

RealViformer: Investigating Attention for Real-World Video Super-Resolution

Yuehan Zhang[✉] and Angela Yao[✉]

National University of Singapore
{zyuehan, ayao}@comp.nus.edu.sg

Abstract. In real-world video super-resolution (VSR), videos suffer from in-the-wild degradations and artifacts. VSR methods, especially recurrent ones, tend to propagate artifacts over time in the real-world setting and are more vulnerable than image super-resolution. This paper investigates the influence of artifacts on commonly used covariance-based attention mechanisms in VSR. Comparing the widely-used spatial attention, which computes covariance over space, versus the channel attention, we observe that the latter is less sensitive to artifacts. However, channel attention leads to feature redundancy, as evidenced by the higher covariance among output channels. As such, we explore simple techniques such as the squeeze-excite mechanism and covariance-based rescaling to counter the effects of high channel covariance. Based on our findings, we propose RealViformer. This channel-attention-based real-world VSR framework surpasses state-of-the-art on two real-world VSR datasets with fewer parameters and faster runtimes. The source code is available at <https://github.com/Yuehan717/RealViformer>.

1 Introduction

Video super-resolution (VSR) recovers a high-resolution (HR) sequence of frames from its low-resolution (LR) counterpart. Recurrent convolutional approaches are commonly used in VSR with standard settings, assuming the LR frames are downsampled from HR frames with known kernels. [3, 10, 15]. However, in real-world VSR (RWVSR), the low-resolution videos are not simply downsampled versions of their high-resolution counterparts. Instead, they feature complex degradations that arise from the camera imaging system, compression, internet transmission, and other factors. These degradations make architecture design for RWVSR challenging, as artifacts and degradations tend to propagate and get exaggerated over the recurrent connection [5, 36].

Recently, transformer architectures have replaced convolutional architectures as state-of-the-art for standard VSR. While attention mechanisms have replaced convolution operations, most methods retain the recurrent connection to aggregate the information over time [22, 30]. Yet, such architectures do not always perform well on RWVSR. For instance, a Swin-based model [21] designed for standard VSR, when applied to a real-world input frame, as shown in Fig. 1a, generates more artificial lines than the convolutional model RealBasicVSR [5].

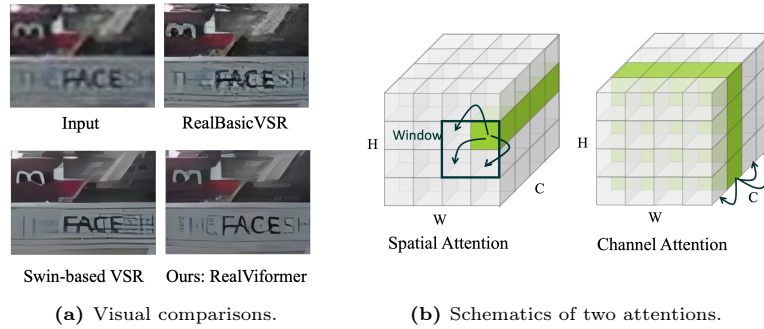


Fig. 1: (a) Designing a RWVSR transformer is not trivial. A Swin-based transformer suited for standard VSR hallucinates more lines than a RealBasicVSR, a convolutional state-of-the-art. We propose RealViformer based on our investigation of attention under the RWVSR setting. RealViformer generates details with fewer artifacts than RealBasicVSR [3] and the Swin-based VSR model. (b) Schematic for spatial and channel attention. Spatial attention aggregates features based on pixel representations. Channel attention takes $H \times W$ feature map for matching across channels.

Why should transformers perform well on (synthetic) standard VSR but poorly in real-world cases? We speculate that standard VSR transformers benefit from the similarity-based matching of the attention mechanism, which accurately aggregates information spatially and temporally [21, 22, 30]. However, when input degradation exists, the aggregated information becomes *less reliable* because the attention queries may be derived from both true source video and artifacts.

