

# PIXART- $\Sigma$ : Weak-to-Strong Training of Diffusion Transformer for 4K Text-to-Image Generation

Junsong Chen<sup>1,2,3\*</sup>, Chongjian Ge<sup>1,3\*</sup>, Enze Xie<sup>1\*†</sup>, Yue Wu<sup>1\*</sup>, Lewei Yao<sup>1,4</sup>, Xiaozhe Ren<sup>1</sup>, Zhongdao Wang<sup>1</sup>, Ping Luo<sup>3</sup>, Huchuan Lu<sup>2</sup>, and Zhenguo Li<sup>1</sup>

<sup>1</sup> Huawei Noah’s Ark Lab   <sup>2</sup> Dalian University of Technology   <sup>3</sup> HKU   <sup>4</sup> HKUST  
Project Page: <https://pixart-alpha.github.io/PixArt-sigma-project/>

**Abstract.** In this paper, we introduce PIXART- $\Sigma$ , a Diffusion Transformer model (DiT) capable of directly generating images at 4K resolution. PIXART- $\Sigma$  represents a significant advancement over its predecessor, PIXART- $\alpha$ , offering images of markedly higher fidelity and improved alignment with text prompts. A key feature of PIXART- $\Sigma$  is its training efficiency. Leveraging the foundational pre-training of PIXART- $\alpha$ , it evolves from the ‘weaker’ baseline to a ‘stronger’ model via incorporating higher quality data, a process we term “weak-to-strong training”. The advancements in PIXART- $\Sigma$  are twofold: (1) High-Quality Training Data: PIXART- $\Sigma$  incorporates superior-quality image data, paired with more precise and detailed image captions. (2) Efficient Token Compression: we propose a novel attention module within the DiT framework that compresses both keys and values, significantly improving efficiency and facilitating ultra-high-resolution image generation. Thanks to these improvements, PIXART- $\Sigma$  achieves superior image quality and user prompt adherence capabilities with significantly smaller model size (0.6B parameters) than existing text-to-image diffusion models, such as SDXL (2.6B parameters) and SD Cascade (5.1B parameters). Moreover, PIXART- $\Sigma$ ’s capability to generate 4K images supports the creation of high-resolution posters and wallpapers, efficiently bolstering the production of high-quality visual content in industries such as film and gaming.

**Keywords:** T2I Synthesis, Diffusion Transformer, Efficient Model

## 1 Introduction

The recent emergence of high-quality Text-to-Image (T2I) models has profoundly impacted the AI Generated Content (AIGC) community. This includes both proprietary models such as DALL-E 3 [31], Midjourney [29], as well as open-source models like Stable Diffusion [36] and PIXART- $\alpha$  [4]. Nonetheless, developing a top-tier T2I model involves considerable resources; for instance, training SD1.5 from scratch necessitates about 6000 A100 GPU days [36], posing a substantial barrier to individual researchers with limited resources and impeding innovation

---

\*Equal contribution. Work done during the students’ internships at Huawei Noah’s Ark Lab.

†Project lead and corresponding author. Thanks to Charan for the 4K dataset collection.



**Fig. 1: Images generated by PIXART-Σ.** The model can output photo-realistic, high aesthetic, extreme aspect ratio, multi-style images, and follow user instructions.

within the AIGC community. Over time, the AIGC community will gain access to continuously updated, higher-quality datasets and more advanced algorithms. A pivotal question is: *how can we efficiently integrate these new elements into an existing model, achieving a more powerful version within the constraints of limited resources?*

To explore this issue, our research focuses on enhancing PIXART- $\alpha$ , an efficient T2I training method. PIXART- $\alpha$  represents an early venture within the DiT framework, a model structure with significant potential, as evidenced by works such as GenTron [6], Sora [33] and Stable Diffusion 3 [42]. To maximize this potential, we build upon the pre-trained foundation of PIXART- $\alpha$ , integrating advanced elements to facilitate its continuous improvement, resulting in a more powerful model, PIXART- $\Sigma$ . We refer to this process of evolving from a relatively weaker baseline to a stronger model through efficient training as “weak-to-strong training”. Specifically, to achieve “weak-to-strong training”, we introduce the following enhancements:

**Higher-Quality Training Data:** We collect a high-quality dataset superior to that used in PIXART- $\alpha$ , focusing on two key aspects: (i) **High-quality images:** The dataset comprises 33M high-resolution images sourced from the Internet, all exceeding 1K resolution, including 2.3M images with resolutions around 4K. These images are predominantly characterized by their high aesthetic and encompass a wide range of artistic styles. (ii) **Dense and accurate captions:** To provide more precise and detailed captions for the aforementioned images, we replace the LLaVA [22] used in PIXART- $\alpha$  with a more powerful image captioner, Share-Captioner [5]. Furthermore, to improve the model’s alignment capacity between the textual and visual concepts, we extend the token length of the text encoder (*i.e.*, Flan-T5 [10]) to approximately 300 words. We observe these improvements effectively eliminate the model’s tendency for hallucination, leading to higher-quality text-image alignment.

**Efficient Token Compression:** To enhance PIXART- $\alpha$ , we expand its generation resolution from 1K to 4K. Generating images at ultra-high resolutions (*e.g.*, 2K/4K) introduces a significant increase in the number of tokens, leading to a substantial rise in computational demand. To address this challenge, we introduced a self-attention module with key and value token compression tailored to the DiT framework. Specifically, we utilize group convolutions with a stride of 2 for local aggregation of keys and values. Additionally, we employ a specialized weight initialization scheme, allowing for a smooth adaptation from a pre-trained model without KV compression. This design effectively reduces training and inference time by  $\sim 34\%$  for high-resolution image generation.

