

CAMixerSR: Only Details Need More “Attention”

Yan Wang^{1,2‡} Yi Liu^{1*} Shijie Zhao^{1*†} Junlin Li¹ Li Zhang¹
¹Bytedance Inc. ²Nankai-Baidu Joint Lab, Nankai University

wyrmymy@foxmail.com {zhaoshijie.0526, liuyi.chester, lijunlin.li, lizhang.idm}@bytedance.com

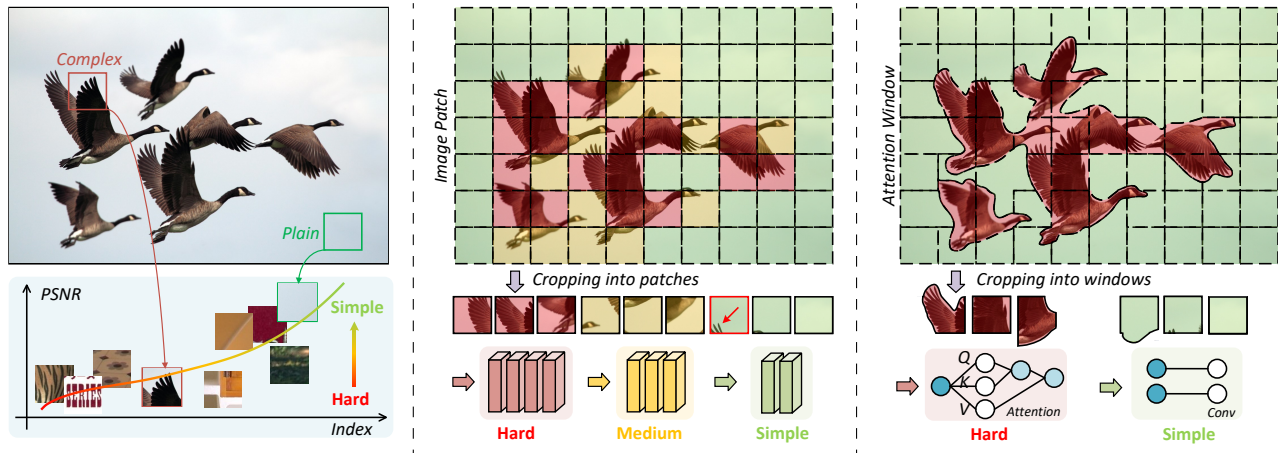


Figure 1. Comparison of ClassSR [19] framework and CAMixer. Left) the plain/complex patches are at varied levels of difficulty to restore. Middle) ClassSR crops input images to sub-images for discriminative processing through models of varied complexities. Right) we introduce a content-aware mixer (CAMixer) to calculate self-attention for complexity regions while convolution for simple context.

Abstract

To satisfy the rapidly increasing demands on the large image (2K-8K) super-resolution (SR), prevailing methods follow two independent tracks: 1) accelerate existing networks by content-aware routing, and 2) design better super-resolution networks via token mixer refining. Despite directness, they encounter unavoidable defects (e.g., inflexible route or non-discriminative processing) limiting further improvements of quality-complexity trade-off. To erase the drawbacks, we integrate these schemes by proposing a content-aware mixer (CAMixer), which assigns convolution for simple contexts and additional deformable window-attention for sparse textures. Specifically, the CAMixer uses a learnable predictor to generate multiple bootstraps, including offsets for windows warping, a mask for classifying windows, and convolutional attentions for endowing convolution with the dynamic property, which modulates attention to include more useful textures self-adaptively and improves the representation capability of convolution. We further in-

troduce a global classification loss to improve the accuracy of predictors. By simply stacking CAMixers, we obtain CAMixerSR which achieves superior performance on large-image SR, lightweight SR, and omnidirectional-image SR.

1. Introduction

Recent research on neural networks has significantly improved the image super-resolution (SR) quality [22, 34, 43]. However, existing methods generate visual-pleasing high-resolution (HR) images but suffer intensive computations in real-world usages, especially for 2K-8K targets. To alleviate the overhead, many **accelerating frameworks** [4, 19] and **lightweight networks** [14, 32] were introduced for practical super-resolution application. However, these approaches are completely independent without cooperation.

The first strategy, the accelerating frameworks [11, 19, 39], is based on the observation that *different image regions require different network complexities*, which tackles the problem from the perspective of content-aware routing of varied models. As depicted in the middle image of Fig. 1, they decomposed large-input images into fixed patches and

*Equal contribution. †Corresponding author.

‡Work done during an internship at Bytedance.

Code is available at: www.github.com/icandle/CAMixerSR




Figure 2. Performance (PSNR-FLOPs) comparison on Test8K. The green dashline indicates the trade-off curve of CAMixerSR.

assigned patches to networks via an extra classification network. ARM [4] further developed the strategy by introducing a LUT-based classifier and parameter-sharing design to improve efficiency. Despite these strategies being generic for all neural networks, two unavoidable defects remain. One is poor classification and inflexible partition. Fig. 1 displays the windows with little details that are improperly sent to a simple model. The other is the limited receptive fields. As shown in Tab. 2, cropping images into patches limits the receptive field, thus influencing the performance.

The second strategy, lightweight model design [7, 8, 17, 44], focuses on refining the neural operators (self-attention or convolution) and backbone structures to enable stronger feature representation capability within limited layers, *i.e.*, using more intra-information to reconstruct images. For instance, NGswin [5] exploited N-Gram for self-attention to reduce the calculations and enlarge the receptive field. IMDN [14] introduced information multi-distillation for efficient block design. Although these lightweight methods reached impressive efficiency on 720p/1080p images, their usages are rarely examined with larger input (2K-8K). Moreover, these approaches ignore intrinsic characteristics that different content can be discriminatively processed.

This paper, firstly integrating the above strategies, is based on the derived observation that *distinct feature regions demand varying levels of token-mixer complexities*. As shown in Tab. 1, simple convolution (Conv) can perform similarly with much more complex convolution + self-attention (SA + Conv) for simple patches. Hence, we propose a content-aware mixer (CAMixer) to route token mixers with different complexities according to the content. As depicted in Fig. 1, our CAMixer uses complex self-attention (SA) for intricate windows and simple convolution for plain windows. Furthermore, to address the limitations of ClassSR, we introduce a more sophisticated predictor. This predictor utilizes multiple conditions to generate ad-

Table 1. PSNR (dB) values obtained by three token mixers.

Method	FLOPs	Simple	Medium	Hard
Conv	517M	43.73	30.96	23.60
SA + Conv	979M	43.80	31.19	23.80
CAMixer	747M	43.81	31.17	23.78

Table 2. Cropping images into smaller tiles results in larger drops.

