

# Inf-DiT: Upsampling Any-Resolution Image with Memory-Efficient Diffusion Transformer

Zhuoyi Yang<sup>1,\*</sup>, Heyang Jiang<sup>1</sup>, Wenyi Hong<sup>1</sup>, Jiayan Teng<sup>1</sup>  
Wendi Zheng<sup>1</sup>, Yuxiao Dong<sup>1</sup>, Ming Ding<sup>2</sup>, and Jie Tang<sup>1,\*\*</sup>

<sup>1</sup> Tsinghua University <sup>2</sup> Zhipu.AI



**Fig. 1:** Selected ultra-high-resolution upsampling examples of Inf-DiT, based on SDXL, DALL-E 3, and real images.

**Abstract.** Diffusion models have shown remarkable performance in image generation in recent years. However, due to a quadratic increase in memory during generating ultra-high-resolution images (e.g.  $4096 \times 4096$ ), the resolution of generated images is often limited to  $1024 \times 1024$ . In this work, we propose a unidirectional block attention mechanism that can adaptively adjust the memory overhead during the inference process and handle global dependencies. Building on this module, we adopt the DiT structure for upsampling and develop an infinite super-resolution model capable of upsampling images of various shapes and resolutions. Comprehensive experiments show that our model achieves SOTA performance in generating ultra-high-resolution images in both machine and

\* Done as intern at Zhipu AI, yangzy22@mails.tsinghua.edu.cn

\*\* Corresponding author jietang@tsinghua.edu.cn

human evaluation. Compared to commonly used UNet structures, our model can save more than  $5\times$  memory when generating  $4096 \times 4096$  images. The project URL is <https://github.com/THUDM/Inf-DiT>.


**Keywords:** diffusion model · ultra-high-resolution generation · super-resolution

## 1 Introduction

Recent years have witnessed rapid advancements in diffusion models, which significantly propelling the field of image generation and editing forward. Despite the advancements, a critical limitation persists: the resolution of images produced by existing image diffusion models is generally confined to  $1024 \times 1024$  pixels or lower, posing a significant challenge for generating ultra-high-resolution images, which are indispensable in various real-world applications including intricate design projects, advertising, and the creation of posters and wallpapers, etc.

A commonly used approach for generating high resolution is cascaded generation [12], which first produces a low-resolution image, then applies multiple upsampling models to increase the image’s resolution step by step. This approach breaks down the generation of high-resolution images into multiple tasks. Based on the results generated in the previous stage, the models in the later stages only need to perform local generation. Building upon the cascaded structure, both DALL-E2 [21] and Imagen [23] can effectively generate images with 1024 resolution.

The biggest challenge for upsampling to much higher resolution images is the significant GPU memory demands. For example, if utilizing the widely-adopted U-Net architecture such as SDXL [18] for image inference (see Figure 2), we observe a dramatic escalation in memory consumption with increasing resolution. Specifically, generating a  $4096 \times 4096$  resolution image, which comprises over 16 million pixels requires more than 80GB of memory, exceeding the capacities of standard RTX 4090 or A100 graphics cards. Furthermore, the process of training models for high-resolution image generation exacerbates these demands, as it necessitates additional memory for storing gradients, optimizer states, etc. LDM [21] reduces the memory consumption by utilizing Variational Autoencoder(VAE) to compress images and generating images in a smaller latent space. However, it is also emphasized that an excessively high compression ratio can



**Fig. 2:** Comparison of memory usage during inference at different resolutions between our model and the SDXL architecture.

substantially deteriorate the quality of generation, imposing severe limitations on the reduction of memory consumption.

We observe that one common and critical problem among the aforementioned models is the necessity to input the entire image to the model, which requires keeping  $O(N^2)$  hidden states in memory, where  $N$  is the width(height) of the image. Based on this observation, we propose a Unidirectional Block Attention (UniBA) algorithm that can dramatically reduce the space complexity of generation from  $O(N^2)$  to  $O(N)$ , increasing the highest available resolution for a large margin (Fig. 2). Specifically, in every diffusion step, we split the image into blocks and perform a sequential batch generation order among them, where each batch simultaneously produces a subset of the blocks, and any amount of blocks can be generated in parallel as long as the memory restriction allows. It is worth noting that, though the model only inputs part of image into the model at the same time, UniBA successfully preserves the ability to interact with hidden states of faraway blocks and maintain high-level semantic consistency. This is different from other block-based generation methods such as [2, 14, 29] that only interact with blocks in pixel space or compressed latent space.