This work investigates and sheds light on the sensitivity of attention in real-world settings. We compare two covariance-based attention mechanisms used in low-level Transformers: spatial attention [21, 30] and channel attention [39]. Spatial attention takes pixel-wise features as keys and queries and estimates their covariances across spatial positions. The most popular form is window-based attention [6]; The shift-window scheme from Swin Transformers [23] enables the model to access distant spatial ranges [2, 21, 30] without computational blow-up. Spatial attention is widely used for video and image super-resolution, albeit in the standard setting. Channel attention [39] estimates covariances across channels (see Fig. 1b) and collapses the spatial extent of a feature map. It defines the number of queries and keys by channel numbers rather than spatial resolution. Although established for deblurring or denoising [39], channel attention’s efficacy in super-resolution remains to be determined.

Our experiments show that channel attention is less sensitive to artifacts than the spatial counterparts, resulting in higher performance gains in RWVSR. However, it is also revealed that channel attention leads to feature channels with higher covariance. From a learning perspective, a high covariance is undesirable because it is a strong indicator of feature redundancy [1, 8, 13]. Therefore, using channel attention naively will have limited improvements over existing RWVSR

state-of-the-art methods. Such a finding has wide-reaching impacts as channel attention is used increasingly for low-level vision.

To verify our findings, we explore established mechanisms to counter the effects of feature redundancy - simple techniques such as squeeze-excite and covariance-based rescaling improve the vanilla channel attention design. From these outcomes, we propose RealViFormer, a new transformer real-world VSR model. RealViFormer performs channel attention between the current frame feature and the propagated hidden state to limit model-produced artifacts. The model then reconstructs features through improved channel attention modules featuring squeeze-and-excite and covariance-based channel rescaling mechanisms. With our effective designs, RealViFormer achieves state-of-the-art performance with fewer parameters on challenging synthetic video datasets and two real-world video datasets collected from different scenes.

Summarizing our contributions in order of importance, our paper

- investigates the differences between spatial and channel attention for RWVSR. Spatial attention, although widely used, is revealed to be highly sensitive to the noise and degradations common in RWVSR sequences, while channel attention is more robust.
- reveals that naively applying channel attention increases channel covariance, which is problematic for learning; this overlooked fact has a wide-reaching impact as channel attention becomes more used in low-level vision.
- empirically verifies the negative effect of high channel covariance by countering it with established techniques, based on which we develop the RealViFormer for RWVSR. Our simple modification surpasses state-of-the-art despite using less compute.

2 Related Work

Standard Video Super-Resolution models focus on architecture design to use temporal information better [2–4, 15]. Previous research starts from slide-window-based [15, 33] to recurrent-based frameworks [3, 4] for using distant-frame information. Recent works introduce Transformer blocks into existing recurrent frameworks to overcome the locality limitation of convolution and accurately match abundant information for feature reconstruction [20, 22, 30].

Real-world video super-resolution focuses on modeling, removing, and limiting the impact of real-world degradations. Existing works are convolutional models and focus on designing losses or modules for degradation processing. DBVSR [28] explicitly estimates the degradation kernel through a sub-network. RealBasicVSR [5] tries to ‘clean’ artifacts through a processing module for BasicVSR [3]. In a similar approach, FastRealVSR [36] borrows an external pool of blur and sharpening filters to ‘clean’ the hidden states. Other recent works [16, 31] advance the synthesis method for paired training data. Instead, we focus on investigating the function of attention in RWVSR rather than architecture design or data synthesis.

Attention mechanisms have been widely applied for low-level vision tasks [9, 25, 43, 43]. Transformers with covariance-based attention are the most prevalent [21, 30, 39] for standard VSR. Most existing methods adopt shift-window-based spatial attention [23] to aggregate information from other positions within or across frames. In contrast, Restormer [39] computes the covariance among channels and shows its effectiveness for multiple image restoration tasks. More recent works stitch spatial and channel attention together to enlarge the receptive field [7, 32] for standard image super-resolution. Instead, we investigate attention mechanisms in terms of their sensitivity to real-world degradation for the first time and develop an effective real-world VSR Transformer.

3 Explorations on Attention

Sec. 3.1 defines the VSR task and the two attention mechanisms. Sec. 3.2 compares the channel and spatial attentions’ sensitivity to query artifacts and effects on real-world VSR performance. Sec. 3.3 reveals that channel attention leads to higher covariance among channels and explores mitigating options.