Method	{Tile, Overlap}		
	{128,8}	{64,4}	{32,2}
RCAN [43]	29.38	29.37	29.32
IMDN [14]	29.03	29.02	28.96
SwinIR-light [21]	29.24	29.22	29.17

ditional valuable information, thereby enhancing CAMixer with improved partition accuracy and better representation. Based on CAMixer, we construct CAMixerSR for super-resolution tasks. To fully examine the performance of CAMixer, we conduct experiments on both lightweight SR, large-input (2K-8K) SR, and omnidirectional-image SR. Fig. 2 illustrates CAMixerSR advances both lightweight SR and accelerating framework by a large margin. Our contributions are summarized as:

- We propose a Content-Aware Mixer (CAMixer) integrating convolution and self-attention, which can adaptively control the inference computation by assigning simple areas to convolution and complex areas to self-attention.
- We propose a powerful predictor to generate the offset, mask, and simple spatial/channel attentions, which modulates CAMixer to capture longer-range correlation with fewer calculations.
- Based on CAMixer, we build CAMixerSR which exhibits state-of-the-art quality-computation trade-offs on three challenging super-resolution tasks: lightweight SR, large-input SR, and omnidirectional-image SR.

2. Related Work

Accelerating framework for SR. As the complexity continuously enlarged for better restoration quality, the practical application for SR models becomes harder, especially for 2K-8K SR. Recent research [4, 19, 39] tackled this problem from a different perspective. Instead of designing a lightweight model, they use content-aware routing to send cropped patches to models with different complexities dynamically. ClassSR [19] leveraged a 3-class classifier to determine the sub-image calculated by complex/medium/simple networks, which saves 50% calculations for RCAN [43] on 8K datasets. PathRestore [39] learned to select feature paths to adapt FLOPs according to context.

Lightweight SR. Since the first work, SRCNN [7] using convolutional neural network (ConvNet) in the super-resolution task, numerous strategies [7, 14, 23, 34] to reduce the complexity have been proposed for a more lightweight

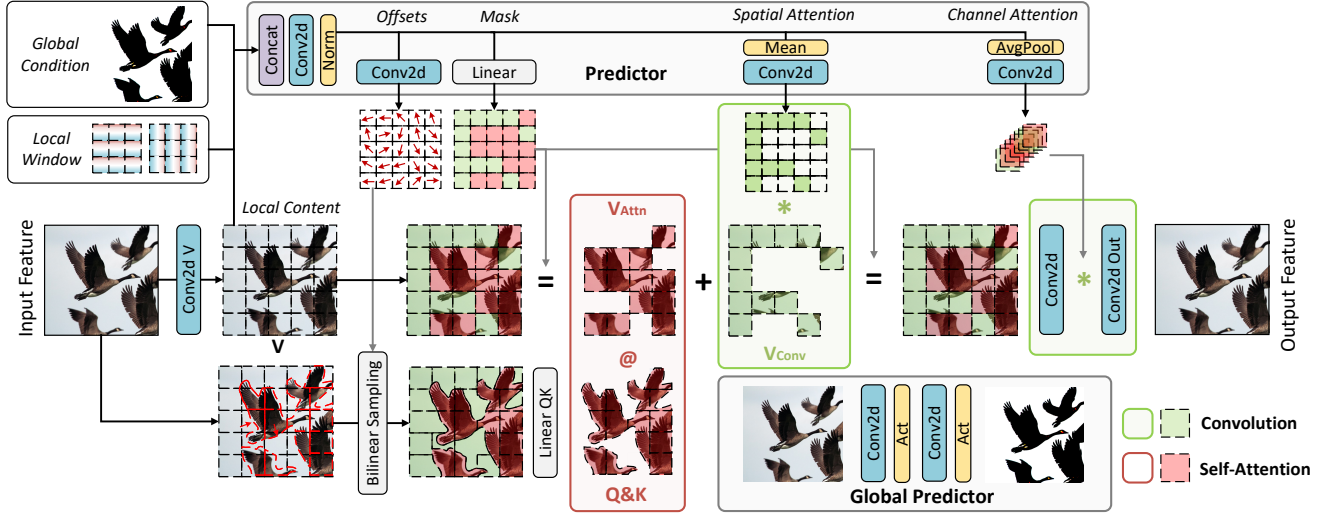


Figure 3. Overview of the proposed CAMixer. CAMixer consists of three parts: Predictor, Self-Attention branch, and Convolution branch.

inference. The early DRCN [16] and DRRN [31] tried to decrease parameters by using recurrent blocks but suffered intensive computations. To alleviate the drawback, IDN [13] and IMDN [14] adopted efficient information fusion structure to reduce both parameters and calculations. The latter ConvNet, *e.g.*, EFDN [32] and RLFN [18] further simplified the information distillation procedure and introduced reparameterization to obtain real-time inference on mobile devices. With the recently rapid development of transformers, varied powerful token mixers, *e.g.*, self-attention [41, 42] and large kernel convolution [33, 45], are introduced to lightweight SR. For example, SwinIR [21] utilized window-based SA, and MAN [33] employed large-kernel convolution, both of which achieved SOTA performance. Although these neural operators are capable of capturing long-range correlations, they cost massive calculations. To this end, this paper integrates the content-aware routing to token mixer design that adopts complex operators for informative areas but simple operators for plain areas.

3. Method

3.1. Content-Aware Mixing

The overview of our proposed CAMixer is shown in Fig. 3. CAMixer consists of three main components: the predictor module, the attention branch, and the convolution branch. Given the input features $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$, it is first projected to attain *value* $\mathbf{V} \in \mathbb{R}^{C \times H \times W}$ via a point-wise convolution:

$$\mathbf{V} = f_{PWConv}(\mathbf{X}), \quad (1)$$

Predictor. Based on the local condition $\mathbf{C}_l = \mathbf{V}$, global condition $\mathbf{C}_g \in \mathbb{R}^{2 \times H \times W}$, and linear positional encoding $\mathbf{C}_w \in \mathbb{R}^{2 \times H \times W}$, the predictor first calculates the shared

intermediate feature map \mathbf{F} and then generate **offsets maps**, **mixer mask**, and simple **spatial/channel attention**:

