation**: We evaluate model performance using a combination of validation loss and a diverse set of established benchmarks. Following SD3 [14] and MovieGen [35], we report validation loss on both our in-house dataset. While validation loss provides a general measure of model fit, it may not accurately capture text alignment performance, particularly in complex generative tasks. As discussed further in Section 5.1, this limitation underscores the importance of incorporating additional metrics that directly evaluate alignment and compositionality. Therefore, we include Fréchet Inception Distance [19] on COCO30k [30], CLIP-Score [18, 36], and PickScore [25] on MJHQ30k [28], along with GenEval [15], T2I CompBench [21] and LAION-Aesthetics Predictor V2 (Aesthetics) [40]. For multi-category benchmarks such as GenEval and T2I CompBench, we report the overall average in the main paper with detailed per-category results deferred to the Appendix.

### 4.2. Adaptive Layer Normalization Sharing

As DiT-Air can be viewed as a simplification of MMDiT by unifying text and image streams, we start with investigating the impact of sharing AdaLN parameters across layers on the quality of generated images by comparing two MMDiT variants: one with shared AdaLN parameters and one without. The results are presented in Table 2.

As shown, the variant with shared AdaLN parameters is more parameter-efficient, reducing the model size from 902M to 631M. The performance differences between the two variants are minimal, with only slight variations in validation loss and other evaluation metrics, all of which fall within the error bounds. Thus, we conclude that sharing AdaLN parameters improves parameter efficiency without significantly affecting image quality.

### 4.3. Scaling and Efficiency

In this section, we evaluate the performance and efficiency of three architectures: PixArt-$\alpha$, MMDiT, and DiT-Air, for

---
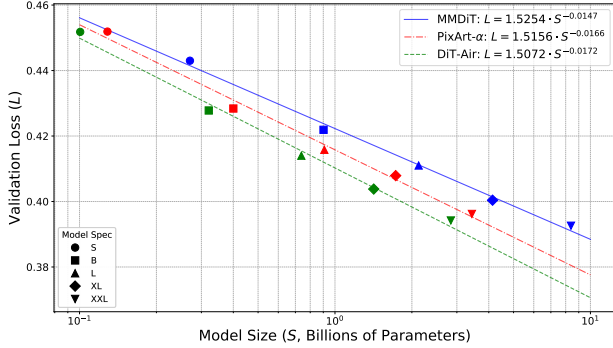
[1] https://github.com/apple/axlearn

Figure 5. **Validation Loss vs. Model Size for PixArt-α, MMDiT, and DiT-Air.** The plot illustrates the scaling behavior of three architectures across model sizes ranging from S to XXL, where the model size refers only to the diffusion transformer component (excluding the text encoder and VAE). The x-axis is in logarithmic scale, and the fitted lines depict the scaling trend using the formula $L = a \cdot S^b$. Among the three, DiT-Air achieves the best parameter efficiency.

various model scales from S to XXL (150M to 8B parameters). Our analysis focuses on validation loss and benchmark performance to highlight the relationship between model size, parameter efficiency, and architecture.

**Validation Loss and Scaling Behavior.** Figure 5 illustrates the validation loss versus model size for three architectures. DiT-Air demonstrates the best parameter efficiency, largely due to its shared AdaLN parameters across layers and its single-stream design for QKVO and MLPs. Among the three models, DiT-Air exhibits the steepest scaling curve, reflecting the most efficient reduction in validation loss as model size increases. Notably, at the S scale, DiT-Air's validation loss is considerably higher than that of MMDiT; however, as the models scale up to XXL, the loss gap vanishes. PixArt-α follows a scaling trend similar to DiT-Air—with a slope substantially steeper than that of MMDiT—but for a fixed parameter budget, DiT-Air consistently achieves a lower validation loss, underscoring its balanced approach to parameter efficiency and performance.

**Benchmark Performance.** Figure 6 compares benchmark performance across all three architectures and model scales. Overall, **DiT-Air** demonstrates consistently strong performance with a significantly higher degree of parameter efficiency. Notably, DiT-Air achieves some of the lowest FID scores while performing on par or slightly better than MMDiT and PixArt-α in PickScore and GenEval. Although MMDiT exhibits higher aesthetics scores at larger scales and PixArt-α achieves strong CLIPScores, these gains often remain within the error bounds or necessitate consider-

Table 3. Comparison of DiT-Air (baseline) with DiT-Air-lite (full) and DiT-Air-lite (attention). The *Full* configuration provides the largest parameter reduction but suffers a noticeable performance drop, while *Attention* strikes a more favorable trade-off between compactness and text alignment.

| Model | Params | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|---|
| DiT-Air/B | 321M | 0.428 | 16.0 | 32.8 | 20.2 | 70.4 | 5.58 | 51.4 |
| DiT-Air/B-lite (full) | 49M | 0.461 | 14.8 | 31.4 | 19.5 | 58.4 | 5.36 | 47.6 |
| DiT-Air/B-lite (attention) | 230M | 0.431 | 17.4 | 32.5 | 20.1 | 66.9 | 5.50 | 49.6 |

ably more parameters. This highlights DiT-Air's strength in maintaining competitive performance without the overhead of excessive scaling.

No single metric fully captures all aspects of text-to-image alignment and quality, underscoring the importance of evaluating models across a diverse set of benchmarks. Overall, the results establish DiT-Air as a compelling choice for large-scale text-to-image generation, offering a well-balanced mix of efficiency and performance.

### 4.4. DiT-Air-Lite Ablation

Having established in Section 4.3 that DiT-Air (DiT with shared AdaLN) compares favorably to PixArt-α and MMDiT, we now evaluate DiT-Air-Lite, the more aggressive parameter-sharing extension introduced in Section 3.3.1. We consider two configurations: *Full Block-Sharing* and *Attention-Sharing*. Full Block-Sharing reuses entire block across all layers, whereas Attention-Sharing keeps distinct MLPs per layer but shares QKVO. Table 3 compares both DiT-Air-lite variants against the baseline.

These results confirm that sharing the entire block minimizes parameters aggressively, but at the cost of lower text alignment and aesthetics. In contrast, *attention*-only sharing still yields substantial parameter savings while incurring only modest performance drops. Overall, DiT-Air-lite (attention) emerges as a favorable option when computational and memory constraints are critical, yet text alignment and image quality must remain high.

## 5. Text Encoders and VAEs

In this section we further investigate the impact of other critical components in the text-conditioned image generation task: text encoders and VAEs.