3.1 Preliminaries

Standard vs. Real-World VSR. Given a low-resolution (LR) video sequence with T frames $I^L \in \mathcal{R}^{T \times H \times W \times K}$, VSR models reconstruct a high-resolution sequence $I^H \in \mathcal{R}^{T \times sH \times sW \times K}$, where $H \times W$ is the input spatial resolution, K is the number of input channels and s is the scaling factor. In the standard setting of VSR, I_t^L , where $t \in \{0, \dots, T\}$, is assumed as the downsampled version of I_t^H , defined as: $I_t^L = (I_t^H) \downarrow_s$. Both training and testing datasets follow this formulation to generate LR-HR pairs given I^H . In a real-world setting, there is no closed formulation for the relationship between I_t^L and I_t^H due to the unknown distribution of real-world degradations. Real-ESRGAN [34] proposed a widely-used training setting that randomly applies synthesized blur, noise, compression, and resizing to the HR frames to generate paired LR frames with complex degradations. The testing datasets are either synthesized by the same pipeline in training or collected from diverse real-world sources [5, 37]. The synthesized testing sets have paired ground-truth sequences and are evaluated by full-reference metrics, *e.g.* PSNR and LPIPS [42]; real-world datasets are always without ground truth and require no-reference metrics, such as NRQM [24].

Attention Definitions. In Transformers, the attention modules project layer-normalized tensor $X \in \mathbb{R}^{C \times H \times W}$ to query Q , and tensor $Y \in \mathbb{R}^{\hat{C} \times \hat{H} \times \hat{W}}$ to key K and value V , where $H \times W$ and $\hat{H} \times \hat{W}$ are the spatial resolution of the normalized tensors.¹ The attention map \mathcal{A} is generated by calculating the covariance between Q and K , followed by a softmax function, before being applied to the value V to produce output O . Spatial and channel attention differ in tensor dimension taken for the covariance calculation.

¹ Note we define mutual attention as the default; self-attention is a special form of mutual attention where $Y = X$.

Spatial attention generates an $\mathbb{R}^{HW \times \hat{H}\hat{W}}$ attention map, \mathcal{A}_s , by computing the covariance between features in the query (Q_s) and key (K_s) at each spatial position. The query, key and values are computed as $Q_s = W_s^Q X$, $K_s = W_s^K Y$, $V_s = W_s^V Y$, where $W_s^Q \in \mathbb{R}^{D_s \times C}$, $\{W_s^K, W_s^V\} \in \mathbb{R}^{D_s \times \hat{C}}$, and D_s is the dimension of projections. The attention map \mathcal{A}_s and output attention features O_s are:

$$\mathcal{A}_s = \text{softmax}(Q_s^T K_s / \sqrt{D_s}), O_s = \mathcal{A}_s V_s^T, \quad (1)$$

where Q_s (K_s) is reshaped to matrix of $\mathbb{R}^{D_s \times HW}$ ($\mathbb{R}^{D_s \times \hat{H}\hat{W}}$). We omit the later reshaping operation for brevity. \mathcal{A}_s exhaustively relates every spatial location to every other location in space. A window-based implementation [21] limits the correlations to windows of $\omega \times \omega$, reducing the map size to $\mathbb{R}^{\omega^2 \times \omega^2}$.

Channel attention applies convolutions on X to get query (Q_c) and Y for key (K_c) and value (V_c) [39], after which features are the same in spatial resolution (assumed to be $H \times W$ in the following context). It estimates covariance across channels to yield a $C \times \hat{C}$ sized map \mathcal{A}_c and output features O_c :

$$\mathcal{A}_c = \text{softmax}(Q_c K_c^T / \alpha), O_c = \mathcal{A}_c V_c, \quad (2)$$

where Q_c is reshaped to $\mathbb{R}^{C \times HW}$, K_c and V_c are reshaped to $\mathbb{R}^{\hat{C} \times HW}$ and α is learnable scaling parameter. Because channel attention computes correlations between features of size $\mathbb{R}^{1 \times HW}$, it has a larger spatial context than the $\mathbb{R}^{D_s \times 1}$ -sized features in spatial attention (see Fig. 1b).

3.2 Attention in Real-World VSR

Sensitivity to Query Artifacts. In standard VSR, attention is used to aggregate information temporally and spatially [22, 30]. The assumption is that queries can match beneficial cues for super-resolution. But what if the queries themselves are unreliable? We speculate this to be true in real-world VSR, where inputs have artifacts and degradations.