$$\begin{aligned} \mathbf{F} &= f_{head}(\mathbf{C}_l, \mathbf{C}_g, \mathbf{C}_w), \quad \hat{\mathbf{F}} = f_{reduce}(\mathbf{F}) \in \mathbb{R}^{\frac{HW}{M^2} \times M^2}, \\ \Delta p &= r \cdot f_{offsets}(\mathbf{F}) \in \mathbb{R}^{2 \times H \times W}, \\ m &= \hat{\mathbf{F}} \mathbf{W}_{mask} \in \mathbb{R}^{\frac{HW}{M^2} \times 1}, \\ \mathbf{A}_s &= f_{sa}(\mathbf{F}) \in \mathbb{R}^{1 \times H \times W}, \\ \mathbf{A}_c &= f_{ca}(\mathbf{F}) \in \mathbb{R}^{C \times 1 \times 1}, \end{aligned} \quad (2)$$

where Δp is the content-related offsets matrix to warp the window with more complex structures. r is a scalar to control the offsets range. $\hat{\mathbf{F}}$ is the reduced and rearranged immediate feature according to the attention window size M . m is the mask to decide whether the cropped window is calculated by attention or convolution. \mathbf{A}_s and \mathbf{A}_c are spatial and channel attention to enhance the convolutional branch. **Attention Branch.** To calculate sparse attention for complex areas, we use the offsets Δp to modulate the original input \mathbf{X} by bilinear interpolation $\phi(\cdot)$ to include more useful content in selected windows:

$$\tilde{\mathbf{X}} = \phi(\mathbf{X}, \Delta p). \quad (3)$$

We subsequently rearrange the $\tilde{\mathbf{X}}$, $\mathbf{V} \in \mathbb{R}^{\frac{HW}{M^2} \times M^2 \times C}$ according to the window shape $M \times M$. During the training stage, we apply the gumble softmax [15, 29] to calculate binary mask $\mathbf{M} = \text{gumble_softmax}(m)$ for hard and simple token sampling. During inference, by descending sorting mask m with $\text{argsort}(m)$, we obtain the index I_{hard} of top K related windows for sparse attention, and the I_{simple} of other $\frac{HW}{M^2} - K$ ones for convolution, where $K = \sum \mathbf{M}$. We denote the ratio of the attention patch as $\gamma = K / \frac{HW}{M^2}$.

Table 3. Complexity comparison between convolution, window-based multi-head self-attention, and our CAMixer. CAMixer semantically integrates convolution and self-attention with changeable calculations. h, w : height and width of the input image. C : input and output channel. k : convolution kernel size. M : window size. γ : the ratio of tokens calculated by SA. ρ : the reduction ratio.

Method	FLOPs ($\times hw$)	Params
Conv	$k^2 C^2$	$k^2 C^2$
W-MSA	$4C^2 + 2M^2 C$	$4C^2$
CAMixer	$k^2 C +$	$k^2 C$
	$2(1 + \gamma)C^2 + 2\gamma M^2 C +$ $\underbrace{\rho C(C + 4)}_{\text{shared}} + \underbrace{M + 2\rho C}_{\text{mask \& offsets}} + \underbrace{\rho K^2 C + \rho C^2 / hw}_{\text{spatial \& channel}} < 3\rho C^2$	$4C^2$

Upon the indices, we split the $\tilde{\mathbf{X}}$ and \mathbf{V} by:

$$\begin{aligned}\tilde{\mathbf{X}}_{hard} &= \tilde{\mathbf{X}}[I_{hard}] \in \mathbb{R}^{K \times M^2 \times C}, \\ \mathbf{V}_{hard} &= \mathbf{V}[I_{hard}] \in \mathbb{R}^{K \times M^2 \times C}, \\ \mathbf{V}_{simple} &= \mathbf{V}[I_{simple}] \in \mathbb{R}^{(\frac{HW}{M^2} - K) \times M^2 \times C}.\end{aligned}\quad (4)$$

After obtaining $\tilde{\mathbf{X}}_{hard}$, the *query* $\tilde{\mathbf{Q}}_{hard}$ and *key* $\tilde{\mathbf{K}}_{hard}$ are generated by linear layers:

$$\tilde{\mathbf{Q}} = \tilde{\mathbf{X}}_{hard} \mathbf{W}_q, \quad \tilde{\mathbf{K}} = \tilde{\mathbf{X}}_{hard} \mathbf{W}_k \in \mathbb{R}^{K \times M^2 \times C}. \quad (5)$$

Based on the above deduction, the self-attention for complex windows can be expressed as:

$$\mathbf{V}_{hard} = \text{softmax}\left(\frac{\tilde{\mathbf{Q}}\tilde{\mathbf{K}}^T}{\sqrt{d}}\right)\mathbf{V}_{hard} \in \mathbb{R}^{K \times M^2 \times C}. \quad (6)$$

For windows \mathbf{V}_{simple} for light operation, we use the rearranged \mathbf{A}_s to implement the simple attention by element-wise multiplication:

$$\mathbf{V}_{simple} = \mathbf{A}_s \cdot \mathbf{V}_{simple} \in \mathbb{R}^{(\frac{HW}{M^2} - K) \times M^2 \times C}. \quad (7)$$

Overall, we integrate the \mathbf{V}_{hard} and \mathbf{V}_{simple} to obtain the output of attention branch $\mathbf{V}_{attn} \in \mathbb{R}^{C \times H \times W}$ with the help of the the indices.

Convolution Branch. We leverage a depth-wise convolution and pre-generated channel attention to capture the local correlation, which can be formulated as:

$$\mathbf{V}_{conv} = f_{DWConv}(\mathbf{V}_{attn}) \cdot \mathbf{A}_c + \mathbf{V}_{attn} \quad (8)$$

Finally, the output of CAMixer is projected by a point-wise convolution as:

$$\mathbf{V}_{out} = f_{PWConv}(\mathbf{V}_{conv}), \quad (9)$$

CAMixer. Overall, by controlling the self-attention ratio γ , we adjust the content-aware mixing. When $\gamma = 1$, CAMixer is a combination of self-attention and convolution, which is similar to ACMix [28]. For $\gamma = 0$, CAMixer

is a pure convolutional token mixer with low complexity. For $\gamma \in (0, 1)$, CAMixer learns the content-aware mixing that uses complex mode for hard areas but simple mode for plain areas.

Complexity Analysis. We theoretically compare the complexity of convolution, window-based self-attention, and CAMixer in Tab. 3. Specifically, the FLOPs of CAMixer consist of three parts: the convolution, the attention, and the predictor. Given the input of $C \times h \times w$, the convolution branch utilizes a depth-wise convolution with a computational cost $K^2 C h w$. For attention branch, four projection operations cost $2(1 + \gamma)C^2 h w$ while the attention calculation costing $2\gamma M^2 C h w$, where $\gamma = K / (\frac{hw}{M^2})$ is the ratio of hard windows. For predictor module, it adds a serial of computations: $\rho C(C + 4)h w$ for shared head, $M h w$ for mask, $2\rho C h w$ for offsets, $\rho k^2 C h w$ for spatial attention, and ρC^2 for channel attention, where $\rho = \frac{1}{8}$ is the reduction ratio to reduce calculation.