**Weak-to-Strong Training Strategy:** we propose several fine-tuning techniques to rapidly adapt from a weak model to a strong model efficiently. That includes (1) replacing with a more powerful Variational Autoencoder (VAE) [36], (2) scaling from low to high resolution, and (3) evolving from a model without Key-Value (KV) compression to one with KV compression. These outcomes confirm the validity and effectiveness of the “weak-to-strong training” approach.



A stylish woman walks down a Tokyo street filled with warm glowing neon and animated city signage. She wears **a black leather jacket a long red dress** and black boots, and carries a black purse. She wears **sunglasses and red lipstick**. She walks confidently and casually. The street is damp and reflective, creating a **mirror effect of the colorful lights**. Many pedestrians walk about.



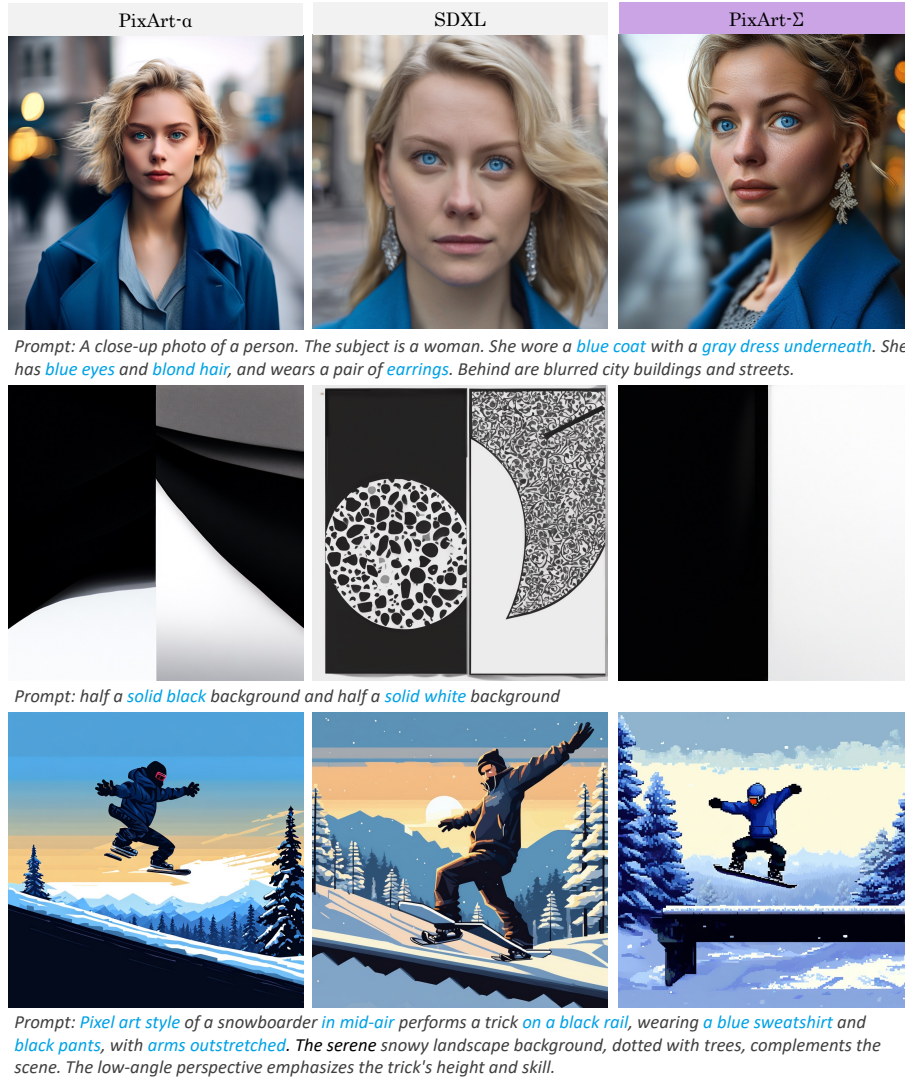
**Fig. 2: 4K image generation with complex dense instructions.** PIXART- $\Sigma$  can directly generate 4K resolution images without post-processing, and accurately respond to the given prompt.

Through the proposed improvements, PIXART- $\Sigma$  achieves high-quality 4K resolution image generation at a minimal training cost and model parameters. Specifically, fine-tuning from a pre-trained model, we additionally utilize only **9%** of the GPU days required by PIXART- $\alpha$  to achieve a strong 1K high-resolution image generation model, which is impressive considering replacing with new training data and a more powerful VAE. Moreover, we use only **0.6B** parameters while SDXL [36] and SD Cascade [35] use 2.6B and 5.1B parameters respectively. Images generated by PIXART- $\Sigma$  possess an aesthetic quality comparable to current top-tier T2I products, such as DALL-E 3 [31] and MJV6 [29] (as in Fig. 4). Additionally, PIXART- $\Sigma$  also demonstrates exceptional capability for fine-grained alignment with textual prompts (as shown in Fig. 2 and 3).