Based on this algorithm, we optimize diffusion transformer (DiT) and train a model named Inf-DiT, which is capable of upsampling images of varying resolution and shape. Furthermore, we design several techniques including providing global image embedding to enhance the global semantic consistency and offer zero-shot text control ability, and provision of all neighboring low-resolution(LR) blocks through cross-attention mechanisms to further enhance the local consistency. Evaluation results show that Inf-DiT achieved significantly superior results compared to other high-resolution generative models in both machine and human evaluation.

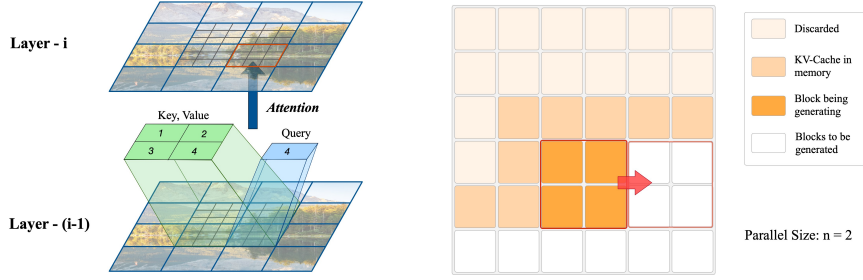
To summarize, our main contributions are as follows:

1. We propose Unidirectional Block Attention (UniBA) algorithm, which reduces the minimum memory consumption from  $O(N^2)$  to  $O(N)$  during inference, where  $N$  represents the edge length. This mechanism is also capable of adapting to various memory restrictions by adjusting the number of blocks generated in parallel, trading off between memory and time overhead.
2. Based on these methods, we train an image upsampling diffusion model, Inf-DiT, a 700M model capable of upsampling images of varying resolutions and shapes. Inf-DiT achieves state-of-the-art performance in both machine (HPDv2 and DIV2K datasets) and human evaluation.
3. We design multiple techniques to further enhance local and global consistency, and offer a zero-shot ability for flexible text control.

## 2 Methodology

### 2.1 Unidirectional Block Attention (UniBA)

We observe that the critical obstacle for generating ultra-high resolution images is memory limitation. As the resolution of the image increases, the corresponding



**Fig. 3:** **Left:** Unidirectional block attention. In our implementation, each block directly depends on the three blocks in each layer: the blocks in the upper left corner, on the left, and on the top. **Right:** The inference process of Inf-DiT. Inf-DiT generates  $n \times n$  block each time, based on the memory size. During this process, only the KV-cache of the blocks that subsequent blocks depend on are stored in memory.

hidden state’s size within the network expands quadratically. For example, a single hidden state with the shape of  $2048 \times 2048 \times 1280$  in only one layer requires 20GB of memory, making it formidable to generate very large images. How to avoid storing the entire image’s hidden state in memory becomes the key issue.

Our main idea is to divide an image  $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$  into blocks  $\mathbf{x}_b \in \mathbb{R}^{h \times w \times B^2 \times C}$ , where  $B$  is block size and  $h = \frac{H}{B}, w = \frac{W}{B}$ . When the image is fed into the network, the channel size and resolution of a block may change, but the layout and the relative positional relationships between blocks will remain unchanged. If there is a way to apply sequential batch generation of blocks where each batch simultaneously produces a subset of the blocks, only a small number of block hidden states have to be kept in memory simultaneously, making it possible to generate ultra-high-resolution images.

Here we define that block<sub>A</sub> is dependent on block<sub>B</sub> if the generation of block<sub>A</sub> involves the hidden state of block<sub>B</sub> in computation. It can be observed that the dependencies between blocks are bidirectional in most previous structures (UNet, DiT, etc.), in which case all blocks in the image must be generated simultaneously. Take UNet as an example: two adjacent elements in neighboring blocks use each other’s hidden state in the convolution operation, therefore all pairs of neighboring blocks must be generated simultaneously. Given the aim to save the memory of blocks’ hidden states, we hope to devise an algorithm that allows the blocks in the same image to be divided into several batches for generation, with each batch only needing to generate a portion of the blocks at the same time, and the batches are generated in sequence. Generally, an image generation algorithm can perform such a sequential batch generation among blocks if it meets the following conditions:

1. The generative dependency between blocks are unidirectional, and can form a directed acyclic graph (DAG).



2. Each block has only a few direct (1st-order) dependencies on other blocks, since the hidden states of the block and its direct dependencies should simultaneously be kept in the memory.