3.2. Network Architecture

Lastly, we construct the CAMixerSR by modifying SwinIR-light [21]. Generally, CAMixerSR consists of four components, three from SwinIR: shallow feature extractor, deep feature extractor, reconstruction module, and additional global predictor module shown in Fig. 3. Furthermore, we replaced the window-based self-attention with CAMixer and reduced the block number.

3.3. Training Loss

We describe the training objectives of our CAMixerSR, including the optimization of the super-resolution framework and the predictor. Following the previous work [22, 34], we adopt the primary ℓ_1 loss to train the backbone. Assuming the input batch with N image pairs, *i.e.*, $\{I_i^{LR}, I_i^{HR}\}_{i=1}^N$, this process can be formulated by:

$$\ell_1 = \frac{1}{N} \sum_{i=1}^N \left\| I_i^{HR} - f_{CAMixerSR}(I_i^{LR}) \right\|_1, \quad (10)$$

where $f_{CAMixerSR}(\cdot)$ is the proposed CAMixerSR.

To supervise the predictors for S CAMixers, following previous works that control SA sparsity [29, 37], we adopt a simple but effective MSE Loss to control the ratio γ_i :

$$\ell_{ratio} = \frac{1}{N} \sum_{i=1}^N \left\| \gamma_{ref} \cdot \left(1 - \frac{2}{S} \sum_{i=1}^S \gamma_i \right) \right\|_2, \quad (11)$$

where γ_{ref} represents the referred overall ratio and γ_i denotes the hard token ratio for i -th CAMixer. To enable one training for dynamic ratio, we pre-train the CAMixerSR with $\gamma_{ref} = 0.5$ and fine-tune with $\gamma_{ref} \in [0, 1]$.

Overall, we train the CAMixerSR by simply combining the above objectives:

$$\ell = \ell_1 + \ell_{ratio}. \quad (12)$$

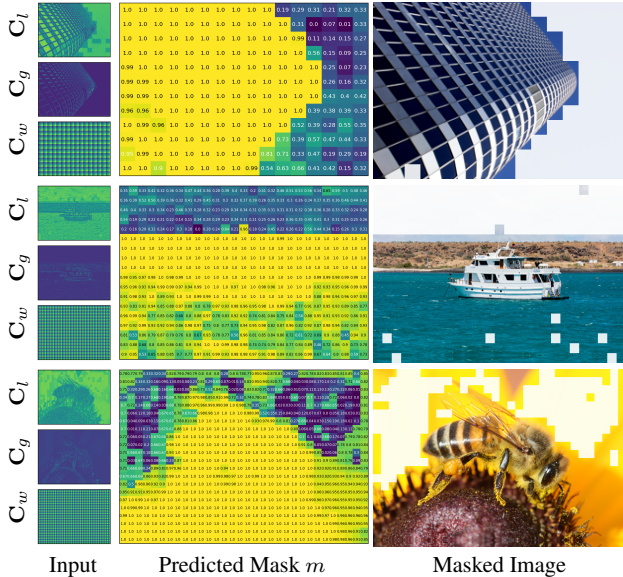


Figure 4. Visualizations of predicted mixer mask m of CAMixerSR. The lighter the color, the larger the magnitude. The scores of attention windows are in black, and the ones of convolution are in white. The unmasked tokens with more complex content (higher score) are processed by self-attention.

4. Experiment

4.1. Implementation Details

Model. Following SwinIR [21] and ELAN [42], we construct the overall backbone with 20 CAMixer and FFN blocks. The channel number is 60. The window size of self-attention is 16 and the convolution branch is implemented by two 3×3 depth-wise convolution. Specifically, we manually set the $\gamma = 1.0$ as the *Original* model (baseline) and $\gamma = 0.5$ ¹ as the *CAMixer* model (target). We provide more results of other settings in *supplementary material*.

Training. We train the proposed framework on three challenging super-resolution (SR) tasks: *lightweight SR*, *large-image SR*, and *omnidirectional-image (ODI) SR*. For the first two tasks, we use DIV2K [1] as the training set. For ODISR, we leverage the cleaned ODI-SR dataset [6]. The loss in Eq. (12) is calculated with batch-size 32 and patch-size 64. The AdamW [24] is adopted with the initial learning rate 5×10^{-4} and 500k iterations training procedure. We halved the learning rate at 250k, 400k, 450k, and 475k.

Testing. We first test CAMixerSR-*Original* with $\gamma = 1.0$ as the baseline, which represents all tokens being processed by self-attention and convolution. Then, we validate the proposed CAMixerSR with $\gamma = 0.5$ which uses self-attention for partial tokens. For *lightweight SR*, we employ five common-used validation datasets: Set5 [3], Set14 [40],

¹For model, γ is the averaged ratio for all CAMixer.

Table 4. Ablation study on window size on Urban100.

Window Size	#Params	<i>Original</i>		<i>CAMixer</i>	
		#MAdds	PSNR	#MAdds	PSNR
8×8	693K	47.5G	26.42	38.6G (81.3%)	26.41
16×16	765K	77.9G	26.65	53.8G (69.1%)	26.63
32×32	1340K	191.2G	26.82	109.7G (57.4%)	26.74




Figure 5. Ablation study on attention ratio γ .

BSD100 [26], Urban100 [12], and Manga109 [27]. For *large-image SR*, we utilize Flickr2K [22] (F2K) and DIV8K [10] (Test2K, Test4K, and Test8K) to generate the testing datasets as ClassSR. For *ODI SR*, we evaluate our models on the ODI-SR [6] testing set and SUN360 [36] dataset. For evaluation, we use PSNR and SSIM [35] and additionally distortion-weighted versions WS-PSNR [30] and WS-SSIM [46] for ODI-SR.

4.2. Ablation Study

Effects of Self-Attention. We first compare the window size M for the SA branch in Tab. 4. The model with window size 32 is about $2.5 \times$ larger than window 16 while improving 0.17dB on Urban100. The model with window 16 advances window 8 by 0.23dB but only increases 16G MAdds. Furthermore, when we set $\gamma = 0.5$, the 32×32 window encounters larger performance drops than the other two models because a large window is hard to classify. To this end, we use window 16×16 for better trade-offs between performance and computations. Moreover, we compare the different self-attention ratios γ in Fig. 5. For both lightweight SR and 2K SR, the computations grow linearly, while the PSNR grows faster when $\gamma < 0.5$ but slower when $\gamma > 0.5$. Hence, we manually select $\gamma = 0.5$ where PSNR is almost the same as $\gamma = 1.0$ but reduce half the computations of SA.