How affected are spatial and channel attention mechanisms by query artifacts? Fig. 2 shows how we compare, using cosine similarity, the attention outputs based on queries from the same frame with and without certain degradations. Using encoding layers from a standard convolutional VSR model, we perform the attention operations defined in Sec. 3.1 by taking embeddings of frame I_t as query and I_{t-1} as key and value. I_t and I_{t-1} are clean without degradation other than downsampling, and we represent the output feature of the attention module as O . When additional degradation D_i is applied to I_t , we represent the corresponding output as O_{D_i} . A smaller deviation of O_{D_i} from O indicates lower sensitivity to artifacts in the query.

Tab. 1 shows the cosine similarity between O and O_{D_i} for spatial and channel attention modules. Curiously, O_{D_i} based on the channel attention module is more similar to the O matched by the degradation-free query, indicating that channel attention is less sensitive to query artifacts. Intuitively, the lower sensitivity of channel attention is related to the larger spatial context used for feature matching. Given a deep feature of size $\mathbb{R}^{C \times H \times W}$, channel attention uses feature sized




Fig. 2: Schematic for sensitivity comparison. I_{t-1}, I_t are downsampled but clean frames at times t and $t - 1$. $D_i(\cdot)$ apply degradations to I_t , where $D_i \in \{\text{blur, noise, compression}\}$. O and O_{D_i} are output features of the attention module. Queries are from the embedding at time t , and keys and values are from time $t - 1$. Higher cosine similarities S between attention output features O and O_{D_i} reflect less sensitivity to artifacts in queries.

Table 1: Cosine similarity between O and O_{D_i} , attention outputs without and with query degradation. Outputs of channel attention change less under query degradation.

Module	D_i		
	blur	noise	compression
Spatial attention	0.75	0.92	0.84
Channel attention	0.98	0.99	0.99

$\mathbb{R}^{1 \times HW}$ to calculate the covariance across channels. As such, feature aggregation is based on global information observed in a large normalized spatial context. Instead, the covariance of spatial attention is for features at each location sized $\mathbb{R}^{C \times 1}$, so it is likely more sensitive to local value changes from artifacts.

Impact on VSR Models. While channel attention output features are less sensitive to query artifacts than spatial attention, how well do they fare for real-world VSR? We experiment by incorporating these attention variants into VSR models. We focus specifically on recurrent pipelines due to their popularity in standard and real-world VSR [5, 22, 36]. In recurrent pipelines, artifacts in hidden states may propagate over time, get exaggerated, and negatively influence overall performance [5, 36]. Therefore, we explicitly add the attention module between features of the current frame and the propagated hidden state. Ideally, the added attention will select relative information from hidden states for the current frame and reduce the propagation of model-produced artifacts.

Fig. 3a shows the *baseline* model which concatenates f_t , the shallow feature of at time t , and \hat{h}_{t-1} , the spatially aligned hidden state at time $t-1$. We experiment with two variants by using a channel-attention module G_c or a spatial-attention module G_s to replace the concatenation in the baseline. As shown in Fig. 3b, the query is generated by f_t , and the key and value are predicted from \hat{h}_{t-1} . The attention output O_t^A is concatenated with f_t as the input for \mathcal{R} .

All models are trained on the REDS dataset [27] with the random degradation pipeline from Real-ESRGAN [34]. We test on REDS4 [27] with different types and extents of degradations, including Gaussian blur, Gaussian noise, and JPEG




Fig. 3: (a) The recurrent baseline in Sec. 3.2 has a shallow mapping module \mathcal{F} , reconstruction module \mathcal{R} , upsampling module \mathcal{U} and warping function W . W aligns the hidden state h_{t-1} to feature at t based on optical flow $s_{(t-1) \rightarrow t}^f$. All residual blocks are convolutional. The concatenation between f_t and \hat{h}_{t-1} are replaced with the spatial or channel attention modules in (b) to compare the effect of attention. (b) The attention module first applies layer normalization to f_t and \hat{h}_{t-1} and then performs channel or spatial attention according to Sec. 3.1. The output feature O_t^A concatenated with f_t is processed by the module \mathcal{R} in (a).