### 5.1. Text Encoder Ablation

We investigate how different text encoders impact the text alignment and overall performance of the DiT-Air architecture. Our study evaluates three primary encoder types: CLIP, T5, and Large Language Models (LLMs). Both the CLIP and LLM models are internal implementations, while
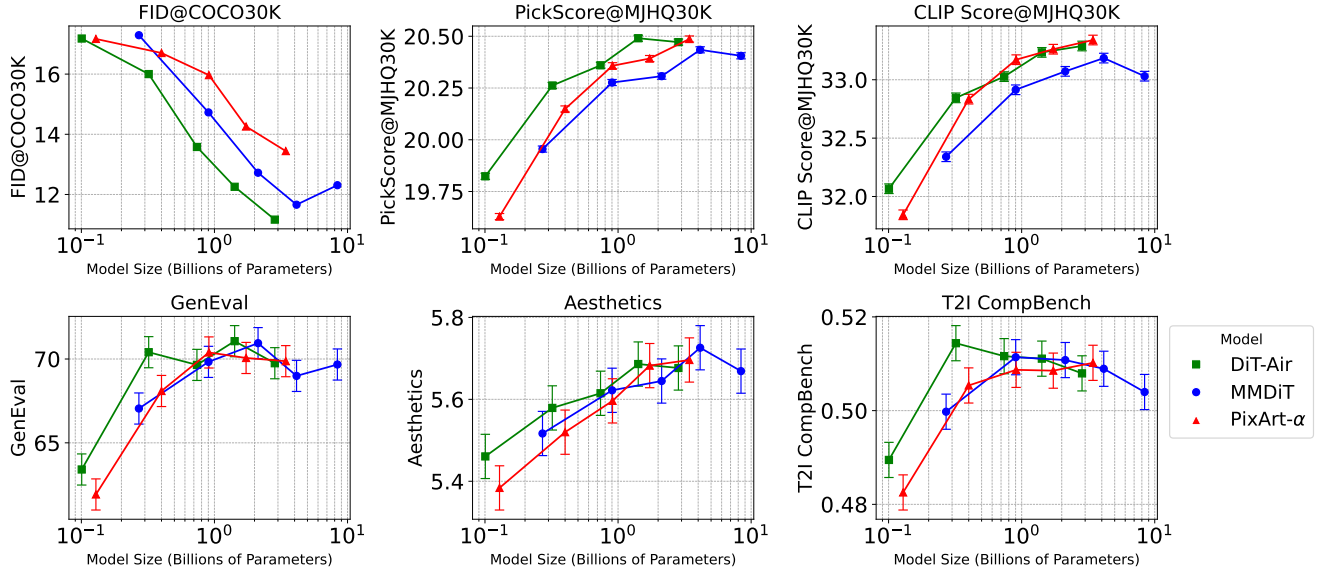
Figure 6. **Benchmark Performance Across Model Scales.** The plots compare PixArt-$\alpha$, MMDiT, and DiT-Air across six evaluation metrics. DiT-Air demonstrates strong parameter efficiency, achieving competitive performance with fewer parameters. The x-axis is in logarithmic scale, and error bounds are indicated where applicable.

the T5 encoder utilizes the open-sourced T5-XXL model. The analysis includes ablations of causal vs. bidirectional CLIP, layer selection strategies for both CLIP and LLMs, and a final comparison of all three encoders. Detailed ablation results are provided in Appendix C.

**Summary of Key Findings.** Our experiments demonstrate that the bidirectional CLIP model consistently outperforms its causal counterpart, showing improved text alignment and image quality across benchmarks. This improvement is attributed to the better synergy between the bidirectional attention in the text encoder and the diffuison transformer. Layer selection experiments, detailed in Appendix C.1.1, indicate that while very shallow layers in the CLIP model underperform, deeper layers yield comparable results.[2]

For LLM-based encoders, we find that text-only LLMs outperform their multimodal counterparts, particularly in the GenEval metric. Layer selection studies in LLMs reveal that using a middle layer combined with a layer at approximately 3/4 of the model's depth offers the best performance. This behavior likely stems from the LLM's pretraining objective, which is typically next-token prediction. Such an objective emphasizes fine-grained, token-level information at deeper layers, potentially at the expense of broader semantic understanding, which is more beneficial for text-to-image alignment.

---

[2]In this paper, "shallow" layers always refer to those near the input of the text encoder, while "deep" layers are closer to the output, with depth consistently measured from input to output.

Table 4. Comparison of CLIP, LLM, and T5 text embeddings. CLIP achieves superior results across most metrics, while LLM excels in GenEval. In contrast, T5 consistently underperforms.

| Model | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|
| CLIP (Bidirectional) | 0.428 | 16.0 | **32.8** | **20.3** | 0.704 | **5.58** | **51.4** |
| LLM (Text-only) | 0.427 | **16.0** | 32.0 | 20.1 | **0.726** | 5.57 | 48.6 |
| T5-XXL | **0.424** | 17.5 | 31.8 | 20.0 | 0.653 | 5.46 | 48.0 |

**Comparison of CLIP, LLM, and T5.** The final comparison of text encoders (Table 4) shows that bidirectional CLIP achieves the best performance across most benchmarks, with text-based LLMs also performing strongly, especially in GenEval. The T5-XXL model, while achieving the lowest validation loss, generally lags behind in benchmark performance. This discrepancy highlights an important observation: validation loss alone is not always indicative of text alignment performance, particularly when comparing different text encoder architectures. Models with lower validation loss, such as T5, may not necessarily offer the best real-world performance on established benchmarks.

**Takeaway for Final Model Design.** Based on these findings, our final model adopts a hybrid strategy, combining bidirectional CLIP with a text-based LLM to leverage both efficient text alignment and deeper semantic understanding. This approach ensures a balanced trade-off between parameter efficiency and robust results across diverse metrics.

Table 5. Comparison with state-of-the-art (SoTA) models. "Total (B)" includes parameters from the text encoder, VAE, and diffusion model, while "Trainable (B)" denotes only those parameters updated during training. The value marked with † is estimated from Figure 8 in [14] as the exact number was unavailable.

| Model | Size (B) | | Metrics | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Total | Trainable | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
| SDXL Base [34] | 3.5 | 2.6 | 268.0 | 22.1 | 17.0 | 55.0 | 4.32 | 40.6 |
| PixArt-$\alpha$ [4] | 5.4 | 0.6 | 120.7 | 27.3 | 17.0 | 55.7 | 5.76 | 44.7 |
| SD3 Medium [14] | 7.7 | 2.0 | 26.0 | 32.0 | 20.7 | 62.0 | 5.99 | 52.4 |
| SD3 [14] | 13.6 | 8.0 | – | – | – | 74.5 | – | 51.4$^\dagger$ |
| Flux-Dev [26] | 16.9 | 12.0 | 68.7 | 30.2 | 19.7 | 66.7 | 6.12 | 49.6 |
| Flux-Schnell [26] | 16.9 | 12.0 | 25.1 | 33.1 | 21.6 | 70.7 | 6.12 | 49.9 |
| JanusPro-7B [7] | 6.9 | 6.9 | **17.2** | 15.5 | 16.6 | 80.3 | 5.95 | 35.2 |
| DiT-Air/L-Lite | 1.2 | 0.7 | 23.1 | 33.9 | 21.5 | 78.4 | 6.06 | 55.4 |
| DiT-Air/XXL | 6.0 | 2.8 | 32.2 | **34.7** | **22.1** | **82.9** | **6.29** | **59.5** |

## 5.2. Progressive VAE Training

The selection of a Variational Autoencoder (VAE) and its training strategy is critical for achieving high image fidelity in text-to-image generation. Although increasing the channel capacity of the VAE generally improves image reconstruction quality, it can inflate the KL divergence, hindering subsequent diffusion training.

To balance this trade-off, we introduce a progressive training pipeline. Our approach starts by training a low-channel VAE (*e.g.*, with 4 channels) from scratch. In a second stage, an intermediate convolutional layer is replaced with a higher channel capacity (*e.g.*, upgraded to 8 channels), followed by continued training. Using this two-stage process, we trained an 8-channel VAE that is employed in all our models. This method not only enhances downstream text-to-image generation but also maintains competitive reconstruction performance compared to the 4-channel or 8-channel VAEs trained from scratch.

Our ablation studies reveal that training an 8-channel VAE from scratch achieves an rFID of 2.59, whereas our progressive approach—starting with a 4-channel VAE and later expanding to 8 channels—attains a comparable rFID of 2.61 while reducing the KL divergence from $9\times10^5$ to $7\times10^4$. This reduction in divergence is critical, as it leads to improved downstream text-to-image performance (GenEval of 70.4 and T2I CompBench of 51.4) compared to 69.4 and 50.7 when using the 8-channel model trained from scratch. More details and comparisons are provided in Appendix D.