## 2 Related Work

**Diffusion Transformers.** The Transformer architecture has achieved remarkable success across various domains, such as language modeling [37,38], computer vision [23, 43, 51, 53], and other areas [3, 14]. In the realm of diffusion models, DiT [34] and UViT [2] pioneer the use of Transformer architecture. Subsequent works, including DiffT [15], SiT [28], and FiT [26], have improved upon DiT’s architecture, while [12, 52] enhance training efficiency through masked modeling techniques. For Text-to-Image (T2I) synthesis, PIXART- $\alpha$  [4] explore efficient T2I training schemes, achieving the first Transformer-based T2I model capable of generating 1024px high-quality images. GenTron [6] explores the flexibility and scalability of diffusion Transformers in both image and video generation. The re-





**Fig. 3: Comparison of PIXART- $\Sigma$  with open-source models, e.g., PIXART- $\alpha$  and SDXL:** Compared with PIXART- $\alpha$ , PIXART- $\Sigma$  improves the realism of portraits and the capability of semantic analysis. Compared with SDXL, our method has a better ability to follow user instructions. The keywords are highlighted as *blue*.

cent advent of the powerful video generation model Sora [33] has further underscored the potential of Diffusion Transformers. In this work, for the first time, we explore using the Transformer architecture to generate 4K ultra-high-resolution images directly, tackling the computational complexity challenges posed by involving long-sequence tokens.

**High Resolution Image Generation** greatly enhances visual quality and is important in various industries such as film and gaming. However, increasing





**Fig. 4: Compare PIXART- $\Sigma$  and four other T2I products:** Firefly 2, Imagen 2, Dalle 3, and Midjourney 6. Images generated by PIXART- $\Sigma$  are very competitive with these commercial products.

image resolution introduces challenges due to the substantial increase in computational demands. Numerous methods have been explored in this direction. For instance, Imagen [40], GigaGAN [18] and Stable Diffusion [39] introduce an additional super-resolution network, while Stable Cascade [35] employs multiple diffusion networks to increase resolution progressively. These combined-model solutions, however, can introduce cumulative errors. On the other hand, works like SDXL [36], DALL·E 2 [30], Playground [19] and PIXART- $\alpha$  [4] have attempted to generate high-resolution images using diffusion models directly. Nevertheless, these efforts are capped at generating images with resolutions up to 1024px due to computational complexity. In this paper, we push this boundary to 4K resolution, significantly enhancing the visual quality of the generated content.

**Efficient Transformer architecture.** The self-attention mechanism in Transformer suffers from quadratic computational complexity with respect to the number of tokens, which hinders the scaling up of token quantity. Many works have sought improvements in this area: (1) Sparse Attention [7, 8, 45, 46, 48], which reduces the overall computational load by selectively processing a subset of tokens. For instance, PVT v2 [46] employs a convolutional kernel to condense the space of the key and value, thus lowering the complexity involved in computing the attention. (2) Local Attention [13, 23, 47, 54] focuses attention within nearby regions; notably, Swin Transformer [23] utilizes window-based attention to limit computations to within a specified window size. (3) Low-rank/Linear Attention [9, 25, 44]. The Linformer [44] reduces the computational complexity of the self-attention mechanism through low-rank approximations. In this paper, inspired by PVT v2 [46], we employ a self-attention mechanism based on key/value compression to mitigate the high complexity of processing 4K images.

### 3 Framework

#### 3.1 Data Analysis

	<b>LLaVA</b> <b>Hallucinations</b>	The image features a large, ornate church with a tall, <b>pointed roof</b> and a large <b>stained-glass window</b> . The church is situated in a city, and it has a white and gray color scheme. The style of the church is Gothic, which is evident from the pointed roof and the intricate details of the <b>stained-glass window</b> . The presence of statues and a <b>clock</b> on the building further emphasizes the grand and historical nature of the structure.
	<b>Share-Captioner</b> <b>Correctness</b>	The image captures the grandeur of a cathedral, painted in a pristine white, stands majestically against the backdrop of a <b>clear blue sky</b> . The entrance to the cathedral is adorned with a <b>large rose window</b> . Flanking the entrance are <b>statues of saints</b> , standing as silent guardians of the sacred space within. The perspective of the photo, <b>taken from a low angle</b> . The <b>lush green trees</b> in the background adding serenity to the scene.
	<b>LLaVA</b> <b>Hallucinations</b>	The image features a <b>woman and a man</b> sitting on a brick walkway near a body of water, which could be a river or a lake. They are both wearing head coverings, and the <b>woman is holding a handbag</b> . The scene is set during the day, with the <b>sun shining brightly</b> , creating a warm and inviting atmosphere. The style of the image is a <b>black and white photo</b> , which adds a <b>timeless and classic</b> feel to the scene.
	<b>Share-Captioner</b> <b>Correctness</b>	The image captures a serene scene at a <b>harbor</b> . <b>Two individuals</b> are seated on a bench, their backs to the camera, engrossed in the view of the water. The water, a deep shade of blue, is dotted with <b>boats of various sizes and colors</b> , including a <b>white boat with a green stripe</b> and a <b>red boat</b> . The sky above is a light blue.

**Fig. 5: Comparative illustration of hallucinations:** Contrasting differences in hallucination occurrences between LLaVA and Share-Captioner, with **red** indicating hallucinations and **green** denoting correctness.

**Higher Aesthetic and higher Resolution.** To enhance the aesthetic quality of our dataset, we expand our internal data from 14M to 33M. For clarity, we name the two datasets Internal- $\alpha$  and Internal- $\Sigma$ , respectively. Note that this expansion still falls short compared to the vast images utilized by currently available open-source models like SD v1.5, which uses 2B data. We demonstrate that effective training strategies with limited data amount can still obtain a strong T2I model.

The images within Internal- $\Sigma$  are above 1K resolution. To facilitate 4K resolution generation, we additionally collect a dataset of 8M real photographic images at 4K resolution. To ensure aesthetic quality, we employ an aesthetic scoring model (AES) [1] to filter these 4K images. This process yields a highly refined dataset of 2M ultra-high-resolution and high-quality images.