Effects of predictor components. Compared to the simple classifiers [4, 19] outputting only decision score, the proposed predictor generates more useful metrics (offsets and spatial/channel attention) for better partition and representation. We examine these extra components in Tab. 5. Removing any of them would result in huge performance drops. In detail, the offsets bring about 0.05dB gains and convolutional attentions jointly obtain 0.1dB improvements.

Effects of mixer mask m . To study the preference of the mixer mask dividing the hard and simple tokens, we visual-

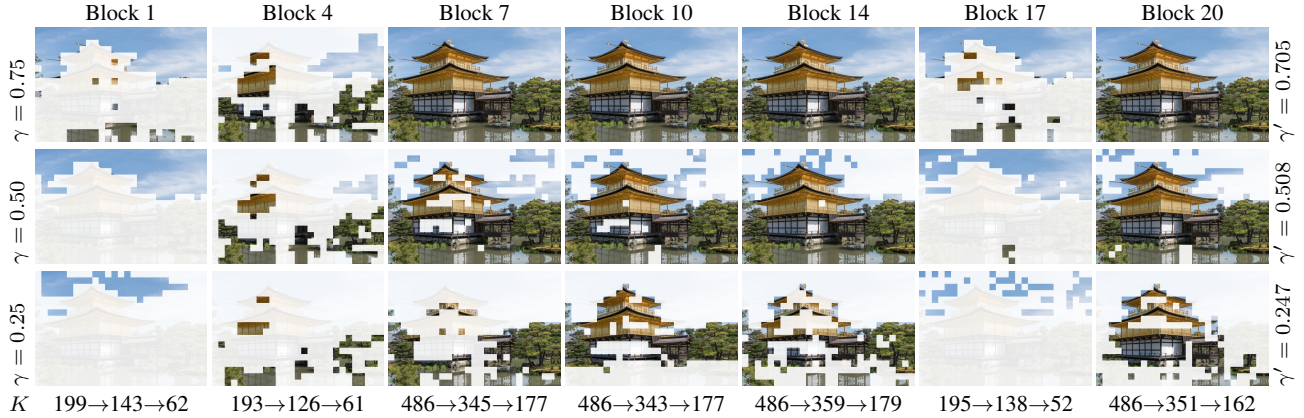


Figure 6. Visualization of the progressively classified tokens under varied ratio γ . The unmasked areas are processed by self-attention while the masked ones are processed by convolutional-attention. $K = \sum M$ is the window number of the SA branch. γ and γ' are the target and actual ratio of SA tokens, respectively.

Table 5. Ablation study on predictor components.

Offsets Δp	SA A_s	CA A_c	#Param	Set5	Urban100
✓	✓	✓	765K	32.51	26.63
✓	✓	✗	745K	32.43	26.50
✓	✗	✓	762K	32.37	26.46
✗	✓	✓	762K	32.46	26.60
✓	✗	✗	742K	32.31	26.45

Table 6. Ablation study on input conditions.

Local C_l	Global C_g	Window C_w	Set5	Urban100
✓	✓	✓	32.51	26.63
✓	✓	✗	32.49	26.55
✓	✗	✓	32.45	26.54
✓	✗	✗	32.42	26.46

ize the learned mask in Fig. 4. The proposed predictor can learn proper masks to assign complex tokens (e.g., building, ship, and bee) to the attention branch while plain tokens (e.g., sky and petal) to the convolution branch. In Fig. 6, we further compare the partition masks for varied blocks under different ratios γ . The token number K for SA is decreased according to the γ and the actual ratio is close to the target. This significantly satisfies our goal to control the inference complexity. Most blocks learn to use SA for complex content. Interestingly, they focus on semantically different contexts, e.g., *Block 7* preserving planets, *Block 14* preserving edges, and *Block 20* preserving buildings. Moreover, in Fig. 4, the C_l illustrates that the output of a convolution is distinctive from the attention. Thus, for *Block 1* & *17*, they employ SA for plain windows to harmonize the feature.

Effects of offsets Δp . In Tab. 7, we validate the performance changes to explore the effectiveness of offsets by adjusting the offset scalar r from 0 to 16. It can be observed that the model with $r = 8$ performs best. In contrast, when r is set to 4 or 16, the model performance negligibly im-

Table 7. Ablation study on offsets.

Offsets Scalar r	B100	Urban100	Manga109
0 (<i>No offsets</i>)	27.69 Δ	26.60 Δ	31.10 Δ
1	27.71 +0.02	26.60 +0.00	31.14 +0.04
4	27.70 +0.01	26.55 -0.05	31.14 +0.04
8	27.72 +0.03	26.63 +0.03	31.18 +0.08
16	26.68 -0.01	26.51 -0.09	31.06 -0.04




Figure 7. Visualization of varied offsets. The offsets are generated according to the Input image. The sampling step is 16.

proves but even drops. To comprehensively understand, we depict the offset vector in Fig. 7. Offsets with $r = 8$ are more reasonable since the most effective shifts are located at the edge areas and provide fine-grained partition.

Effects of input conditions. We ablate the input conditions (local, global, and window) in Tab. 6. Since our method is content-aware, we maintain the local content for all models. The absence of the global/window condition will cause 0.08dB drops on Urban100. Disabling both of them further decreases the PSNR to 26.46dB, which is 0.17dB lower than the initial model. The results indicate the combinations of multiple conditions can improve the predictor accuracy.

4.3. Large-Image SR

Following ClassSR [19] and ARM [4], we validate the efficiency of CAMixerSR with 2K-8K large-image SR task. Unlike previous work omitting lightweight models, we also add IMDN [14] and SwinIR-light [21] for reference. All

Table 8. Quantitative comparison (PSNR) on 2K-8K testsets. The accelerating framework, lightweight models, and proposed network are reported for general comparison. For CAMixerSR, the “+ ClassSR” uses CAMixerSR-S/M/B-Original as three branches with different complexities. For mixed strategies of Content-Aware and ClassSR, “⊕ ClassSR” uses CAMixerSR-B ($\gamma = 0.25/0.3/0.5$) as three branches for ClassSR. “⊕ CAMixer” represent using CAMixerSR-S/M/B ($\gamma = 0.5$) for ClassSR. For ClassSR-like method, the {tile, overlap} is {32, 2}, while for lightweight models, the {tile, overlap} is {64, 4}, to maintain the max FLOPs under 5G.