## 6. Final Models

In this section, we introduce our two final models, **DiT-Air/XXL** and **DiT-Air/L-Lite (attention)**, developed through our progressive training pipeline. Both models undergo a multi-stage training process: initial training at $256^2$ resolution, further training at $512^2$, followed by supervised fine-tuning (SFT) on a curated subset following Dai et al. [9], and finally refined with reward fine-tuning using an approach similar to DRaFT [8] using the HPSv2 [44] reward model.[3]

**DiT-Air/XXL.** DiT-Air/XXL represents our high-capacity model, combining a bidirectional CLIP text encoder for efficient text alignment and a text-based LLM for rich semantic understanding, as detailed in Section 5.1. We opt for the XXL configuration to further push the boundaries of image quality.

**DiT-Air/L-Lite (attention).** For a more parameter-efficient solution, we introduce DiT-Air/L-Lite (attention), which relies solely on the bidirectional CLIP text encoder and an L-sized architecture. This design has a total of **1.15B** parameters—including the text encoder, VAE, and diffusion transformer—while maintaining competitive performance.

Image generation results for both models are provided in Appendix H.

**Comparison with State-of-the-Art.** As shown in Table 5, our models compare favorably with current state-of-the-art text-to-image systems. In particular, DiT-Air/XXL achieves an exceptional GenEval overall score of **82.9** and a T2I CompBench average score of **59.5**, outperforming many larger competitors while maintaining a relatively

---

[3]Detailed training procedures are provided in Appendix A.2.2 (Configurations), E (SFT), and F (Reward fine-tuning).

Table 6. Comparison of different training strategies. Resolutions ($256^2$ vs. $512^2$) refer to the image size used during pretraining.

| Training Stage | FID $\downarrow$ | CLIP $\uparrow$ | Pick $\uparrow$ | GenE. $\uparrow$ | Aesth. $\uparrow$ | T2I. $\uparrow$ |
|---|---|---|---|---|---|---|
| Pretrain $256^2$ | 12.0 | 33.4 | 20.5 | 71.1 | 5.57 | 50.7 |
| Pretrain $512^2$ | 13.0 | 33.5 | 20.7 | 74.2 | 5.62 | 51.7 |
| Supervised fine-tuning | 22.5 | 34.2 | 21.5 | 79.0 | 5.89 | 55.3 |
| Reward fine-tuning | 32.2 | 34.7 | 22.1 | 82.9 | 6.21 | 59.5 |

compact size of 5.95B parameters. Meanwhile, DiT-Air/L-Lite offers a compelling balance between efficiency and quality, with strong performance across metrics such as CLIPScore, PickScore, and T2I CompBench scores. Although FID is included for consistency, we note that its reliability diminishes post fine-tuning due to distribution shifts.

**Multi-stage Training.** Table 6 details the metrics of the multi-stage training pipeline for DiT-Air/XXL. Increasing the pretraining resolution from $256^2$ to $512^2$ slightly increases FID, indicating that higher-resolution data alone does not fundamentally alter the model's capacity for coherent image generation under similar training conditions. At the same time, other metrics improve modestly, suggesting that a larger input resolution helps capture finer semantic and aesthetic details before any fine-tuning. Subsequent supervised and reward-based fine-tuning produce a more noticeable shift in FID—likely driven by distribution changes introduced by specialized or narrower data—yet these stages yield a marked improvement in text–image alignment and overall quality.

## 7. Conclusion

In this work, we systematically study the performance of Diffusion Transformers (DiTs) for text-to-image generation, focusing on architectural choices, text-conditioning strategies, and training protocols. We find that the standard DiT architecture, when enhanced with shared AdaLN parameters and configured to directly process concatenated text and noise inputs, achieves superior parameter efficiency compared to alternative approaches, particularly at scale.

We introduce **DiT-Air** and **DiT-Air-Lite**, two models that enhance the parameter efficiency of the standard DiT backbone while carefully balancing model size and performance for text-to-image generation. Through comprehensive ablation studies of text encoders and variational autoencoders (VAEs), we identified design choices that significantly improve text-conditioned image generation quality. By applying a multi-stage training process—including supervised and reward fine-tuning—our final models set new state-of-the-art performance across key text-to-image generation benchmarks, outperforming existing models.

Our findings offer valuable insights into the development of more efficient and expressive text-to-image models, underscoring the potential for further optimizing diffusion architectures and training practices.

## References

[1] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. Improving image generation with better captions. *Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf*, 2(3):8, 2023. 5

[2] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pretraining to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3558–3568, 2021. 3

[3] Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart-$\sigma$: Weak-to-strong training of diffusion transformer for 4k text-to-image generation, 2024. 1

[4] Junsong Chen, YU Jincheng, GE Chongjian, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart-$\alpha$: Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. 1, 2, 8

[5] Junsong Chen, Yue Wu, Simian Luo, Enze Xie, Sayak Paul, Ping Luo, Hang Zhao, and Zhenguo Li. Pixart-$\delta$: Fast and controllable image generation with latent consistency models, 2024. 1

[6] Shoufa Chen, Mengmeng Xu, Jiawei Ren, Yuren Cong, Sen He, Yanping Xie, Animesh Sinha, Ping Luo, Tao Xiang, and Juan-Manuel Perez-Rua. Gentron: Diffusion transformers for image and video generation, 2024. 1

[7] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Januspro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 8

[8] Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet.

Directly fine-tuning diffusion models on differentiable rewards, 2024. 8, 4

[9] Xiaoliang Dai, Ji Hou, Chih-Yao Ma, Sam Tsai, Jialiang Wang, Rui Wang, Peizhao Zhang, Simon Vandenhende, Xiaofang Wang, Abhimanyu Dubey, et al. Emu: Enhancing image generation models using photogenic needles in a haystack. *arXiv preprint arXiv:2309.15807*, 2023. 8, 4

[10] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. In *International Conference on Machine Learning*, pages 7480–7512. PMLR, 2023. 1

[11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2

[12] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. *Advances in neural information processing systems*, 34:19822–19835, 2021. 1

[13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2

[14] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 1, 2, 3, 5, 8

[15] Dhruba Ghosh, Hanna Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment, 2023. 2, 5

[16] Xinyu Gong, Wuyang Chen, Tianlong Chen, and Zhangyang Wang. Sandwich Batch Normalization: A Drop-In Replacement for Feature Distribution Heterogeneity. In *Winter Conference on Applications of Computer Vision (WACV)*, 2022. 1

[17] Ali Hatamizadeh, Jiaming Song, Guilin Liu, Jan Kautz, and Arash Vahdat. Diffit: Diffusion vision transformers for image generation, 2024. 1

[18] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning, 2022. 5, 2

[19] 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, 2018. 5, 2

[20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 2

[21] Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation, 2023. 5, 2

[22] Kaiyi Huang, Chengqi Duan, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench++: An enhanced and comprehensive benchmark for compositional text-to-image generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2

[23] Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021. 5, 1

[24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022. 2

[25] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation, 2023. 5, 2

[26] Black Forest Labs. Flux. https://github.com/black-forest-labs/flux, 2024. 2, 8

[27] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*, 2020. 1, 3, 4

[28] Daiqing Li, Aleks Kamko, Ehsan Akhgari, Ali Sabet, Linmiao Xu, and Suhail Doshi. Playground v2.5: Three insights towards enhancing aesthetic quality in text-to-image generation, 2024. 5