Interestingly, we have observed that as the resolution of the images increases, there is an improvement in the model’s fidelity (Fréchet Inception Dis-



**Table 1: Statistics of noun concepts for different datasets.** VN: valid distinct nouns (appearing more than 10 times); DN: total distinct nouns; **Average**: average noun count per image; **ACL**: Average Caption length.

Dataset	Volume	Caption	VN/DN	Total Noun	ACL	Average
Internal- $\alpha$	14M	Raw	187K/931K	175M	25	11.7/Img
Internal- $\alpha$	14M	LLaVA	28K/215K	536M	98	29.3/Img
Internal- $\alpha$	14M	Share-Captioner	51K/420K	815M	184	54.4/Img
Internal- $\Sigma$	33M	Raw	294K/1512K	485M	35	14.4/Img
Internal- $\Sigma$	33M	Share-Captioner	77K/714K	1804M	180	53.6/Img
4K- $\Sigma$	2.3M	Share-Captioner	24K/96K	115M	163	49.5/Img

tance (FID) [17]) and semantic alignment (CLIP Score), which underscores the importance of the capabilities of generating high-resolution images.

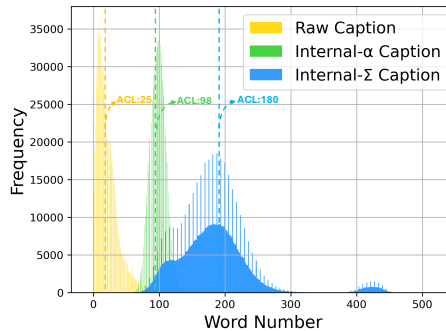
#### Better Text-Image Alignment.

Recent works such as PIXART- $\alpha$  [4] and DALL-E 3 [31] emphasize the significance of text-image alignment. Strengthening this alignment is crucial for boosting model capabilities. To refine our collected “raw” descriptions further, we focus on improving both the length and accuracy of our captions. Notably, our captions (Internal- $\Sigma$ ) show several advantages over the one used in PIXART- $\alpha$  (Internal- $\alpha$ ) in the following aspects:

1. **Enhanced caption accuracy:** As depicted in Fig. 5, LLaVa used in PIXART- $\alpha$  encounters certain hallucination. We leverage a more powerful model, i.e., Share-Captioner [5], to generate detailed and correct captions, augmenting the collected raw prompts.
2. **Increased caption length:** As shown in Tab. 1 and Fig. 6, the average caption length increased significantly to 180 words, highly enhancing the descriptive power of the captions. Additionally, we extend the token processing length of the text encoder from 120 tokens (as in Internal- $\alpha$ ) to 300 tokens. Our model is trained on a mix of long (Share-Captioner) and short (raw) captions with a ratio of 60% and 40%, respectively. This approach enhances the diversity of textual descriptions and mitigates potential biases that might arise from solely relying on generative captions.

Tab. 1 demonstrates a summary for both Internal- $\alpha$  and - $\Sigma$ , where we assess the diversity of the datasets through various metrics, including the noun variety, total noun count, average caption length, and average nouns per image.

**High-Quality Evaluation Dataset.** Most SoTA T2I models chose MSCOCO [21] as the evaluation set to assess the FID and CLIP Scores. However, we observe



**Fig. 6: Histogram Visualization of the Caption Length.** We randomly select 1M captions from the raw captions, Internal- $\alpha$ , and Internal- $\Sigma$  to draw the corresponding histogram. ACL denotes the average caption length.

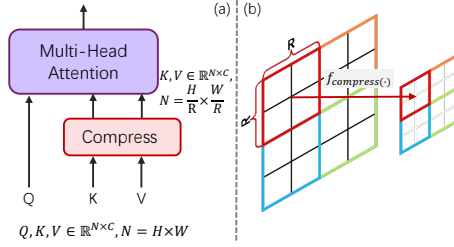
evaluations conducted on the MSCOCO dataset may not adequately reflect a model’s capabilities in aesthetics and text-image alignment. Therefore, we propose a curated set comprising 30,000 high-quality, aesthetically pleasing text-image pairs to facilitate the assessment. The selected samples of the dataset are presented in the appendix. This dataset is designed to provide a more comprehensive evaluation of a model’s performance, particularly in capturing the intricacies of aesthetic appeal and the fidelity of alignment between textual descriptions and visual content. Unless otherwise specified, the evaluation experiments in the paper are conducted in the collected High-Quality Evaluation Dataset.

### 3.2 Efficient DiT Design

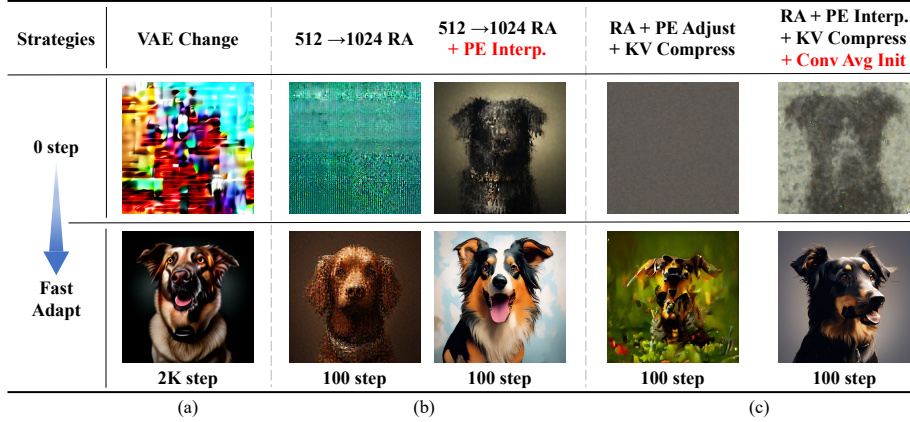
An efficient DiT network is essential since the computational demand significantly increases when generating images at ultra-high resolutions. The attention mechanism plays a pivotal role in the efficacy of Diffusion Transformers, yet its quadratic computational demands significantly limit model scalability, particularly at higher resolutions *e.g.*, 2K and 4K. Inspired by PVT v2 [46], we incorporate KV compression within the original PIXART- $\alpha$ ’s framework to address the computational challenges. This design adds a mere 0.018% to the total parameters yet achieves efficient reduction in computational costs via token compression, while still preserving both spatial and semantic information.