Furthermore, to ensure consistency across the whole image, it also required that the blocks have a large enough receptive field to manage long-range dependencies.

According to the conditions and analysis above, we choose an efficient implementation illustrated in Fig. 3, Unidirectional Block Attention(UniBA). For each layer, every block directly depends on three 1st-order neighboring blocks: the block on the top, on the left, and in the upper left corner. For example, if we adopt Diffusion Transformer (DiT) architecture which is the base architecture of Inf-DiT, the dependency between blocks is attention operation, where the query vectors of each block interact with the key, value vectors of 4 blocks: 3 blocks located to its upper left and itself, illustrated in Fig. 3.

Formally, the UniBA process in the transformer can be formulated as

$$\mathbf{Q} = \mathbf{z}_{(i,j)}^n \mathbf{W}^Q, \quad (1)$$

$$\mathbf{K} = [\mathbf{z}_{(i,j)}^n + \mathbf{P}_1; \mathbf{z}_{(i-1,j)}^n + \mathbf{P}_2; \mathbf{z}_{(i,j-1)}^n + \mathbf{P}_3; \mathbf{z}_{(i-1,j-1)}^n + \mathbf{P}_4] \mathbf{W}^K, \quad (2)$$

$$\mathbf{V} = [\mathbf{z}_{(i,j)}^n; \mathbf{z}_{(i-1,j)}^n; \mathbf{z}_{(i,j-1)}^n; \mathbf{z}_{(i-1,j-1)}^n] \mathbf{W}^V, \quad (3)$$

$$\mathbf{z}_{(i,j)}^{n+1} = \text{FFN}(\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V})), \quad (4)$$

where  $\mathbf{z}_{(i,j)}^n$  is the hidden states of the block at row  $i$ , column  $j$  in layer  $n$ , and  $\mathbf{P}_i$  is the block-level relative position encoding. We also implement an efficient approach to apply UniBA for full image in pytorch style, which is attached in the Appendix A.

Note that, though each block only attends to a few number of neighboring blocks in each layer, as features propagate layer by layer, blocks can indirectly interact with faraway blocks, thereby capturing both long- and short-range relationships. Our design shares a similar spirit with the natural language model Transformer-XL [4], which can be viewed as a special form of ours in one dimension case.

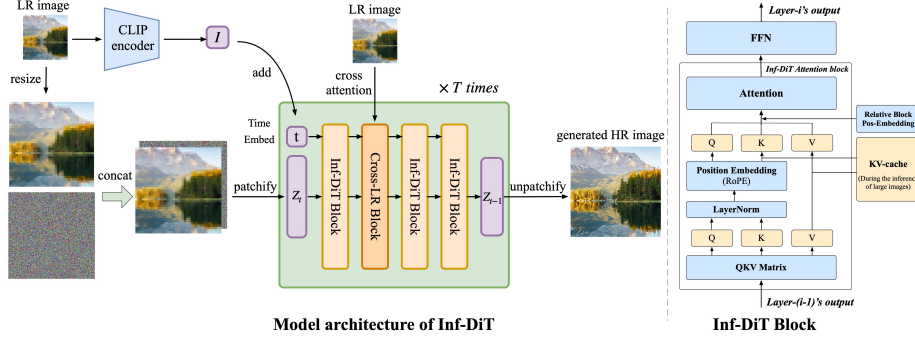
**Inference process with  $\mathcal{O}(N)$  Memory Consumption** Although our method can generate each block sequentially, it differs from auto-regressive generative models, in which the next block depends on the final output of the previous block. Any number of blocks can also be generated in parallel in our model, as long as the union of their dependent blocks has been generated. Based on this property, we implement a simple but effective inference process. As illustrated in Fig. 3, we generate  $n \times n$  blocks at once, from the top-left to bottom-right. After generating a set of blocks, we discard hidden states i.e. KV-cache that are no longer used and append newly generated KV-cache to the memory.

It can be easily proved that the number of block KV-cache retained in memory during the process is always  $\leq w + n$ . Assume the space needed by the

model when generating a single block is  $M_1$ , the space for one block’s KV-cache is  $M_2$ , and other essential space consumption (such as storing the raw input image) is  $C$ , then the maximum space usage of the inference process is  $n^2M_1 + (w+n)M_2 + C$ . When  $n$  is much smaller than  $w$ , the memory consumption is directly proportional to  $w$ . If  $w$  is bigger than  $h$ , we can easily change the trajectory of generation to column-major.