[29] Zhengyang Liang, Hao He, Ceyuan Yang, and Bo Dai. Scaling laws for diffusion transformers. *arXiv preprint arXiv:2410.08184*, 2024. 3

[30] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context, 2015. 5

[31] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 2

[32] Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, pages 23–40. Springer, 2024. 2

[33] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 1, 2

[34] 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. *arXiv preprint arXiv:2307.01952*, 2023. 2, 8

[35] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, David Yan, Dhruv Choudhary, Dingkang Wang, Geet Sethi, Guan Pang, Haoyu Ma, Ishan Misra, Ji Hou, Jialiang Wang, Kiran Jagadeesh, Kunpeng Li, Luxin Zhang, Mannat Singh, Mary Williamson, Matt Le, Matthew Yu, Mitesh Kumar Singh, Peizhao Zhang, Peter Vajda, Quentin Duval, Rohit Girdhar, Roshan Sumbaly, Sai Saketh Rambhatla, Sam Tsai, Samaneh Azadi, Samyak

Datta, Sanyuan Chen, Sean Bell, Sharadh Ramaswamy, Shelly Sheynin, Siddharth Bhattacharya, Simran Motwani, Tao Xu, Tianhe Li, Tingbo Hou, Wei-Ning Hsu, Xi Yin, Xiaoliang Dai, Yaniv Taigman, Yaqiao Luo, Yen-Cheng Liu, Yi-Chiao Wu, Yue Zhao, Yuval Kirstain, Zecheng He, Zijian He, Albert Pumarola, Ali Thabet, Artsiom Sanakoyeu, Arun Mallya, Baishan Guo, Boris Araya, Breena Kerr, Carleigh Wood, Ce Liu, Cen Peng, Dimitry Vengertsev, Edgar Schonfeld, Elliot Blanchard, Felix Juefei-Xu, Fraylie Nord, Jeff Liang, John Hoffman, Jonas Kohler, Kaolin Fire, Karthik Sivakumar, Lawrence Chen, Licheng Yu, Luya Gao, Markos Georgopoulos, Rashel Moritz, Sara K. Sampson, Shikai Li, Simone Parmeggiani, Steve Fine, Tara Fowler, Vladan Petrovic, and Yuming Du. Movie gen: A cast of media foundation models, 2025. 5, 1

[36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 2, 5

[37] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022. 2

[38] 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. 2

[39] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 2

[40] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models, 2022. 5, 2, 3

[41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2

[42] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 1

[43] Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023. 2

[44] Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis, 2023. 8, 4

[45] Jiazheng 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. *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023. 2

# DiT-Air: Revisiting the Efficiency of Diffusion Model Architecture Design in Text to Image Generation

## Supplementary Material

## A. Implementation Details

Our model builds upon a diffusion transformer framework with design choices that enhance training stability and performance.

### A.1. Model

#### A.1.1. Diffusion Transformer Variants

We define five variants of the diffusion transformer: S, B, L, XL, and XXL. These correspond to models with 12, 18, 24, 30, and 38 transformer layers, respectively. The hidden dimension scales proportionally with the number of layers as $d = 64 \times n_{\text{layer}}$, and the number of attention heads is set to equal the transformer depth.

#### A.1.2. Stability Enhancements

To improve training stability, we integrate several techniques. First, we apply QK-normalization [10] following the SD3 methodology to stabilize query-key interactions within the attention mechanism. We also employ sandwich normalization [12, 16] in both the attention blocks and MLP modules, a method proven effective for large-scale model training. Additionally, rather than using static positional embeddings for visual tokens, we incorporate a 2D rotary positional embedding [42] within the attention mechanism to dynamically capture spatial relationships.

#### A.1.3. Conditional Inputs via AdaLN

In our implementation, we adopt the MMDiT approach by incorporating a pooled text embedding through Adaptive Layer Normalization (AdaLN) alongside the time embedding. We believe that this pooled embedding provides high-level semantic information with only a minimal increase in computational cost.

#### A.1.4. Textual and Language Components

For the textual component, we utilize an internal CLIP/H text encoder consisting of 24 transformer layers, a hidden dimension of 1024, and 16 attention heads, totaling approximately 335 million parameters. The pooled embedding is generated via last token pooling for causal CLIP or average pooling for bidirectional CLIP. Additionally, our internal language model (LLM) consists of 56 transformer layers, a hidden dimension of 6,656, and 16 attention heads, with around 2.8 billion parameters, and employs last token pooling to produce the pooled embedding for AdaLN.

### A.2. Training and Inference Details

#### A.2.1. Training Objective

The model is trained using a *flow-matching* objective:

$$\min_\theta \mathbb{E}_{(\mathbf{z}_0, \mathbf{c}),\, \epsilon \sim \mathcal{N}(0, I),\, t \sim \mathcal{P}(t)} \left[ \| f_\theta(\mathbf{z}_t, \mathbf{c}, t) - \mathbf{z}_0 + \epsilon \|_2^2 \right].$$

In this formulation, the timestep distribution $\mathcal{P}(t)$ follows a logit-Normal distribution as in SD3 [14], which emphasizes intermediate steps during the flow-matching process.

#### A.2.2. Training Setup

Training is conducted on TPU v5p hardware. We use the AdaFactor optimizer with a constant learning rate of $1 \times 10^{-4}$, and momentum parameters $b_1 = 0.9$ and $b_2 = 0.999$.

For ablation studies, the model is trained for 1 million steps with a batch size of 4,096 at a resolution of $256^2$. In contrast, the final models are trained in multiple stages:

1. An initial stage of 500k steps at $256^2$ resolution with a batch size of 4,096.
2. A subsequent stage of 100k steps at $512^2$ resolution with a batch size of 2,048.
3. A supervised fine-tuning (SFT) stage for 2.5k steps with a batch size of 64, where the timestep distribution $\mathcal{P}(t)$ is shifted so that the log-SNR aligns with that of the low-resolution training.
4. A reward fine-tuning stage for an additional 4.8k steps with a batch size of 64.

Further details regarding the SFT and reward finetuning stages are provided in Appendix E and Appendix F, respectively.

#### A.2.3. Inference

During inference, we employ a second-order Heun SDE solver [23] with 50 sampling steps, combined with a classifier-free guidance scale of 7.5 to steer the sampling process.

## B. Evaluation

We evaluate model performance using a combination of validation loss and several established benchmarks.

### B.1. Validation Loss

Recent works such as SD3 [14] and MovieGen [35] have proposed using validation loss as a performance estimate. We follow this trend. More specifically, in a flow-matching paradigm, the validation loss measures how well the model

learns the velocity field induced by the transport equation. With a fixed VAE across all experiments, the source and target distributions remain constant (*i.e.*, the standard normal distribution and the VAE-encoded latents, respectively). Consequently, the validation loss directly reflects the model's ability to predict the rectified flow trajectory, which often aligns with human perceptual preferences.

## B.2. Benchmark Metrics

We report performance on the following metrics:
- **Fréchet Inception Distance (FID):** Quantifies the similarity between the generated and real image distributions by comparing their Inception-v3 embeddings [19].
- **CLIPScore:** Assesses the semantic alignment between images and text using CLIP embeddings. Higher scores denote better correspondence [18, 36].
- **PickScore:** Similar to CLIPScore, but based on a CLIP model trained on an open dataset of text-to-image prompts and real user preferences, thereby achieving compelling performance in predicting human preferences [25].
- **GenEval:** Provides an overall evaluation of image generation performance. The original implementation, using 4 samples per prompt, tends to exhibit larger uncertainty bounds; to mitigate this, we increase the number of samples to 64 [15].
- **T2I CompBench:** A comprehensive benchmark for assessing text-to-image synthesis quality [21].
- **LAION-Aesthetics Predictor V2:** Predicts the aesthetic quality of images, with higher scores indicating superior visual appeal [40]. In our evaluation, images are generated using ImageReward prompts [45] and subsequently assessed with this aesthetic model.