**Key-Value (KV) Token Compression.** Our motivation stems from the intriguing observation that applying key-value (KV) token compression directly to the pre-trained PIXART- $\alpha$  can still generate reasonable images. This suggests a redundancy in the features. Considering the high similarity within adjacent  $R \times R$  patches, we assume that feature semantics within a window are redundant and can be compressed reasonably. We propose KV token compression, which is denoted as  $f_c(\cdot)$ , to compress token features within a  $R \times R$  window through a compression operator, as depicted in Fig. 7.

Furthermore, to mitigate the potential information loss caused by KV compression in self-attention computation, we opt to retain all the tokens of queries (Q). This strategy allows us to utilize KV compression effectively while mitigating the risk of losing crucial information. By employing KV compression, we enhance the efficiency of attention computations and reduce the computation complexity from  $O(N^2)$  to  $O\left(\frac{N^2}{R^2}\right)$ , thereby making the computational cost of directly generating high-resolution images manageable.



**Fig. 7: Design of KV Token Compression.** We merge KV tokens in spatial space to reduce the computation complexity.



**Fig. 8:** This illustration demonstrates how our training strategy designs expedite the model’s convergence during the transition to VAEs, adjustment to higher resolutions, and the KV compression process, facilitating rapid learning from weak to strong.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{Q \cdot f_c(K)^T}{\sqrt{d_k}}\right) f_c(V) \quad (1)$$

We compress deep layers using the convolution operator “Conv2 $\times$ 2” with specific initialization. Detailed experiments on other design variants are discussed in Sec. 5. Specifically, we design a specialized convolution kernel initialization “Conv Avg Init” that utilizes group convolution and initializes the weights  $w = \frac{1}{R^2}$ , equivalent to an average operator. This initialization strategy can initially produce coarse results, accelerating the fine-tuning process while only introducing 0.018% additional parameters.

### 3.3 Weak-to-Strong Training Strategy

We propose several efficient training strategies to enhance the transition from a “weak” model to a “strong” model. These strategies encompass VAE rapid adaptation, high-resolution fine-tuning, and KV Token compression.

#### Adapting model to new VAEs.

As VAEs develop, training T2I models from scratch is resource-intensive. We replace PIXART- $\alpha$ ’s VAE with SDXL’s VAE and continue fine-tuning the diffusion model. We observe a rapid convergence phenomenon that fine-tuning quickly converges at 2K training steps as shown in Fig 8 (a). Fine-tuning is more efficient when dealing with VAE model transferring and negates the necessity of training from scratch.

**Table 2:** We fine-tune a high-resolution model from a low-resolution model and observe that even fine-tuning for a relatively short duration, such as 1K steps, can still yield high-quality results.

Resolution	Iterations	FID $\downarrow$	CLIP $\uparrow$
256	20K	16.56	0.270
256 $\rightarrow$ 512	1K	9.75	0.272
256 $\rightarrow$ 512	100K	8.91	0.276

#### Adapting to Higher-Resolution.

When we fine-tune from a low-resolution (LR) model to a high-resolution (HR) model, we observe a performance degradation in Fig. 8 (b), which we attribute to discrepancies in positional embeddings



(PE) between different resolutions. To address this, we utilize the “PE Interpolation” trick [4, 49]: initializing the HR model’s PE by interpolating the LR model’s PE, significantly enhancing the HR model’s initial status and expediting fine-tuning. We can obtain visually pleasing images even within only 100 training iterations. Besides, we quantitatively compare the model’s performance as illustrated in Tab. 2. The fine-tuning quickly converges at 1K steps, and further training slightly improves the performance. This illustrates that using the “PE Interpolation” trick enables rapid convergence of higher resolution generation, obviating the need for training from scratch for generating at higher resolutions.

**Adapting model to KV compression.** We can use KV compression directly when fine-tuning from LR pre-trained models without KV compression. As shown in Fig. 8 (c), with our “Conv Avg Init.” strategy, PIXART- $\Sigma$  starts from a better initial state, making converging easier and faster. Notably, PIXART- $\Sigma$  performs satisfied visual results even within 100 training steps. Finally, through the KV compression operators and compression layers design in Sec 3.2, we can reduce  $\sim 34\%$  of the training and inference time.

## 4 Experiment

### 4.1 Implementation Details

**Training Details.** We follow Imagen [40] and PIXART- $\alpha$  [4] to employ the T5 [10]’s encoder (*i.e.*, Flan-T5-XXL) as the text encoder for conditional feature extraction, and use PIXART- $\alpha$  [4] as our base diffusion model. Unlike most works that extract fixed 77 text tokens, we adjust the length of text tokens from PIXART- $\alpha$ ’s 120 to 300, as the caption curated in Internal- $\Sigma$  is much denser to provide highly fine-grained details. To capture the latent features of input images, we employ a pre-trained and frozen VAE from SDXL [36]. Other implementation details are the same as PIXART- $\alpha$ . Models are finetuned on the PIXART- $\alpha$ ’s 256px pre-trained checkpoint with the position embedding interpolation trick [4]. Our final models, including 1K resolution, are trained on 32 V100 GPUs. We additionally use 16 A100 GPUs to train the 2K and 4K image generation models. For further information, please refer to the appendix.