Fig. 4: Comparison of spatial and channel attention through impact on the performance of real-world VSR model. The Y-axis shows improvements compared to the convolutional baseline. A lower LPIPS score is better. The channel attention module is the best except for the PSNR score of highly blurred inputs.

compression. Fig. 4 displays changes in PSNR and LPIPS compared with the baseline model. Trained on the same dataset and degradation setting, channel attention between temporal information yields better objective and perceptual reconstruction quality than spatial attention and baseline for noise and JPEG compression inputs. Channel attention still achieves better scores for blurred inputs, except for the PSNR performance of severely blurred inputs.

3.3 Limitations of Channel Attention

The results in Sec. 3.2 indicate that channel attention is better suited for propagating information over time. However, it also has an inherent flaw – the correlation among output channels will increase since each channel in the attention output is a weighted summation over channels of the value V_c . We calculate the covariance following VICReg [1] for a quantitative measurement. Given deep

feature $z^n \in \mathbb{R}^{C \times H \times W}$, where $n \in \{1, N\}$ is the index of a sample, we reshape z^n to $\mathbb{R}^{C \times HW}$ and define the covariance matrix over $Z = \{z^1, \dots, z^N\}$ as:

$$\text{Cov}(Z) = \frac{1}{N-1} \sum_{n=1}^N (z_n - \bar{z})(z_n - \bar{z})^T, \quad (3)$$

where $\bar{z} = \frac{1}{N} \sum_{n=1}^N z_n$. We then define the indicator for the covariance matrix as $ac(Z) = \frac{1}{d} \sum_{i \neq j} |\text{Cov}(Z)|_{i,j}$, i.e. the average of absolute *off-diagonal* coefficients of $\text{Cov}(Z)$, where d is the number of off-diagonal coefficients. Function $ac(Z)$ encodes the covariance among feature channels. Taking O in Fig. 2 for comparison, $ac(O)$ with $N = 400$ based on channel attention is 0.87 and significantly higher than the input features (≈ 0.15). Instead, $ac(O)$ of spatial attention remains similar to the input features. Previous works on representation learning [1] and overfitting [8] propose that the high covariance of feature channels indicates redundancy and tends to hinder informative predictions.

Similarly, we speculate the redundancy effects will hinder the prediction of HR outputs when adopting channel attention in building blocks. For a closer look, we investigate the standard VSR task, which is artifact-free. Specifically, we build M_c and M_s by replacing the *Residual Block* in Fig. 3a with channel attention blocks [39] and spatial attention blocks [21]. Models are trained on the REDS dataset without extra degradation. For evaluation, we choose SSIM [35], which focuses on structural information. The SSIM score of M_c (0.8338) is lower than M_s (0.8432); $ac(\cdot)$ of the last features before the upsampling module is higher for channel attention than for spatial attention, i.e. 0.199 vs. 0.147.

We adopt two simple modifications to channel attention and boost informative features to verify the redundancy effect empirically. Fig. 5 shows the Improved Channel Attention (ICA) module. Our approach features two key steps to enhance the quality. First, we use the squeeze-and-excite mechanism to predict new information. The features are squeezed channel-wise to extract meaningful

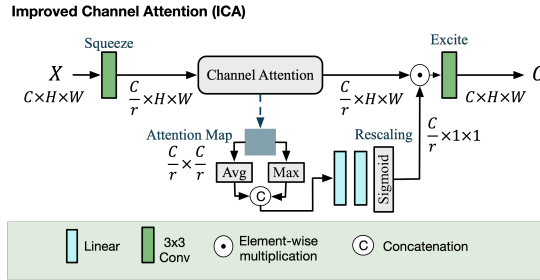


Fig. 5: Improved Channel Attention Module (ICA), showing self-attention for simplicity. The ‘squeeze’ convolution compresses the number of input feature channels $X \in \mathbb{R}^{C \times H \times W}$ by ratio r . The features are then rescaled by weights predicted from the $\frac{C}{r} \times \frac{C}{r}$ attention map before being expanded by the ‘excite’ convolution back to the original number of input channels.