## B.3. Abbreviation Key

For clarity, we list the abbreviations used in tables throughout the paper:
- **Val.** — Validation Loss.
- **FID** — Fréchet Inception Distance.
- **CLIP** — CLIPScore.
- **Pick** — PickScore.
- **GenE.** — GenEval.
- **Aesth.** — LAION-Aesthetics Predictor V2.
- **T2I.** — T2I CompBench.

## C. Detailed Text Encoder Ablation Studies

This section provides a detailed analysis of the text encoder ablation experiments, including the impact of causal vs. bidirectional attention in CLIP, the effect of layer selection in both CLIP and LLMs, and a comparison between text LLM and multi-modal LLM. The results presented here supplement the summary findings discussed in Section 5.1 of the main paper.

Table 7. Zero-Shot performance of causal vs. bidirectional CLIP models on ImageNet and COCO5k.

| CLIP Model | ImageNet | | COCO5k | | | |
|---|---|---|---|---|---|---|
| | Acc@1 ↑ | Acc@5 ↑ | I2T@1 ↑ | I2T@5 ↑ | T2I@1 ↑ | T2I@5 ↑ |
| Causal | 80.6 | 96.5 | 74.4 | 91.5 | 53.6 | 77.7 |
| Bidirectional | 80.6 | 96.5 | 74.6 | 91.5 | 53.8 | 78.3 |

Table 8. Performance comparison of causal vs. bidirectional CLIP models as text embedding models for text-to-image generation.

| CLIP Model | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|
| Causal | 0.429 | 16.4 | 32.8 | 20.2 | 0.683 | 5.61 | 50.6 |
| Bidirectional | **0.428** | **16.0** | 32.8 | **20.3** | **0.704** | 5.58 | **51.4** |

## C.1. CLIP

### C.1.1. Causal vs. Bidirectional Attention in CLIP

To assess the impact of attention mechanisms in the text encoder, we compared causal and bidirectional variants of the CLIP/H model. In the causal configuration, the CLIP model employs causal attention with last-token pooling, whereas the bidirectional variant uses global average pooling during contrastive loss training. Both models exhibit comparable performance in standard zero-shot classification and retrieval tasks, as summarized in Table 7. However, as shown in Table 8, the bidirectional CLIP consistently outperforms its causal counterpart in terms of text alignment and image quality benchmarks. We hypothesize that the observed improvements in diffusion models arise specifically from enhanced attention alignment, rather than from intrinsic differences in the pretrained text encoder performance.

### C.1.2. Layer Selection in CLIP

We investigated how the selection of different layers in the bidirectional CLIP text encoder (24 layers) affects performance. Embeddings from layers 6, 12, 18, 23, and 24 were tested, along with a concatenation of multiple layers (6, 12, 18, 24) followed by a linear projection. The 23rd layer was included as part of this study due to its common use as the penultimate layer in open-source text-to-image models. As shown in Table 9, our results indicate that all deeper layers (12, 18, 23, 24) exhibit comparable performance, while the shallow layer (6) underperforms. Concatenating embeddings from multiple layers did not yield significant improvements, suggesting that a single mid-to-deep layer is sufficient for robust text alignment.

## C.2. LLM

### C.2.1. Text LLM vs. Multimodal LLM

We further investigate the impact of the text encoder by comparing a text-only LLM with a multimodal LLM

Table 9. CLIP Layer Selection Performance.

| Model | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|
| Layer 6 | 0.428 | 15.9 | 32.3 | 20.1 | 68.4 | 5.46 | 51.2 |
| Layer 12 | 0.428 | 16.0 | 32.7 | 20.2 | 68.3 | 5.56 | 50.5 |
| Layer 18 | 0.428 | 16.1 | **32.8** | 20.2 | 69.1 | 5.56 | 51.2 |
| Layer 23 | 0.428 | 15.6 | **32.8** | 20.3 | **70.6** | **5.59** | 50.9 |
| Layer 24 | 0.428 | 16.0 | **32.8** | 20.3 | 70.4 | 5.58 | **51.4** |
| Layer 6 + 12 + 18 + 24 | 0.426 | **15.4** | **32.8** | 20.3 | 70.4 | 5.60 | 51.2 |

Table 10. Performance Comparison of text LLM vs. Multimodal LLM.

| Model | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|
| LLM | 0.427 | **16.0** | **32.0** | **20.1** | **72.6** | **5.57** | 48.6 |
| MLLM | 0.427 | 16.4 | 31.9 | 20.0 | 70.0 | 5.54 | **49.2** |

Table 11. LLM Layer Selection Performance.

| Model | Val. ↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|
| Layer 14 | 0.427 | **15.8** | 32.2 | 20.1 | 68.1 | 5.56 | 0.494 |
| Layer 28 | 0.427 | 16.5 | **32.3** | 20.2 | 71.7 | **5.64** | 0.502 |
| Layer 42 | 0.427 | 17.2 | **32.3** | 20.1 | 71.6 | 5.61 | 0.501 |
| Layer 56 (last) | 0.427 | 16.0 | 32.0 | 20.1 | 72.6 | 5.57 | 0.486 |
| Layer 28 + 42 | 0.427 | 15.9 | **32.3** | 20.2 | **73.0** | 5.60 | **0.507** |

(MLLM). Our experiments, summarized in Table 10, reveal that text-only LLMs tend to outperform their multimodal counterparts, particularly on the GenEval metric.

### C.2.2. Layer Selection in LLMs

We also evaluated different layers in the 56-layer text-only LLM to determine the optimal choice for text embeddings. Specifically, we compared embeddings from layer 14 (early), layer 28 (middle), layer 42 (deeper), and a concatenation of layers 28 and 42. Our results, summarized in Table 11, indicate that both the middle (28) and deeper (42) layers offer a strong balance between preserving low-level token details and capturing high-level semantic representations. In contrast, the final layer, while performing well on GenEval, provides less balanced representations for text-to-image generation tasks, possibly due to over-specialization in the pretraining objective.

## D. Progressive VAE Training Studies

To validate our progressive training approach described in Section 5.2, we conducted experiments on three VAE variants, all employing an 8× compression factor and trained on the OpenImages 9M dataset. The variants are defined as follows:

- **Variant A:** A VAE with 4 channels trained from scratch.

Table 12. Comparison of VAE on reconstruction and generation.

| Model | KL | rFID↓ | FID ↓ | CLIP ↑ | Pick ↑ | GenE. ↑ | Aesth. ↑ | T2I. ↑ |
|---|---|---|---|---|---|---|---|---|
| A | $7 \times 10^4$ | 4.62 | 17.2 | 32.7 | 20.2 | 69.8 | 5.52 | 50.4 |
| B | $9 \times 10^5$ | **2.59** | 16.3 | 32.8 | 20.2 | 69.4 | 5.56 | 50.7 |
| C | $7 \times 10^4$ | 2.61 | **16.0** | 32.8 | **20.3** | 70.4 | **5.58** | **51.4** |

- **Variant B:** A VAE with 8 channels trained from scratch.
- **Variant C:** A VAE with 8 channels trained using our proposed Progressive training approach. This variant is initially trained with 4 channels (as in Variant A) and subsequently refined by replacing an intermediate convolutional layer with one that uses 8 channels.