In practical terms, despite the total FLOPs of generation remaining constant for various values of  $n$ , due to the overhead such as operator initialization time and memory allocation time, the generation time decreases when  $n$  increases. Therefore, it’s optimal to choose the largest  $n$  allowed by the memory limitation.

## 2.2 Basic Model Architecture



**Fig. 4:** (Left) Overall architecture of Inf-DiT. (Right) The inner structure of Inf-DiT block. We do not depict the LayerNorm that originally existed in the DiT for simplicity.

Figure 4 provides an overview of our model, Inf-DiT’s architecture. The model uses a similar backbone as DiT [17], which applies Vision Transformer (ViT) [6] to diffusion models and proves its efficacy and scalability. In addition to its superior performance, compared to convolution-based architectures such as UNet [22], DiT solely utilizes attention as the mechanism of interaction among patches, which can conveniently implement unidirectional block attention. To adapt to unidirectional block attention and enhance the performance of upsampling, we make several modifications and optimizations detailed as follows.

**Model input** Inf-DiT first partitions input images into multiple non-overlapping *blocks*, which are further divided into *patches* with a side length equal to the *patch size*. Unlike DiT, considering the compression loss such as color shifting and detail loss, the patchifying of Inf-DiT is conducted in **RGB pixel space** instead of latent space. In the case of super-resolution by  $f$  times, Inf-DiT first upsamples the low-resolution RGB image condition by a factor of  $f$ , then concatenates it

with the diffusion’s noised input on the feature dimension before feeding it to the model.

**Position Encoding** Unlike UNet-based diffusion models [21] that can perceive positional relationships through convolution operations, all operations including self-attention, FFN in transformers are permutation invariant functions. Therefore, transformer-based models require auxiliary input of explicit positional information. We refer to the design of Rotary Positional Encoding (RoPE) [28], which performs well in the long context generation, and adapt it into two-dimensional form for image generation. Specifically, we divide channels of hidden states in half, one for encoding the x-coordinate and the other for the y-coordinate, and apply RoPE in both halves. We create a sufficiently large rope positional encoding table to ensure it meets the requirements during generation. To ensure all parts of the positional encoding table can be seen by the model during training, we employ the *Random Starting Point*: For each training image, we randomly assign a position  $(x, y)$  for the top-left corner of the image, instead of the default  $(0, 0)$ . In addition, considering the difference in interactions within the same block and between different blocks, we additionally introduce block-level relative position encoding  $\mathbf{P}_{1\sim 4}$ , which assigns a distinct learnable embedding based on the relative position before attention.

### 2.3 Global and Local Consistency

**Global Consistency with CLIP Image Embedding** The global semantic information within low-resolution (LR) images, such as artistic style and object material, plays a crucial role during upsampling. However, compared to text-to-image generation models, the upsampling model has an additional task: understanding and analyzing the semantic information of LR images, which significantly increases the model’s burden. This is particularly challenging when training without text data, as high-resolution images rarely have high-quality paired texts, making these aspects difficult for the model. We utilize the image encoder from pre-trained CLIP [19] to extract image embedding  $I_{LR}$  from low-resolution images, which we refer to as *Semantic Input*. Since CLIP is trained on massive image-text pairs from the Internet, its image encoder can effectively extract global information. We add the global semantic embedding to the time embedding of the diffusion transformer and input it into each layer, enabling the model to learn directly from high-level semantic information.

Another interesting advantage of global semantic embedding is that, using the aligned image-text latent space in CLIP, *we can use text to guide the direction of generation, even if our model has not been trained on any image-text pairs*. Given a positive prompt  $C_{pos}$  and a negative prompt  $C_{neg}$ , we can update our image embedding:

$$\tilde{I}_{LR} = \text{norm}(\mathbf{I}_{LR} + \alpha(\text{TextEnc}(C_{pos}) - \text{TextEnc}(C_{neg}))), \quad (5)$$

where  $\alpha$  can control the intensity of the guidance. During inference, we can simply use  $\tilde{I}_{LR}$  in place of  $\mathbf{I}_{LR}$  as the global semantic embedding to conduct

the control. For example, to get a more clear result set,  $C_{pos} = \text{“clear”}$  and  $C_{neg} = \text{“blur”}$  sometimes help.