Method		#Params	F2K	#FLOPs	Test2K	#FLOPs	Test4K	#FLOPs	Test8K	#FLOPs
SRResNet	Original [20]	1.5M	29.01	5.20G (100%)	26.19	5.20G (100%)	27.65	5.20G (100%)	33.50	5.20G (100%)
	+ ClassSR [19]	3.1M	29.02	3.43G (66%)	26.20	3.62G (70%)	27.66	3.30G (63%)	33.50	2.70G (52%)
	+ ARM-L [4]	1.5M	29.03	4.23G (81%)	26.21	4.00G (77%)	27.66	3.41G (66%)	33.52	3.24G (62%)
	+ ARM-M [4]	1.5M	29.01	3.59G (69%)	26.20	3.48G (67%)	27.65	3.24G (62%)	33.50	2.47G (48%)
	+ ARM-S [4]	1.5M	28.97	2.74G (53%)	26.18	2.87G (55%)	27.63	2.77G (53%)	33.46	1.83G (35%)
	+ CAMixer	925K	29.15	3.24G	26.31	3.24G	-	-	-	-
RCAN	Original [43]	15.6M	-	-	26.39	32.60G (100%)	27.89	32.60G (100%)	33.76	32.60G (100%)
	+ ClassSR [19]	30.1M	-	-	26.39	21.22G (65%)	27.88	19.49G (60%)	33.73	16.36G (50%)
IMDN	Original [14]	715K	29.03	1.46G	26.19	1.46G	27.65	1.46G	33.57	1.46G
SwinIR-light	Original [21]	930K	29.24	2.10G	26.33	2.10G	27.79	2.10G	33.67	2.10G
CAMixerSR-S	Original	351K	29.12	894M (100%)	26.26	894M (100%)	27.73	894M (100%)	33.66	894M (100%)
	+ CAMixer	351K	29.08	652M (73%)	26.24	652M (73%)	27.70	652M (73%)	33.63	652M (73%)
CAMixerSR-M	Original	535K	29.20	1.37G (100%)	26.32	1.37G (100%)	27.80	1.37G (100%)	33.72	1.37G (100%)
	+ CAMixer	535K	29.18	1.03G (75%)	26.30	1.03G (75%)	27.79	1.03G (75%)	33.71	1.03G (75%)
CAMixerSR-B	Original	765K	29.31	1.96G (100%)	26.39	1.96G (100%)	27.89	1.96G (100%)	33.81	1.96G (100%)
	+ CAMixer	765K	29.30	1.49G (76%)	26.38	1.49G (76%)	27.87	1.49G (76%)	33.81	1.49G (76%)
	⊕ ClassSR [19]	820K	29.19	1.35G (69%)	26.32	1.37G (70%)	-	-	-	-
	+ ClassSR [19]	1711K	29.18	1.44G (73%)	26.28	1.52G (77%)	27.76	1.53G (78%)	33.65	1.48G (75%)
	⊕ CAMixer	1711K	29.17	1.04G (53%)	26.26	1.11G (57%)	-	-	-	-

Table 9. Quantitative comparison (average PSNR/SSIM, Parameters, and Mult-Adds) with state-of-the-art approaches for efficient image SR. The best results and the second best results are in **bold** and underline, respectively. Mult-Adds (MAdds) are measured under the setting of upscaling the image to 1280×720. More results are available in the *supplementary material*.

Method	Scale	#Params	#MAdds	Set5 [3]		Set14 [40]		BSD100 [26]		Urban100 [12]		Manga109 [27]	
				PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
IMDN [14]	×4	715K	40.9G	32.21	0.8948	28.58	0.7811	27.56	0.7353	26.04	0.7838	30.45	0.9075
LatticeNet [25]	×4	777K	43.6G	32.18	0.8943	28.61	0.7812	27.57	0.7355	26.14	0.7844	-	-
FDIWN [9]	×4	664K	28.4G	32.23	0.8955	28.66	0.7829	27.62	0.7380	26.28	0.7919	-	-
SwinIR-light [21]	×4	930K	61.7G	32.44	0.8976	28.77	0.7858	27.69	0.7406	26.47	0.7980	30.92	0.9151
ELAN-light [42]	×4	601K	43.2G	32.43	0.8975	28.78	0.7858	27.69	0.7406	26.54	0.7982	30.92	0.9150
NGswin [5]	×4	1019K	36.4G	32.33	0.8963	28.78	0.7859	27.66	0.7396	26.45	0.7963	30.80	0.9128
SwinIR-NG [5]	×4	1201K	63.0G	32.44	0.8980	28.83	0.7870	27.71	0.7411	26.54	0.7998	31.09	0.9161
DiVANet [2]	×4	939K	57.0G	32.41	0.8973	28.70	0.7844	27.65	0.7391	26.42	0.7958	30.73	0.9119
CAMixerSR	×4	765K	53.8G	32.51	0.8988	<u>28.82</u>	0.7870	27.72	0.7416	26.63	0.8012	31.18	0.9166

models are running under the same FLOPs restriction.

Quantitative results. In Tab. 8, we implement CAMixerSR-Small/Medium/Base with 36/48/60 channels. Generally, CAMixerSR remarkably surpasses other methods. Against classic SRResNet and RCAN accelerated by ClassSR [19] or ARM [4], the CAMixerSR-Original delivers similar or better restoration quality but is 3.2×-11× lighter. Moreover, our CAMixer models further save about 25% calculations. Overall, CAMixerSR (765K/747M) can compete with RCAN (15.6M/32.6G) for 2K-8K image restoration. In comparison with lightweight models SwinIR-light, our CAMixerSR renders 0.14dB PSNR improvement with fewer FLOPs or 51% fewer FLOPs with higher PSNR.

To fairly compare the accelerating framework, *i.e.*, ClassSR with the proposed Content-Aware Mixing, we apply ClassSR to CAMixerSR with three classes. The

ClassSR helps to reduce similar computations as Content-Aware but falls behind 0.12dB on F2K. This PSNR gap may result from to limited receptive field of small cropping size, which we will explore in further research. We also illustrate two potential combination ways for ClassSR and CAMixer. “⊕ ClassSR” indicates applying ClassSR on CAMixerSR-B, *i.e.*, using the same backbone but adjusting $\gamma = 0.25/0.3/0.5$ as three branches for ClassSR. “⊕ CAMixer” employs CAMixer with $\gamma = 0.5$ for CAMixerSR-S/M/B as three branches of ClassSR. As expected, the former approach maintains better restoration quality while the latter reduces more calculations. Conclusively, Content-Aware mixer is a better choice for large-image tasks and can work with other strategies, without any difficulty.