Note that we use CAME optimizer [27] with a weight decay of 0 and a constant learning rate of  $2e-5$ , instead of the regular AdamW [24] optimizer. This helps us reduce the dimension of the optimizer’s state, leading to lower GPU memory without performance degradation.

**Evaluation Metrics.** To better illustrate aesthetics and semantic ability, we collect 30K high-quality text-image pairs (as mentioned in Sec. 3.1) to benchmark T2I models. We mainly evaluate PIXART- $\Sigma$  via human and AI preference study since FID [39] may not adequately reflect the generation quality. However, we still provide the FID results on the collected dataset in the appendix.

### 4.2 Performance Comparisons

**Image Quality Assessment.** We qualitatively evaluated our methodology against both closed-source text-to-image products and open-source models. As

illustrated in Fig. 1, our model can produce high-quality, photo-realistic images with intricate details over diverse aspect ratios and styles. This capability underscores the superior performance of our approach in generating visually compelling content from textual descriptions. In Fig. 3, we compare PIXART- $\Sigma$  with open-source models SDXL [36] and PIXART- $\alpha$  [4], our method enhances the realism of portraits and boosts the capacity for semantic analysis. In contrast to SDXL, our approach demonstrates superior proficiency in following instructions.

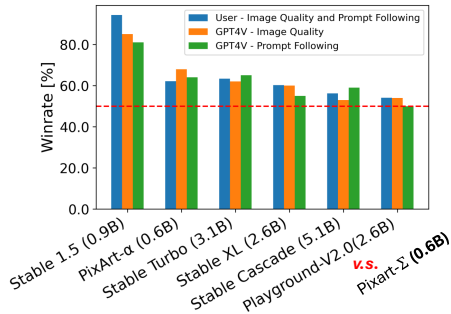
Not only superior to open-source models, but our method is also very competitive with current T2I closed-source products, as depicted in Fig. 4. PIXART- $\Sigma$  produces photo-realistic results and adheres closely to user instructions, which is on par with contemporary commercial products.

### High-resolution Generation.

Our method is capable of directly generating 4K resolution images without the need for any post-processing. Additionally, it excels in accurately following complex, detailed, and long text provided by users, as demonstrated in Fig. 2. Thus, users do not need prompt engineering to achieve satisfactory results.

Our approach enables direct 4K image generation. In parallel, studies [11, 16] have introduced tuning-free post-processing techniques aimed at generating HR images from LR models or employing super-resolution models [50] to produce HR images. However, their corresponding results often exhibit artifacts for two primary reasons: (1) Accumulative error may arise due to the cascade pipeline. (2) These methods do not capture the true distribution of 4K images nor learn the alignment between text and 4K images. We argue that our method might be a more promising way to generate high-resolution images. Our method yields superior results, and more visual comparison is included in the supplement.

**Human/AI (GPT4V) Preference Study.** We evaluate the well-trained model in both the human and AI preference study using a subset of 300 captions randomly collected from the High-Quality Evaluation Dataset mentioned in Sec. 3.1. We collect images generated by overall six open-source models, including PIXART- $\alpha$ , PIXART- $\Sigma$ , SD1.5 [39], Stable Turbo [41], Stable XL [36], Stable Cascade [35] and Playground-V2.0 [20]. We develop a website for the human preference study to display the prompts and their corresponding images. This website was distributed to trained evaluators, who were asked to assess the images, ranking them according to quality and how well they matched the text prompts. The results, illustrated by the blue bar in Fig. 9, indicate a marked preference for PIXART- $\Sigma$  over the other six T2I generators. PIXART- $\Sigma$  gener-



**Fig. 9: Comparisons on FID and Clip-Score with Open-sourced T2I Models.** PixArt- $\Sigma$  demonstrates enhanced performance in terms of FID and Clip-Score on the curated 30K High-Quality Evaluation Dataset.

Layers	FID ↓ CLIP-Score ↑	
N/A	8.244	0.276
Shallow (1-14)	9.278	0.275
Middle (7-20)	9.063	0.276
Deep (14-27)	8.532	0.275

(a) Compression layers.

Operator	FID ↓ CLIP-Score ↑	
N/A	8.244	0.276
Token Discarding	8.918	0.275
Token Pooling	9.415	0.275
Conv2×2	8.505	0.274

(b) Compression operators.

Res.	Ratio	FID ↓ CLIP-Score ↑		Train Latency ↓
512	1	8.244	0.276	2.3
512	2	9.063	0.276	2.2 (-4%)
512	4	9.606	0.276	2.1 (-9%)
1024	1	5.685	0.277	27.5
1024	2	5.512	0.273	22.5 (-18%)
1024	4	5.644	0.276	20.0 (-27%)
1024	9	5.712	0.275	17.8 (-35%)

(c) Compression ratios on different resolutions.

Res.	Ratio	Train Latency ↓		Test Latency ↓
		(s/Iter@32BS)		(s/Img)
2K	1	56		58
2K	4	37 (-34%)		38 (-34%)
4K	1	191		91
4K	4	125 (-35%)		60 (-34%)

(d) Speed of different resolutions.

**Table 3: KV-Token Compression Settings in Image Generation.** This study employs FID, CMMD, and CLIP-Score metrics to assess the impact of various token compression components, such as compression ratio, positions, operators, and varying resolutions. Speed calculation in Tab. 3c is Second/Iteration/384 Batch-size.

ates superior high-quality images that closely follow user prompts, using a much smaller size (0.6B parameters) compared to existing T2I diffusion models like SDXL (2.6B parameters) and SD Cascade (5.1B parameters).