**Local Consistency with Nearby LR Cross Attention** Although concatenating the LR image with noised input already provides a good inductive bias for models to learn the local correspondence between LR and HR images, there can still be continuity issues. The reason is that, there are several possibilities of upsampling for a given LR block, which require analysis in conjunction with several nearby LR blocks to select one solution, but the LR block to the right and below is unseen. To handle this problem, we introduce Nearby LR Cross Attention. In the first layer of the transformer, each block conducts cross-attention on the surrounding  $3 \times 3$  LR blocks to capture nearby LR information. This approach significantly reduces the probability of generating discontinuous images. It is worth noting that this operation will not change our inference process since the entire LR image is known before generation.

We further devise techniques including continuity class-free guidance, LR-based noise initialization, QK Normalization, etc., which are elaborated in detail in the Appendix [A](#).

### 3 Experiments


In this section, we first introduce the detailed training process of Inf-DiT, then comprehensively evaluate Inf-DiT by both machine and human evaluation. The results show that Inf-DiT surpasses all baselines, excelling in both ultra-high-resolution image generation and upsampling tasks. Finally, we conduct ablation studies to validate the effectiveness of our design.

#### 3.1 Training Details

**Datasets** Our dataset comprises a subset of LAION-5B [\[25\]](#) with a resolution higher than  $1024 \times 1024$  and aesthetic score higher than 5, and 100 thousand high-resolution wallpapers from the Internet. Following the previous works [\[20, 23, 30\]](#), we use fixed-size image crops of  $512 \times 512$  resolution during training. Since upsampling can be conducted with local information only, it can be directly employed at higher resolution during inference, which is not easy for most generation models. The data processing details are illustrated in Appendix [A.2](#).

**Training settings** During training, We set *block size* = 128 and *patch size* = 4, which means every training image is divided into  $4 \times 4$  blocks and every block has  $32 \times 32$  patches. We employ the framework of EDM [\[15\]](#) for training, and set the upsampling factor to  $4\times$ . Because the upsampling task is more concerned with the high-frequency details of images, we adjusted the mean and std of training noise distribution to  $-1.0$  and  $1.4$ . To address overflow problem during training, we employed BF16 format due to its broader numerical range. Our

CLIP model is a ViT-L/16 pre-trained on Datacomp dataset [13]. Since the CLIP can only process images with  $224 \times 224$  resolution, we first resize LR images to  $224 \times 224$  and then input them to CLIP. Other hyperparameters are listed in the Appendix A.




Direct Patch Scalecrafter BSRGAN Demofusion Inf-DiT  
**Fig. 5:** Qualitative comparison of different methods in detail at  $2048 \times 2048$  resolution

### 3.2 Machine Evaluation

In this part, we conduct quantitative comparisons of Inf-DiT against state-of-the-art methodologies on ultra-high-resolution image generation tasks. The baselines encompass two main categories of high-resolution generation: 1. Direct

**Table 1:** Quantitative comparison results with state-of-the-art methods of super resolution on DIV2K dataset. The best results are marked in **bold**.

Setting	Method	FID ↓	FID <sub>crop</sub> ↓	PSNR ↑	SSIM ↑
Variable resolution	BSRGAN	35.0	143.2	25.9	0.73
	Real-ESRGAN	35.5	127.6	25.2	0.72
	Inf-DiT(Ours)	<b>20.2</b>	<b>76.5</b>	<b>26.3</b>	<b>0.74</b>
Fix resolution (512x512)	StableSR	74.8	112.6	22.0	0.60
	Inf-DiT(Ours)	<b>38.6</b>	<b>83.3</b>	<b>24.6</b>	<b>0.67</b>
Fix resolution (256x256)	LDM	152.6	-	23.9	0.66
	Inf-DiT(Ours)	<b>86.8</b>	-	<b>24.6</b>	<b>0.67</b>



**Fig. 6:** Qualitative comparison of different methods in detail at  $4096 \times 4096$  resolution

high-resolution image generation, including Direct Inference of SDXL [18], MultiDiffusion [2], ScaleCrafter [9], and 2. High-resolution image generation based on super-resolution techniques, including BSRGAN [33], DemoFusion [7]. We employ the FID (Fréchet Inception Distance) [10] to evaluate the quality of ultra-high-resolution generation, which is widely used to evaluate the perceptual quality of images in image generation tasks [17, 21, 23]. To further validate the super-resolution ability of our model, we additionally benchmark it against renowned super-resolution models over classic super-resolution tasks.