Qualitative results. In Fig. 8, we present the visual results of the CAMixerSR against other methods. CAMixerSR

Table 10. Quantitative comparison on ODI-SR, SUN 360, under Fisheye downsampling. “†” represents using augmentation dataset.

Method	Scale	#Params	ODI-SR [6]				SUN 360 Panorama [36]			
			PSNR	SSIM	WS-PSNR	WS-SSIM	PSNR	SSIM	WS-PSNR	WS-SSIM
Bicubic	×2	-	28.21	0.8215	27.61	0.8156	28.14	0.8118	28.01	0.8321
RCAN [43]	×2	15.4M	30.08	0.8723	29.49	0.8714	30.56	0.8712	31.18	0.8969
SRResNet [20]	×2	-	30.16	0.8717	29.59	0.8697	30.64	0.8714	31.20	0.8953
EDSR [22]	×2	40.7M	30.32	0.8770	29.68	0.8727	30.89	0.8784	31.42	0.8995
OSRT-light† [38]	×2	1.26M	30.42	0.8775	29.79	0.8735	31.00	0.8792	31.55	0.9004
CAMixerSR	×2	1.30M	30.46	0.8789	29.83	0.8747	31.04	0.8810	31.60	0.9020
Bicubic	×4	-	25.59	0.7118	24.95	0.6923	25.29	0.6993	24.90	0.7083
RCAN [43]	×4	15.6M	26.85	0.7621	26.15	0.7485	27.10	0.7660	26.99	0.7856
SRResNet [20]	×4	-	26.91	0.7597	26.24	0.7457	27.10	0.7618	26.99	0.7812
EDSR [22]	×4	43.1M	26.97	0.7589	26.30	0.7458	27.19	0.7633	27.10	0.7827
OSRT-light† [38]	×4	1.28M	27.17	0.7667	26.49	0.7526	27.48	0.7718	27.41	0.7911
CAMixerSR	×4	1.32M	27.19	0.7691	26.49	0.7538	27.48	0.7736	27.36	0.7916



Figure 8. Visual comparison of CAMixerSR with other methods for ×4 task on Test4K dataset.

recovers significantly clearer lattice content than other approaches. Moreover, the CAMixerSR with Content-Aware strategy induces less performance drop than ClassSR.

4.4. Lightweight SR

To evaluate the generality of CAMixer, we compare the proposed CAMixerSR with numerous SOTA lightweight models, including IMDN [14], LatticeNet [25], FDIWN [9], SwinIR-light [21], ELAN-light [42], NGswin [5], and DiVANet [2]. Tab. 9 shows the quantitative comparison. Inclusively, our CAMixerSR obtains superior restoration quality on all five benchmark datasets with moderate parameters and fewer computations. In particular, compared to the newest SwinIR-NG [5] using self-attention for all tokens, the CAMixerSR can attain better performance while saving 9.2G computations. The results strongly indicate that our CAMixer can work not only for large input images with plenty of plain areas but also for classic SR tasks.

4.5. Omni-Directional-Image SR

To understand the effectiveness of CAMixer under practical utilization with large-resolution inputs, we test our CAMixerSR on the Omni-Directional-Image (ODI) SR task, which

has 2K output and complex distortion. Specifically, we add the distortion map as an extra condition for predictors. In Tab. 10, we exhibit the qualitative comparison of our CAMixerSR with other methods. Our CAMixerSR obtains better restoration quality for almost all validation metrics. In detail, compared with 30× larger EDSR [22], CAMixerSR obtains 0.26dB improvement on SUN 360 benchmark. Moreover, CAMixerSR advances OSRT-light [38] by a maximum of 0.05dB/0.0028 (WS-PSNR/WS-SSIM) improvements without using additional training sets.

5. Conclusion

In this paper, we propose a content-aware mixer (CAMixer) that integrates the model accelerating scheme and token mixer design by routing neural operators (self-attention and convolution) of varied complexities according to the difficulty of content recovery. Particularly, the simple tokens are captured by convolution while the complex tokens are additionally processed by deformable self-attention. To improve the accuracy of routing, we introduce an effective predictor, which uses rich input conditions to generate multiple useful guiding information. Based on CAMixer, we build CAMixerSR, which reaches remarkable performance-calculation trade-offs on three SR tasks.