Our evaluation employs the reconstruction FID (rFID) metric on the COCO validation set to assess image reconstruction quality, along with an evaluation of the downstream diffusion model using DiT-Air/B. The experimental results, summarized in Table 12, indicate that although increasing the channel size significantly enhances reconstruction quality, it also leads to a higher KL divergence in the latent features. This elevated KL divergence can impede the latent diffusion model's learning, resulting in only marginal gains in final visual generation quality. In contrast, our progressive training pipeline mitigates this issue by first training a smaller VAE and then gradually increasing its channel capacity. This approach achieves notable improvements in text-to-image generation while maintaining competitive reconstruction performance.

## E. Supervised Fine-Tuning and Data Curation

In the supervised fine-tuning (SFT) stage, our goal is to refine the pretrained model using a very high-quality but relatively small dataset of image-text pairs. To this end, we curated a dataset of 1,033 pairs, ensuring that the images and their corresponding captions meet stringent quality standards.

The curation process involved several key steps:

1. **Automated Filtering:** Initially, we applied both the LAION image aesthetics model [40] and our internal photo aesthetics model to the pretraining data. This step allowed us to filter out images that did not meet our high aesthetic standards, ensuring that only the best images were considered.
2. **Manual Selection:** From the automatically filtered subset, we manually reviewed the images to further refine the selection. The focus here was on achieving diversity across object categories and image styles, with special attention given to important verticals such as people and animals.
3. **Caption Curation:** For the selected images, we crafted precise captions in the style of our automatic captioning model. This manual curation ensured that each caption

was not only accurate but also semantically well-aligned with the corresponding image.

Fine-tuning with the resulting dataset, as demonstrated in [9], can lead the model to converge to a state where the generated images surpass the average quality of the pre-training data.

Overall, our SFT strategy emphasizes quality over quantity. By leveraging a meticulously curated dataset, we ensure that the fine-tuning process yields improved image generation performance, achieving both higher aesthetic quality and better semantic alignment.

## F. Reward Fine-tuning

We adopt an approach similar to DRaFT [8] to fine-tune our models using the HPSv2 [44] reward model. In our setup, the models receive a prompt $p$ and an initial latent noise $\mathbf{z}_T$ as inputs, which are then denoised over $T$ timesteps to generate the final image $I_0$. The HPSv2 model computes a human preference score for the generated image, denoted as $r(p, I_0)$, with scores normalized between 0 and 1. Consequently, the loss backpropagated through the sampling chain is defined as $1 - r(p, I_0)$.

To specifically target areas where the model underperforms, we selected 2,000 of the lowest-scoring prompts as the training data for reward fine-tuning. During this stage, we fine-tune the full set of model parameters, using a total sampling timestep $T = 50$ and a stop gradient timestep $T_s = 25$ to ensure that gradients are propagated only from timestep $T$ down to $T_s$, with no gradient updates before $T_s$.

Despite these precautions, we observed a reward model hacking phenomenon [8], where HPSv2 occasionally assigns very high scores to poor-quality images. To mitigate this issue, we implemented an early stopping strategy to prevent overfitting and reward hacking. Ultimately, this approach effectively reduced structural artifacts, enhanced text alignment, and improved the overall visual appeal of the generated images.

## G. State-of-the-Art Model Size Breakdown

In Table 13, we present a detailed breakdown of the architectural components and parameter counts for various state-of-the-art text-to-image generation models. For JanusPro-7B, the SigLIP encoder is omitted for simplicity. The comparison demonstrates that DiT-Air is relatively compact—both in total parameter count and in trainable parameters—when compared with existing models.

## H. Generation Examples

A selection of model-generated images illustrating different capabilities are shown in Figure 7 (generations by DiT-Air/XXL) and Figure 8 (generations by DiT-Air/L-Lite).

Through qualitative observation, we find that the DiT-Air/XXL consistently excels in generating complex scenes, capturing intricate structural details, and delivering superior visual quality. It demonstrates a strong ability to produce highly detailed and realistic images, even when handling lengthy or complex prompts. Its integrated LLM encoder contributes significantly to its robust text rendering capabilities, maintaining clarity and accuracy even in challenging scenarios.

In contrast, the DiT-Air/L-Lite offers a well-balanced approach, prioritizing efficiency while still delivering strong performance across a variety of tasks. It is particularly well-suited for scenarios with limited computational resources, providing high-quality images and effective handling of most prompts. While the DiT-Air/L-Lite maintains an excellent balance of efficiency and quality, our qualitative observations indicate that it may occasionally struggle with more ambiguous prompts and exhibit limitations in rendering complex text or achieving the same level of visual fidelity as the DiT-Air/XXL. These observations highlight the intended trade-off of the DiT-Air/L-Lite: a streamlined design that prioritizes efficiency, making it a compelling choice when balancing performance with resource constraints.

Overall, these observations highlight the trade-offs between the two models. The DiT-Air/XXL is ideal for tasks that demand high-quality, detailed images and strong text-image alignment, whereas the DiT-Air/L-Lite serves as a compelling, resource-efficient alternative for more lightweight use cases.

Table 13. Breakdown of architectural components for various text-to-image generation models. Trainable components are marked with $\star$.

| Model | Text Encoder 1 | Text Encoder 2 | Diffusion Model$^\star$ | Autoencoder | Total (B) |
|---|---|---|---|---|---|
| SDXL Base | CLIP/L (123M) | OpenCLIP/g (694M) | U-Net (2.6B) | 8-ch (84M) | 3.50 |
| PixArt-$\alpha$ | Flan-T5-XXL (4.7B) | – | DiT (0.6B) | 4-ch (80M) | 5.38 |
| SD3 Medium | CLIP/L + bigG (817M) | T5-v1.1-XXL (4.7B) | DiT (2.0B) | 16-ch (85M) | 7.65 |
| SD3 | CLIP/L + bigG (817M) | T5-v1.1-XXL (4.7B) | DiT (8.0B) | 16-ch (85M) | 13.60 |
| Flux-Dev | CLIP/L (123M) | T5-XXL (4.7B) | DiT (12B) | 16-ch (85M) | 16.91 |
| Flux-Schnell | CLIP/L (123M) | T5-XXL (4.7B) | DiT (12B) | 16-ch (85M) | 16.91 |
| JanusPro-7B | – | – | LLM (6.9B) | 16-ch VQ (85M) | 6.91 |
| DiT-Air/L-Lite | CLIP/H (335M) | – | DiT (0.7B) | 8-ch (84M) | 1.15 |
| DiT-Air/XXL | CLIP/H (335M) | LLM (2.8B) | DiT (2.8B) | 8-ch (84M) | 5.95 |

Figure 7. Sample images from our DiT-Air/XXL illustrating different capabilities.

**Complex Scenes**

A detailed steampunk laboratory with bubbling potions, mechanical contraptions, and papers scattered across an old wooden desk.

A dragon's hoard hidden inside a deep cave— glittering gold coins, gemstones in every color, and ancient scrolls scattered among the treasure. A pair of glowing eyes peek from the darkness.

**Creative Drawings**

A tiny mouse wearing a wizard hat, holding a magical staff and casting sparks in a forest at twilight.

A cat with ethereal, shimmering fur, floating through a dreamy, star-filled galaxy.

**Intricate Structures**

A close-up of an aged leather-bound book, highlighting the intricate embossing and weathered texture.
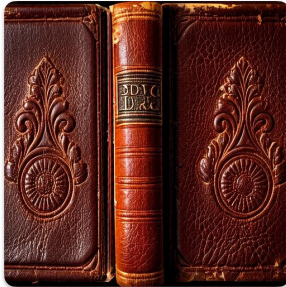
A dewdrop-covered spiderweb at sunrise, with each strand glistening in the light.