**Ultra-High-Resolution Generation** We use the test set of HPDv2 [31] for evaluation, which allows for a comprehensive evaluation of the model’s generative capabilities (3200 prompts) across various domains and styles (“Animation”, “Concept-art”, “Painting”, and “Photo”). We test on two resolutions:  $2048^2$  and  $4096^2$ . For super-resolution-based models, we first use SDXL to generate a  $1024^2$  resolution image and upsample it without text. Although Inf-DiT was trained in a setting with  $4\times$  upsample, we find that it can generalize well to upsampling ratios lower than training. Therefore, for  $2048^2$  generation, we directly resize the LR image from  $1024^2$  to  $2048^2$  and concatenate it with the noise input. As for the metrics, we use FID to assess how closely generated images mimic real image distributions, and the real images are randomly sampled from LAION-5B. However, the original implementation of FID downsamples the input image to  $299 \times 299$  before the feature extraction, which ignores critical high-resolution details. Therefore, inspired by [3, 7], we additionally adopted the idea of  $FID_{\text{crop}}$



**Table 2:** Quantitative comparison results with state-of-the-art methods of ultra-high-resolution generation on HPDV2 dataset. The best results are marked in **bold**, and the second best results are marked by underline.

Method	2048 × 2048		4096 × 4096		Mean ↓
	FID ↓	FID <sub>crop</sub> ↓	FID ↓	FID <sub>crop</sub> ↓	
Direct Inference	92.2	92.3	OOM	OOM	92.3
MultiDiffusion	99.7	109.3	OOM	OOM	104.5
SDXL+BSRGAN	<u>66.3</u>	81.7	<b>66.0</b>	85.0	74.8
ScaleCrafter	79.9	95.5	115.1	153.3	111.0
DemoFusion	67.5	<u>72.0</u>	70.8	<u>82.6</u>	<u>73.2</u>
SDXL+Inf-DiT(Ours)	<b>66.0</b>	<b>71.9</b>	<u>67.0</u>	<b>76.2</b>	<b>70.3</b>

and randomly cropped  $299 \times 299$  patches from high-resolution images for FID evaluation.


The results in Tab. 2 demonstrate that our model achieves state-of-the-art (SOTA) in three out of four metrics. This showcases our model’s exceptional ability to generate high-resolution details and harmonious global information. Note that our model is capable of being applied to all generative models. Examples from other models are listed in the Appendix D. We have also evaluated the model’s time-space consumption in Appendix A.

**Super-Resolution** In addition to its ability to generate high-resolution images, Inf-DiT can also be used as a classic super-resolution model. We conduct the evaluation on DIV2k valid dataset [1] which contains multiple real-world high-resolution images in different scenarios. Following [21, 24], we fix image degradation to bicubic interpolation with  $4\times$  downsampling. Before comparing with fix-resolution models LDM [21] and StableSR [29], we center-crop specific-sized patches from the high-resolution image as ground truth. Throughout this process, we utilize both perceptual (FID, FID<sub>crop</sub>) and fidelity (PSNR, SSIM) metrics to ensure a detailed and comprehensive assessment.

As demonstrated in Tab. 1, our model achieves state-of-the-art across all metrics. This shows that, as a super-resolution model, ours not only excels in performing super-resolution at arbitrary scales but also preserves global and detailed information while restoring results that closely resemble the LR images.

### 3.3 Human Evaluation

To further evaluate Inf-DiT and more accurately reflect its generative quality from a human perspective, we conduct a human evaluation. The whole evaluation process is detailed in Appendix C. We ultimately collected 3,600 comparisons. As depicted in Fig. 7, our model outperforms the other 3 methods in all the 3 criteria. It is particularly noteworthy that each of the other three models ranks relatively lower on at least one of the three evaluation criteria, while Inf-DiT achieves the highest score on all three criteria: detail authenticity, global



**Fig. 7:** Human evaluation results. After participants rank the images from different models, We assign scores from 4 to 1 to models in order and finally calculate the average of all the results. Inf-DiT gets the highest scores in all three categories.

coherence, and consistency with low-resolution input. This indicates that our model is the only one capable of excelling in both high-resolution generation and super-resolution tasks simultaneously.