For AI preference study, we employ the advanced multimodal model, GPT-4 V [32], as the evaluator. For each trial, we supply GPT-4V with two images: one from PIXART- $\Sigma$  and another from a competing T2I model. We craft distinct prompts guiding GPT-4V to vote based on image quality and image-and-text alignment. The results, represented by orange and green bars in Fig. 9, demonstrate consistent outcomes in both human and AI preference studies. Specifically, PIXART- $\Sigma$  surpasses the baseline, PIXART- $\alpha$ , in effectiveness. Compared to advanced models such as Stable Cascaded, PIXART- $\Sigma$  exhibits competitive or superior performance on image quality and instruction-following abilities.

## 5 Ablation Studies

We conduct ablation studies on generation performance on various KV compression designs. Unless specified, the experiments are conducted on 512px generation. The detailed settings of each ablation trial are included in the appendix.

### 5.1 Experimental settings

We use the test set described in Sec. 3.1 for evaluation. We employ FID to compute the distributional difference between the collected and generated data for comparative metrics. Furthermore, we utilize CLIP-Score to assess the alignment between prompts and the generated images.



## 5.2 Compression Designs

**Compression positions.** We implemented KV compression at different depths within the Transformer structure: in the shallow layers (1~14), the intermediate layers (7~20), and the deep layers (14~27). As indicated in Tab. 3a, employing KV compression on deep layers notably achieves superior performance. We speculate this is because shallow layers typically encode detailed texture content, while deep layers abstract high-level semantic content. Since compression affects image quality rather than semantic information, compressing deep layers achieves the least information loss, making it practical for accelerating training without compromising generation quality.

**Compression operators.** We explored the impact of different compression operators. We employed three techniques, random discarding, average pooling, and parametric convolution, to compress  $2 \times 2$  tokens into a single token. As illustrated in Table 3b, the “Conv  $2 \times 2$ ” method outperforms the others, underscoring the advantage of using a learnable kernel to more effectively reduce redundant features than simple discarding methods.

**Compression ratios on different resolutions.** We investigated the influence of varying compression ratios on different resolutions. As shown in Tab. 3c, remarkably, we find that token compression does not affect the alignment between textual and generated images (CLIP Score) but influences the image quality (FID) across resolutions. Although there is a slight degradation in image quality with increasing compression ratios, our strategy brings a training speedup of 18% to 35%. This suggests that our proposed KV compression is both effective and efficient for achieving high-resolution T2I generation.

**Speed comparisons on different resolutions.** We further comprehensively validate the speed acceleration in both training and inference in Tab. 3d. By ‘latency’, we refer to the duration required for inference on a single sample. Our method speeds up training and inference by about 35% in the 4K generation. Notably, training acceleration increases with resolution, accelerating from 18% to 35% as the resolution increases from 1K to 4K. This indicates the effectiveness of our method with increasing resolution, demonstrating its potential applicability to even higher-resolution image generation tasks.

## 6 Conclusion

In this paper, we introduce PIXART- $\Sigma$ , a Text-to-Image (T2I) diffusion model capable of directly generating high-quality images at 4K resolution. Building upon the pre-trained foundation of PIXART- $\alpha$ , PIXART- $\Sigma$  achieves efficient training through a novel “weak-to-strong training” methodology. This approach is characterized by the incorporation of higher-quality data and the integration of efficient token compression. PIXART- $\Sigma$  excels at producing high-fidelity images while adhering closely to textual prompts, surpassing the high standards set by its predecessor, PIXART- $\alpha$ . We believe that the innovations presented in PIXART- $\Sigma$  will not only contribute to advancements in the AIGC community but also pave the way for entities to access more efficient, and high-quality generative models.

**Acknowledgements.** We would like to thank Zeqiang Lai, Fei Chen, and Shuchen Xue for discussing and helping implement the DMD part.

## References

1. Aesthetic predictor (2023), <https://github.com/christophschuhmann/improved-aesthetic-predictor> 7
2. Bao, F., Nie, S., Xue, K., Cao, Y., Li, C., Su, H., Zhu, J.: All are worth words: A vit backbone for diffusion models. In: CVPR (2023) 4
3. Chandra, A., Tünnermann, L., Löfstedt, T., Gratz, R.: Transformer-based deep learning for predicting protein properties in the life sciences. eLife Sciences Publications, Ltd (2023) 4
4. Chen, J., Yu, J., Ge, C., Yao, L., Xie, E., Wu, Y., Wang, Z., Kwok, J., Luo, P., Lu, H., Li, Z.: PixArt- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In: ICLR (2024) 1, 4, 6, 8, 11, 12
5. Chen, L., Li, J., Dong, X., Zhang, P., He, C., Wang, J., Zhao, F., Lin, D.: ShareGPT4V: Improving large multi-modal models with better captions. In: arXiv (2023) 3, 8
6. Chen, S., Xu, M., Ren, J., Cong, Y., He, S., Xie, Y., Sinha, A., Luo, P., Xiang, T., Perez-Rua, J.M.: Gentron: Delving deep into diffusion transformers for image and video generation. arXiv preprint arXiv:2312.04557 (2023) 3, 4
7. Chen, T., Cheng, Y., Gan, Z., Yuan, L., Zhang, L., Wang, Z.: Chasing sparsity in vision transformers: An end-to-end exploration. In: NeurIPS (2021) 7
8. Chen, X., Liu, Z., Tang, H., Yi, L., Zhao, H., Han, S.: SparseViT: Revisiting activation sparsity for efficient high-resolution vision transformer. In: CVPR (2023) 7
9. Choromanski, K., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., et al.: Rethinking attention with performers. In: ICLR (2021) 7
10. Chung, H.W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., et al.: Scaling instruction-finetuned language models. In: arXiv (2022) 3, 11
11. Du, R., Chang, D., Hospedales, T., Song, Y.Z., Ma, Z.: Demofusion: Democratising high-resolution image generation with no \$\$\$\$. In: CVPR (2024) 12
12. Gao, S., Zhou, P., Cheng, M.M., Yan, S.: Masked diffusion transformer is a strong image synthesizer. In: ICCV (2023) 4
13. Ge, C., Ding, X., Tong, Z., Yuan, L., Wang, J., Song, Y., Luo, P.: Advancing vision transformers with group-mix attention. In: arXiv (2023) 7
14. von Glehn, I., Spencer, J.S., Pfau, D.: A self-attention ansatz for ab-initio quantum chemistry. In: ICLR (2023) 4
15. Hatamizadeh, A., Song, J., Liu, G., Kautz, J., Vahdat, A.: Diffit: Diffusion vision transformers for image generation. In: arXiv (2023) 4
16. He, Y., Yang, S., Chen, H., Cun, X., Xia, M., Zhang, Y., Wang, X., He, R., Chen, Q., Shan, Y.: ScaleCrafter: Tuning-free higher-resolution visual generation with diffusion models. In: ICLR (2024) 12
17. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) 8