**Lighting & Shades**

A honeycomb dripping with golden honey, with glistening droplets catching the warm light.

A freshly painted oil canvas, with thick impasto strokes catching the light.

**Artistic Styles**

A boy is on his way to school with birds singing in 3d animation style.
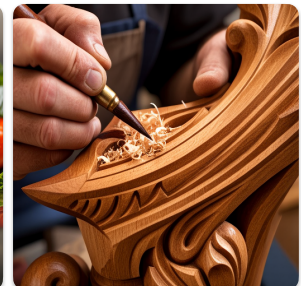
A painting of panda wearing a crown in Van Gogh style.

**Hands & Gestures**

A chef slicing a juicy ripe tomato, with glistening droplets of juice and knife reflections captured in detail.

A close-up of a craftsman carving an intricate wooden sculpture, the wood grain and shavings vividly detailed.

**Text Writing**

A digital clock on a spaceship dashboard reading "MISSION TIME/ 2034-07-16 14/35/22."

An adult cat teaches a kitty in a classroom with "How to say Meow" on the whiteboard.

**Typography**

Think Different made of colorful feathers.

a graffiti wall in an urban alley with vibrant street art and the phrase "art is freedom" integrated seamlessly into the design.

Figure 8. Sample images from our DiT-Air/L-Lite illustrating different capabilities.

7

Table 14. Supplementary Table: Detailed GenEval and T2I CompBench Breakdown.