### 3.4 Iterative Upsampling


Since our model can upsample images of arbitrary resolution, it’s a natural idea to test if the model can iteratively upsample images generated by itself. In this study, we experiment on generating a  $2048^2$  resolution image from a  $32^2$  resolution image,  $64\times$  upscale, through three times of iterative upsampling. Figure 8 illustrates two cases of this process. In the first case, the model successfully produces a high-resolution image after three stages of upsampling. It generates details of different frequencies in different resolution upsampling: the contour of the face, the shape of the eyeball, and individual eyelashes. However, it is hard for the model to correct inaccuracies generated in earlier stages, leading to the accumulation of errors. An example of this issue is demonstrated in the second sample. We leave this problem for our future work.

**Table 3:** Ablation study on  $2048 \times 2048$  resolution.

Method	FID	FID <sub>crop</sub>	Block Size	32	64	128	256	512
Ours	<b>66.0</b>	<b>71.9</b>	Memory Usage (GB)	23.3	23.3	23.3	23.5	23.7
w/o Nearby LR Attention	66.7	73.6	FID	276	271	268	265	246
w/o Semantic Input	66.2	72.8	Inference Time (s)	57	61	69	112	269
Inner Block Attention	66.3	83.7						

### 3.5 Ablation Study

We conduct all ablation studies with  $2048^2$  resolution on HPDV2 dataset, with results in Tab. 3. To verify our method, we separately 1. remove the Global



**Fig. 8:** Generated samples of iterative upsampling. **Top:** Inf-DiT can upsample an image generated by itself several times, and generate details of different frequencies at corresponding resolution. **Bottom:** After failing to generate a pupil at the  $128^2$  resolution, it is difficult for subsequent upsampling stages to correct this error.

Input, 2. remove Nearby LR attention and 3. remove the attention between blocks (Inner Block Attention setting). As a result, both the FID and  $\text{FID}_{\text{crop}}$  of the model get worse to varying degrees. We also conduct an ablation study to decide the optimal block size. The results show that as block size increases, the performance improves, but the time also increases. To trade-off, we chose 128 as the default setting. Further case studies and analysis are listed in the Appendix B.

## 4 Related Work

### 4.1 Diffusion Image Generation

Diffusion models have emerged as a spotlight in the realm of image generation, boasting an array of groundbreaking advancements in recent years. Initially introduced in 2015 [26], and further developed through works such as DDPM [11] and DDIM [27], these models leverage a stochastic diffusion process, conceptualized as a Markov chain, to convert a simple prior distribution, like Gaussian noise, into a complex data distribution. This methodology yield impressive outcomes in terms of the quality and diversity of generated images.

Recent enhancements have markedly elevated the generation capabilities of diffusion models. CDM [12] creates a cascade generation pipeline with multistage super-resolution models, which can be applied to large pretrained models [20, 23].

The introduction of Latent Diffusion Models (LDMs) [18, 21] represents a pivotal extension, incorporating a latent space to boost both efficiency and scalability. Alongside these advancements, there has been significant progress in

the optimization of network architectures. The advent of Diffusion Transformers (DiT) [17] Replaces U-Net with ViT [6] for noise prediction.

## 4.2 Image Super-Resolution

Given a low-resolution (LR) image  $I_{LR}$  degraded from high-resolution (HR) observation  $I_{HR}$ , image super-resolution (SR) aims to reconstruct a HR one  $\hat{I}_{HR}$ :  $I_{LR} = D(I_{HR}; \delta)$ ,  $\hat{I}_{HR} = F(I_{LR}; \theta)$ . Here  $D$  and  $F$  refer to degradation process and the super-resolution model.  $\delta$  and  $\theta$  represent the parameters. In recent years, Blind SR has been a major focus, in which the degradation process is unknown and learnable. The representative methods are BSRGAN [33] and Real-ESRGAN [30]. Recently, diffusion-based SR methods have yielded exciting results. These works focus on finetuning pre-trained text-to-image diffusion models to take advantage of their excellent generative ability. Specifically, DiffBir [16] employs ControlNet [34] on pre-trained stable-diffusion models, whereas PASD [32] enhances it by executing pixel-aware cross-attention. Both approaches have garnered considerable success in fixed resolution super-resolution but cannot be directly used for higher resolution.

## 4.3 Ultra-High-Resolution Image Upsampler

Currently, image generation methods fall short in generating ultra-high-resolution images owing to the constraints of memory and issues with training efficiency. Under these circumstances, MultiDiffusion [2] and Mixture of Diffusers [14] bind together multiple diffusion generation processes by dividing images into overlapping blocks, processing each separately, and then stitching them together, aiming to maintain continuity between blocks. However, because they only used the local weighted averaging for aggregation, it leads to a low interaction efficiency and makes it difficult to ensure the global consistency of the images.