References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, pages 1122–1131, 2017. [5](#)
- [2] Parichehr Behjati, Pau Rodriguez, Carles Fernández, Isabelle Hupont, Armin Mehri, and Jordi González. Single image super-resolution based on directional variance attention network. *Pattern Recognition*, 133:108997, 2023. [7](#), [8](#)
- [3] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie-Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *BMVC*, pages 1–10, 2012. [5](#), [7](#)
- [4] Bohong Chen, Mingbao Lin, Kekai Sheng, Mengdan Zhang, Peixian Chen, Ke Li, Liujuan Cao, and Rongrong Ji. Arm: Any-time super-resolution method. In *ECCV*, pages 254–270, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [5] Haram Choi, Jeongmin Lee, and Jihoon Yang. N-gram in swin transformers for efficient lightweight image super-resolution. In *CVPR*, pages 2071–2081, 2023. [2](#), [7](#), [8](#)
- [6] Xin Deng, Hao Wang, Mai Xu, Yichen Guo, Yuhang Song, and Li Yang. Lau-net: Latitude adaptive upscaling network for omnidirectional image super-resolution. In *CVPR*, pages 9189–9198, 2021. [5](#), [8](#)
- [7] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE TPAMI*, 38(2):295–307, 2016. [2](#)
- [8] Zongcai Du, Ding Liu, Jie Liu, Jie Tang, Gangshan Wu, and Lean Fu. Fast and memory-efficient network towards efficient image super-resolution. In *CVPRW*, pages 853–862, 2022. [2](#)
- [9] Guangwei Gao, Wenjie Li, Juncheng Li, Fei Wu, Huimin Lu, and Yi Yu. Feature distillation interaction weighting network for lightweight image super-resolution. In *Proceedings of the AAAI conference on artificial intelligence*, pages 661–669, 2022. [7](#), [8](#)
- [10] Shuhang Gu, Andreas Lugmayr, Martin Danelljan, Manuel Fritsche, Julien Lamour, and Radu Timofte. Div8k: Diverse 8k resolution image dataset. In *ICCVW*, pages 3512–3516, 2019. [5](#)
- [11] Xiaotao Hu, Jun Xu, Shuhang Gu, Ming-Ming Cheng, and Li Liu. Restore globally, refine locally: A mask-guided scheme to accelerate super-resolution networks. In *ECCV*, pages 74–91. Springer, 2022. [1](#)
- [12] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *CVPR*, pages 5197–5206, 2015. [5](#), [7](#)
- [13] Zheng Hui, Xiumei Wang, and Xinbo Gao. Fast and accurate single image super-resolution via information distillation network. In *CVPR*, pages 723–731, Salt Lake City, USA, 2018. Computer Vision Foundation / IEEE Computer Society. [3](#)
- [14] Zheng Hui, Xinbo Gao, Yunchu Yang, and Xiumei Wang. Lightweight image super-resolution with information multi-distillation network. In *ACM MM*, pages 2024–2032, 2019. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [15] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2016. [3](#)
- [16] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *CVPR*, pages 1637–1645, 2016. [3](#)
- [17] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, pages 1646–1654, 2016. [2](#)
- [18] Fangyuan Kong, Mingxi Li, Songwei Liu, Ding Liu, Jingwen He, Yang Bai, Fangmin Chen, and Lean Fu. Residual local feature network for efficient super-resolution. In *CVPRW*, pages 766–776, 2022. [3](#)
- [19] Xiangtao Kong, Hengyuan Zhao, Yu Qiao, and Chao Dong. Classsr: A general framework to accelerate super-resolution networks by data characteristic. In *CVPR*, pages 12016–12025, 2021. [1](#), [2](#), [5](#), [6](#), [7](#)
- [20] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, pages 4681–4690, 2017. [7](#), [8](#)
- [21] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *ICCVW*, pages 1833–1844, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [22] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, pages 1132–1140, 2017. [1](#), [4](#), [5](#), [8](#)
- [23] Jie Liu, Jie Tang, and Gangshan Wu. Residual feature distillation network for lightweight image super-resolution. In *ECCVW*, pages 41–55, 2020. [2](#)
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. [5](#)
- [25] Xiaotong Luo, Yuan Xie, Yulun Zhang, Yanyun Qu, Cuihua Li, and Yun Fu. Latticenet: Towards lightweight image super-resolution with lattice block. In *ECCV*, pages 272–289, 2020. [7](#), [8](#)
- [26] David R. Martin, Charless C. Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, pages 416–425, 2001. [5](#), [7](#)
- [27] Yusuke Matsui, Kota Ito, Yuji Aramaki, Azuma Fujimoto, Toru Ogawa, Toshihiko Yamasaki, and Kiyoharu Aizawa. Sketch-based manga retrieval using manga109 dataset. *Multim. Tools Appl.*, 76(20):21811–21838, 2017. [5](#), [7](#)
- [28] Xuran Pan, Chunjiang Ge, Rui Lu, Shiji Song, Guanfu Chen, Zeyi Huang, and Gao Huang. On the integration of self-attention and convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 815–825, 2022. [4](#)
- [29] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *NeurIPS*, 34:13937–13949, 2021. [3](#), [4](#)

- [30] Yule Sun, Ang Lu, and Lu Yu. Weighted-to-spherically-uniform quality evaluation for omnidirectional video. *IEEE Sign. Process. Letters*, 24(9):1408–1412, 2017. 5
- [31] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *CVPR*, pages 2790–2798, 2017. 3
- [32] Yan Wang. Edge-enhanced feature distillation network for efficient super-resolution. In *CVPRW*, pages 777–785, 2022. 1, 3
- [33] Yan Wang, Yusen Li, Gang Wang, and Xiaoguang Liu. Multi-scale attention network for single image super-resolution. *arXiv preprint arXiv:2209.14145*, 2022. 3
- [34] Yan Wang, Tongtong Su, Yusen Li, Jiuwen Cao, Gang Wang, and Xiaoguang Liu. Ddistill-sr: Reparameterized dynamic distillation network for lightweight image super-resolution. *IEEE Transactions on Multimedia*, 2022. 1, 2, 4
- [35] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 5
- [36] Jianxiong Xiao, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Recognizing scene viewpoint using panoramic place representation. In *CVPR*, pages 2695–2702, 2012. 5, 8
- [37] Yifan Xu, Zhijie Zhang, Mengdan Zhang, Kekai Sheng, Ke Li, Weiming Dong, Liqing Zhang, Changsheng Xu, and Xing Sun. Evo-vit: Slow-fast token evolution for dynamic vision transformer. In *AAAI*, pages 2964–2972, 2022. 4
- [38] Fanghua Yu, Xintao Wang, Mingdeng Cao, Gen Li, Ying Shan, and Chao Dong. Osrt: Omnidirectional image super-resolution with distortion-aware transformer. In *CVPR*, pages 13283–13292, 2023. 8
- [39] Ke Yu, Xintao Wang, Chao Dong, Xiaoou Tang, and Chen Change Loy. Path-restore: Learning network path selection for image restoration. *IEEE TPAMI*, 44(10):7078–7092, 2021. 1, 2
- [40] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Curves and Surfaces - 7th International Conference*, pages 711–730, 2010. 5, 7
- [41] Dafeng Zhang, Feiyu Huang, Shizhuo Liu, Xiaobing Wang, and Zhezhu Jin. Swinfir: Revisiting the swinir with fast fourier convolution and improved training for image super-resolution. *arXiv preprint arXiv:2208.11247*, 2022. 3
- [42] Xindong Zhang, Hui Zeng, Shi Guo, and Lei Zhang. Efficient long-range attention network for image super-resolution. In *ECCV*, pages 649–667, 2022. 3, 5, 7, 8
- [43] Yulun Zhang, Kungpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, pages 294–310, 2018. 1, 2, 7, 8
- [44] Hengyuan Zhao, Xiangtao Kong, Jingwen He, Yu Qiao, and Chao Dong. Efficient image super-resolution using pixel attention. In *ECCVW*, pages 56–72, Glasgow, UK, 2020. Springer. 2
- [45] Lin Zhou, Haoming Cai, Jinjin Gu, Zheyuan Li, Yingqi Liu, Xiangyu Chen, Yu Qiao, and Chao Dong. Efficient image super-resolution using vast-receptive-field attention. In *EC-CVW*, pages 256–272, 2022. 3
- [46] Yufeng Zhou, Mei Yu, Hualin Ma, Hua Shao, and Gangyi Jiang. Weighted-to-spherically-uniform ssim objective quality evaluation for panoramic video. In *IEEE International Conference on Signal Processing*, pages 54–57, 2018. 5