18. Kang, M., Zhu, J.Y., Zhang, R., Park, J., Shechtman, E., Paris, S., Park, T.: Scaling up gans for text-to-image synthesis. In: CVPR (2023) 6
19. Li, D., Kamko, A., Akhgari, E., Sabet, A., Xu, L., Doshi, S.: Playground v2.5: Three insights towards enhancing aesthetic quality in text-to-image generation. In: arXiv (2024) 6
20. Li, D., Kamko, A., Sabet, A., Akhgari, E., Xu, L., Doshi, S.: Playground v2, <https://huggingface.co/playgroundai/playground-v2-1024px-aesthetic> 12
21. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: ECCV (2014) 8
22. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: arXiv (2023) 3
23. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021) 4, 7
24. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: arXiv (2017) 11
25. Lu, J., Yao, J., Zhang, J., Zhu, X., Xu, H., Gao, W., Xu, C., Xiang, T., Zhang, L.: Soft: Softmax-free transformer with linear complexity. In: NeurIPS (2021) 7
26. Lu, Z., Wang, Z., Huang, D., Wu, C., Liu, X., Ouyang, W., Bai, L.: Fit: Flexible vision transformer for diffusion model. In: arXiv (2024) 4
27. Luo, Y., Ren, X., Zheng, Z., Jiang, Z., Jiang, X., You, Y.: CAME: Confidence-guided adaptive memory efficient optimization. In: ACL (2023) 11
28. Ma, N., Goldstein, M., Albergo, M.S., Boffi, N.M., Vanden-Eijnden, E., Xie, S.: SiT: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In: arXiv (2024) 4
29. Midjourney: Midjourney (2023), <https://www.midjourney.com> 1, 4
30. OpenAI: Dalle-2 (2023), <https://openai.com/dall-e-2> 6
31. OpenAI: Dalle-3 (2023), <https://openai.com/dall-e-3> 1, 4, 8
32. OpenAI: Gpt-4v(ision) system card. In: OpenAI (2023) 13
33. OpenAI: Sora (2024), <https://openai.com/sora> 3, 5
34. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: ICCV (2023) 4
35. Pernias, P., Rampas, D., Richter, M.L., Pal, C., Aubreville, M.: Würstchen: An efficient architecture for large-scale text-to-image diffusion models. In: ICLR (2023) 4, 6, 12
36. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: SDXL: Improving latent diffusion models for high-resolution image synthesis. In: arXiv (2023) 1, 3, 4, 6, 11, 12
37. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training. OpenAI blog (2018) 4
38. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog (2019) 4
39. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) 6, 11, 12
40. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. In: NeurIPS (2022) 6, 11
41. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. In: arXiv (2023) 12
42. Stability.AI: Stable diffusion 3 (2024), <https://stability.ai/news/stable-diffusion-3> 3



43. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: ICML (2021) 4
44. Wang, S., Li, B.Z., Khabsa, M., Fang, H., Ma, H.: Linformer: Self-attention with linear complexity. In: arXiv (2020) 7
45. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: ICCV (2021) 7
46. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pvt v2: Improved baselines with pyramid vision transformer. Computational Visual Media (2022) 7, 9
47. Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: Vision transformer with deformable attention. In: CVPR (2022) 7
48. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: SegFormer: Simple and efficient design for semantic segmentation with transformers. In: NeurIPS (2021) 7
49. Xie, E., Yao, L., Shi, H., Liu, Z., Zhou, D., Liu, Z., Li, J., Li, Z.: DiffFit: Unlocking transferability of large diffusion models via simple parameter-efficient fine-tuning. In: ICCV (2023) 11
50. Xue, Z., Song, G., Guo, Q., Liu, B., Zong, Z., Liu, Y., Luo, P.: Raphael: Text-to-image generation via large mixture of diffusion paths. In: NeurIPS (2023) 12
51. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.H., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token ViT: Training vision transformers from scratch on imagenet. In: ICCV (2021) 4
52. Zheng, H., Nie, W., Vahdat, A., Anandkumar, A.: Fast training of diffusion models with masked transformers. In: arXiv (2023) 4
53. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: CVPR (2021) 4
54. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable transformers for end-to-end object detection. ICLR (2021) 7