Given this concern, DemoFusion [7] and ScaleCrafter [9] adapt dilated policies including dilated sampling and dilated convolution kernels, aiming at acquiring more global information. These methods indeed achieve improvement on a global semantic level without additional training. However, the huge difference between training and generation leads these methods to easily produce illogical images.

## 5 Conclusion

In this work, we observe that the major obstacle to generating ultra-high-resolution images is the substantial memory occupied by model hidden states. Based on this, We propose Unidirectional Block Attention mechanism (UniBA) which can lower the space complexity by performing batch generation among blocks. With UniBA, we train Inf-DiT, a  $4\times$  memory-efficient image upsampler which achieves state-of-the-art performance in both generation and super-resolution tasks.

## Acknowledgements

The authors would like to thank Zhipu.AI for their GPU and human evaluation support, and the reviewers of ECCV for their valuable suggestions.

## References

1. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 126–135 (2017)
2. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation (2023)
3. Chai, L., Gharbi, M., Shechtman, E., Isola, P., Zhang, R.: Any-resolution training for high-resolution image synthesis. In: European Conference on Computer Vision. pp. 170–188. Springer (2022)
4. Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q.V., Salakhutdinov, R.: Transformer-xl: Attentive language models beyond a fixed-length context. arXiv preprint arXiv:1901.02860 (2019)
5. Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A.P., Caron, M., Geirhos, R., Alabdulmohsin, I., et al.: Scaling vision transformers to 22 billion parameters. In: International Conference on Machine Learning. pp. 7480–7512. PMLR (2023)
6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
7. Du, R., Chang, D., Hospedales, T., Song, Y.Z., Ma, Z.: Demofusion: Democratising high-resolution image generation with no \$\$\$\$. arXiv preprint arXiv:2311.16973 (2023)
8. Guttenberg, N.: Diffusion with offset noise (2023), <https://www.crosslabs.org/blog/diffusion-with-offset-noise>
9. 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: The Twelfth International Conference on Learning Representations (2023)
10. 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. Advances in neural information processing systems **30** (2017)
11. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
12. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. The Journal of Machine Learning Research **23**(1), 2249–2281 (2022)
13. Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., Hajishirzi, H., Farhadi, A., Schmidt, L.: Openclip (Jul 2021). <https://doi.org/10.5281/zenodo.5143773> <https://doi.org/10.5281/zenodo.5143773> if you use this software, please cite it as below.
14. Jiménez, Á.B.: Mixture of diffusers for scene composition and high resolution image generation. arXiv preprint arXiv:2302.02412 (2023)

15. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* **35**, 26565–26577 (2022)
16. Lin, X., He, J., Chen, Z., Lyu, Z., Fei, B., Dai, B., Ouyang, W., Qiao, Y., Dong, C.: Diffbir: Towards blind image restoration with generative diffusion prior. *arXiv preprint arXiv:2308.15070* (2023)
17. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4195–4205 (2023)
18. 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. *arXiv preprint arXiv:2307.01952* (2023)
19. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
20. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
21. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* **18**. pp. 234–241. Springer (2015)
23. 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. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)
24. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4713–4726 (2022)
25. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C.W., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S.R., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: LAION-5b: An open large-scale dataset for training next generation image-text models. In: *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track* (2022), <https://openreview.net/forum?id=M3Y74vmsMcY>
26. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. pp. 2256–2265. PMLR (2015)
27. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
28. Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., Liu, Y.: Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* **568**, 127063 (2024)
29. Wang, J., Yue, Z., Zhou, S., Chan, K.C., Loy, C.C.: Exploiting diffusion prior for real-world image super-resolution. *arXiv preprint arXiv:2305.07015* (2023)
30. Wang, X., Xie, L., Dong, C., Shan, Y.: Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1905–1914 (2021)



31. Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., Li, H.: Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. arXiv preprint arXiv:2306.09341 (2023)
32. Yang, T., Ren, P., Xie, X., Zhang, L.: Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. arXiv preprint arXiv:2308.14469 (2023)
33. Zhang, K., Liang, J., Van Gool, L., Timofte, R.: Designing a practical degradation model for deep blind image super-resolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4791–4800 (2021)
34. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)