| Method | GenEval | | | | | | | T2I Compbench (%) | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Overall | Single Object | Two Objects | Counting | Colors | Position | Color Attribution | Average | Color | Shape | Texture | Spatial | Non-Spatial | Complex |
| **Table 2** | | | | | | | | | | | | | | |
| Per-layer Adaln | 69.8 | 99.3 | 88.7 | 68.6 | 81.0 | 22.0 | 59.3 | 51.1 | 81.8 | 62.2 | 74.9 | 19.6 | 31.3 | 37.1 |
| Shared Adaln | 69.8 | 99.3 | 90.0 | 68.2 | 79.1 | 23.5 | 58.7 | 51.0 | 81.8 | 61.9 | 73.3 | 19.9 | 31.3 | 37.7 |
| **Table 3** | | | | | | | | | | | | | | |
| DiT-Air/B | 70.4 | 99.7 | 88.6 | 68.1 | 80.0 | 25.2 | 60.8 | 51.4 | 83.0 | 61.7 | 73.4 | 21.6 | 31.3 | 37.7 |
| DiT-Air/B-lite (full) | 58.4 | 96.8 | 69.3 | 49.8 | 75.7 | 14.1 | 44.4 | 47.6 | 79.1 | 57.4 | 69.4 | 13.5 | 30.8 | 35.6 |
| DiT-Air/B-lite (attention) | 66.9 | 99.1 | 84.3 | 65.2 | 79.2 | 18.5 | 54.9 | 49.6 | 79.9 | 60.3 | 71.2 | 18.0 | 31.2 | 37.1 |
| **Figure 4** | | | | | | | | | | | | | | |
| Pixart-$\alpha$/S | 61.9 | 98.7 | 81.2 | 58.3 | 79.4 | 15.3 | 52.0 | 48.3 | 80.2 | 57.5 | 68.8 | 15.8 | 31.0 | 36.2 |
| Pixart-$\alpha$/B | 68.1 | 99.5 | 89.4 | 67.3 | 79.6 | 21.2 | 57.7 | 50.5 | 82.7 | 60.5 | 72.5 | 18.9 | 31.4 | 37.1 |
| Pixart-$\alpha$/L | 70.4 | 99.2 | 89.1 | 69.5 | 80.4 | 24.2 | 58.2 | 50.9 | 82.7 | 61.0 | 73.0 | 19.5 | 31.4 | 37.7 |
| Pixart-$\alpha$/XL | 70.1 | 99.2 | 88.6 | 68.2 | 83.4 | 26.1 | 57.9 | 50.9 | 81.9 | 61.3 | 72.7 | 20.4 | 31.3 | 37.5 |
| Pixart-$\alpha$/XXL | 69.9 | 99.2 | 88.6 | 71.2 | 81.9 | 26.8 | 56.5 | 51.0 | 82.3 | 63.0 | 72.8 | 19.0 | 31.2 | 37.8 |
| MMDiT/S | 67.0 | 99.5 | 89.4 | 73.6 | 83.1 | 29.1 | 58.8 | 50.0 | 80.6 | 60.8 | 72.7 | 18.2 | 31.0 | 36.5 |
| MMDiT/B | 69.8 | 99.3 | 88.7 | 68.6 | 81.0 | 22.0 | 59.3 | 51.1 | 81.8 | 62.2 | 74.9 | 19.6 | 31.3 | 37.1 |
| MMDiT/L | 70.9 | 99.3 | 90.9 | 71.2 | 82.3 | 23.8 | 58.3 | 51.1 | 82.3 | 62.2 | 74.4 | 19.1 | 31.3 | 37.2 |
| MMDiT/XL | 69.0 | 98.9 | 87.3 | 66.3 | 81.1 | 26.4 | 53.9 | 50.9 | 82.6 | 63.0 | 73.1 | 18.1 | 31.2 | 37.4 |
| MMDiT/XXL | 69.7 | 99.3 | 87.4 | 67.1 | 83.0 | 24.3 | 56.9 | 50.4 | 82.3 | 62.7 | 73.2 | 15.7 | 31.3 | 37.2 |
| DiT-Air/S | 63.4 | 98.0 | 78.2 | 57.2 | 79.0 | 15.0 | 53.1 | 48.9 | 80.1 | 59.1 | 71.6 | 15.2 | 31.0 | 36.6 |
| DiT-Air/B | 70.4 | 99.7 | 88.6 | 68.1 | 80.0 | 25.2 | 60.8 | 51.4 | 83.0 | 61.7 | 73.4 | 21.6 | 31.3 | 37.7 |
| DiT-Air/L | 69.6 | 99.4 | 88.4 | 67.1 | 79.5 | 24.7 | 58.9 | 51.2 | 82.6 | 62.5 | 73.7 | 19.6 | 31.1 | 37.5 |
| DiT-Air/XL | 71.1 | 99.3 | 90.0 | 71.4 | 82.1 | 26.8 | 56.8 | 51.1 | 82.6 | 62.0 | 73.4 | 19.9 | 31.2 | 37.6 |
| DiT-Air/XXL | 69.7 | 99.0 | 89.0 | 70.1 | 81.3 | 24.9 | 54.3 | 50.8 | 82.7 | 62.6 | 73.2 | 17.9 | 31.5 | 36.8 |
| **Table 4** | | | | | | | | | | | | | | |
| CLIP (Bidirectional) | 70.4 | 99.3 | 88.7 | 68.6 | 81.0 | 22.0 | 59.3 | 51.1 | 81.8 | 62.2 | 74.9 | 19.6 | 31.3 | 37.1 |
| LLM (Text-only) | 72.6 | 99.2 | 88.7 | 65.7 | 83.6 | 34.7 | 63.6 | 48.6 | 78.8 | 56.8 | 69.8 | 18.4 | 30.9 | 37.1 |
| T5-XXL | 66.9 | 65.3 | 82.7 | 63.9 | 78.3 | 19.5 | 49.6 | 48.0 | 76.3 | 57.5 | 69.1 | 17.6 | 31.1 | 36.5 |
| **Table 5** | | | | | | | | | | | | | | |
| SDXL | 55.7 | 98.0 | 74.0 | 39.0 | 85.0 | 15.0 | 23.0 | 40.6 | 58.8 | 46.9 | 53.0 | 21.3 | 31.2 | 32.4 |
| PixArt-$\alpha$ | 47.8 | 98.0 | 50.0 | 44.0 | 80.0 | 8.00 | 7.00 | 44.6 | 66.9 | 49.3 | 64.8 | 20.6 | 32.0 | 34.3 |
| SD3-medium | 62.0 | 98.0 | 74.0 | 63.0 | 67.0 | 34.0 | 36.0 | 52.4 | 81.3 | 58.9 | 73.3 | 32.0 | 31.4 | 37.7 |
| SD3 | 74.5 | 99.0 | 94.0 | 72.0 | 89.0 | 33.0 | 60.0 | 51.4 | – | – | – | – | – | – |
| Flux-dev | 66.7 | 99.0 | 81.0 | 79.0 | 74.0 | 20.0 | 47.0 | 49.6 | 74.1 | 57.2 | 69.2 | 28.6 | 31.3 | 37.0 |
| Flux-schnell | 70.7 | 99.0 | 92.0 | 73.0 | 78.0 | 28.0 | 54.0 | 49.9 | 73.9 | 55.8 | 68.5 | 31.2 | 31.5 | 38.6 |
| Janus-pro | 80.3 | 99.0 | 89.0 | 59.0 | 90.0 | 79.0 | 66.0 | 35.2 | 52.0 | 33.1 | 40.6 | 15.4 | 31.3 | 39.2 |
| DiT-Air/L-Lite | 77.6 | 99.9 | 96.5 | 80.8 | 86.8 | 30.5 | 71.3 | 55.4 | 86.1 | 66.6 | 79.5 | 27.7 | 31.2 | 41.3 |
| DiT-Air/XXL | 82.9 | 99.9 | 98.8 | 83.8 | 86.4 | 50.5 | 77.7 | 59.5 | 88.9 | 69.4 | 82.0 | 40.6 | 31.9 | 44.1 |
| **Table 6** | | | | | | | | | | | | | | |
| Pretrain $256^2$ | 71.1 | 99.3 | 88.7 | 69.0 | 82.0 | 32.0 | 55.7 | 50.7 | 81.9 | 61.2 | 72.4 | 19.5 | 31.3 | 37.9 |
| Pretrain $512^2$ | 74.2 | 99.4 | 91.7 | 73.6 | 84.0 | 35.4 | 61.4 | 51.7 | 81.9 | 61.2 | 73.9 | 22.6 | 31.1 | 39.6 |
| Supervised fine-tuning | 79.0 | 99.9 | 96.3 | 78.7 | 86.1 | 44.6 | 68.5 | 55.3 | 85.5 | 63.4 | 76.9 | 32.7 | 31.6 | 41.9 |
| Reward fine-tuning | 82.9 | 100.0 | 98.8 | 83.8 | 86.6 | 50.5 | 77.7 | 59.5 | 88.9 | 69.4 | 82.0 | 40.6 | 31.9 | 44.2 |
| **Table 8** | | | | | | | | | | | | | | |
| Bidirectional | 70.4 | 99.7 | 88.6 | 68.1 | 80.0 | 25.2 | 60.8 | 51.4 | 83.0 | 61.7 | 73.4 | 21.6 | 31.3 | 37.7 |
| Causal | 68.3 | 99.5 | 83.9 | 67.3 | 79.0 | 21.9 | 58.5 | 50.60 | 81.7 | 61.5 | 72.3 | 19.3 | 31.4 | 37.2 |
| **Table 9** | | | | | | | | | | | | | | |
| Layer 6 | 68.4 | 99.0 | 85.6 | 64.9 | 79.2 | 23.6 | 58.2 | 51.2 | 81.2 | 61.6 | 74.0 | 22.4 | 31.1 | 37.2 |
| Layer 12 | 68.3 | 99.5 | 87.0 | 66.0 | 80.0 | 20.9 | 56.7 | 50.5 | 81.7 | 61.0 | 73.2 | 18.6 | 31.2 | 37.5 |
| Layer 18 | 69.1 | 99.2 | 89.1 | 66.7 | 79.5 | 22.1 | 58.2 | 51.2 | 82.4 | 61.5 | 74.3 | 20.5 | 31.0 | 37.4 |
| Layer 23 | 70.6 | 99.3 | 88.9 | 69.7 | 80.9 | 24.1 | 60.8 | 50.9 | 82.1 | 61.1 | 73.7 | 19.5 | 31.3 | 37.7 |
| Layer 24 | 70.4 | 99.7 | 88.6 | 68.1 | 80.0 | 25.2 | 60.8 | 51.4 | 83.0 | 61.7 | 73.4 | 21.6 | 31.3 | 37.7 |
| Layer 6 + 12 + 8 + 24 | 70.4 | 99.6 | 91.0 | 67.9 | 79.6 | 26.1 | 58.5 | 51.2 | 81.8 | 61.3 | 74.1 | 21.0 | 31.2 | 37.5 |
| **Table 10** | | | | | | | | | | | | | | |
| LLM | 70.2 | 99.4 | 87.9 | 67.4 | 81.2 | 24.9 | 60.1 | 51.4 | 83.0 | 61.0 | 74.2 | 21.5 | 31.2 | 37.8 |
| MLLM | 69.6 | 99.5 | 88.5 | 67.1 | 80.9 | 25.2 | 56.6 | 51.3 | 82.0 | 62.2 | 73.6 | 20.7 | 31.4 | 37.8 |
| **Table 11** | | | | | | | | | | | | | | |
| Layer 14 | 68.1 | 98.9 | 82.0 | 59.2 | 80.7 | 27.3 | 60.6 | 49.4 | 80.4 | 59.3 | 68.7 | 19.8 | 31.3 | 37.1 |
| Layer 28 | 71.7 | 99.0 | 87.6 | 62.6 | 82.6 | 34.3 | 64.0 | 50.2 | 81.1 | 60.0 | 72.0 | 19.9 | 31.1 | 37.1 |
| Layer 42 | 71.6 | 98.4 | 86.1 | 63.3 | 82.8 | 37.1 | 61.8 | 50.1 | 79.5 | 60.0 | 72.0 | 20.7 | 31.0 | 37.6 |
| Layer 56 (last) | 70.2 | 99.4 | 88.0 | 67.4 | 81.2 | 24.9 | 60.1 | 51.4 | 83.0 | 61.0 | 74.2 | 21.5 | 31.2 | 37.8 |
| Layer 28 + 42 | 73.0 | 99.2 | 87.7 | 66.3 | 85.0 | 33.0 | 66.6 | 50.7 | 81.1 | 61.1 | 73.0 | 20.3 | 31.1 | 37.7 |
| **Table 12** | | | | | | | | | | | | | | |
| A | 69.8 | 99.2 | 88.5 | 69.9 | 79.1 | 25.7 | 55.9 | 50.4 | 80.3 | 61.8 | 72.9 | 19.3 | 31.3 | 37.0 |
| B | 69.4 | 99.6 | 89.9 | 67.4 | 79.8 | 23.0 | 56.9 | 50.7 | 81.1 | 61.0 | 73.1 | 21.2 | 31.1 | 36.5 |
| C | 70.4 | 99.7 | 88.6 | 68.1 | 80.0 | 25.2 | 60.8 | 51.4 | 83.0 | 61.7 | 73.4 | 21.6 | 31.3 | 37.7 |