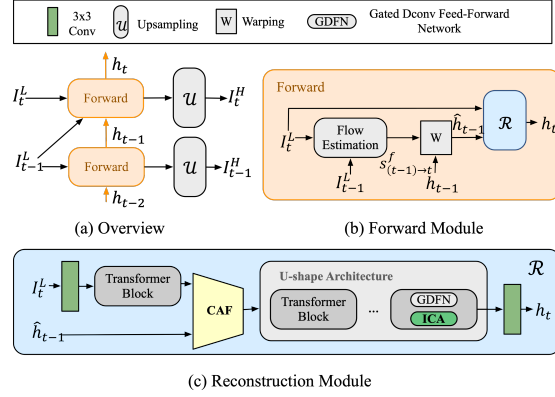


Fig. 6: The framework of RealViformer. (a) Overview of RealViformer, following a unidirectional recurrent framework. The outputs of the Forward module are propagated to the next time step and upscaled by module \mathcal{U} to get HR frames. (b) Explanation of the Forward module in (a), where W denotes the warping function. The reconstruction module \mathcal{R} takes current frame I_t^L and warped hidden state \hat{h}_{t-1} as inputs. (c) Reconstruction module \mathcal{R} . The shallow feature of I_t^L and \hat{h}_{t-1} are fused by CAF and then forwarded to Transformer blocks with U-shape connection [39]. Module GDFN follows Restormer [39]. Details of CAF and ICA modules are stated in Fig. 7 and Fig. 5.

information and then expanded into new channels based on the attention outputs. Secondly, we rescale channels in attention outputs by scalar weights predicted from the attention map. The attention map measures relationships across channels and encourages the associated convolutions in the ‘excite’ operation to yield more precise and discriminating features. Our designs are inspired by the SE Network [12] but with two key distinctions. First, we apply the squeeze-and-excite mechanism to generate new information, while SE Network uses it to help with scalar weight prediction. Second, the rescaling weights prediction takes the attention map as a cue rather than the naïve pooling results of channels. In Sec. 4, we show the efficacy of new designs on the real-world VSR task and compare channel correlation with and without new designs in ablations.

4 RealViformer

4.1 Overview

We apply our findings in Sec. 3.2 to the real-world VSR task as further support. To that end, we design RealViformer to incorporate channel attention as the basic processing module and the modification to boost informative features. Our emphasis here is not novel architecture design but a showcase of effectiveness brought by applying our findings in Sec. 3.2.

RealViformer is a recurrent Transformer network with channel attention modules. Fig. 6 shows the model architecture. To reduce the computational




Fig. 7: Details of Channel Attention Fusion (CAF) module. CAF gets the query from current frame feature f_t and {key, value} from hidden state \hat{h}_{t-1} . The attention output is concatenated with f_t to process for module output O_t .

cost, we follow a widely used recurrent architecture [3] in a unidirectional setting. The model first estimates optical flow $s_{(t-1) \rightarrow t}^f$ from I_{t-1}^L to I_t^L through Spynet [29] and warps previous hidden state h_{t-1} to current time step based on the flow. Frame I_t^L and spatially aligned hidden state \hat{h}_{t-1} are processed in the reconstruction module \mathcal{R} . The hidden state is updated with outputs of \mathcal{R} and further processed by the upsampling module \mathcal{U} to output high-resolution frames.

The reconstruction module \mathcal{R} uses channel attention in two ways. First, the Channel Attention Fusion (CAF) module fuses the temporal information to limit the produced artifacts in the hidden state. CAF queries the aligned hidden state \hat{h}_{t-1} by the shallow feature f_t . Secondly, we take the Improved Channel Attention Module (ICA) in Fig. 5 to build the Transformer blocks for better HR reconstructions.

4.2 Channel Attention Fusion (CAF) Module

As explored in Sec. 3.2, adding channel attention promotes the real-world VSR performance compared to spatial attention and the simple concatenation baseline. Thus, we keep this design in our model and put it as the Channel Attention Fusion (CAF) module for temporal aggregation. Fig. 7 show details of how CAF perform channel attention between f_t , the feature of the current frame, and \hat{h}_{t-1} , the aligned propagated hidden state. The query (Q_t) is generated as $Q_t = K_{3 \times 3} * \text{LayerNorm}(f_t)$, where $K_{3 \times 3} *$ refers to 3×3 convolution operation. Similarly, we process hidden state \hat{h}_{t-1} by $\tilde{h}_{t-1} = K_{3 \times 3}^d * K_{1 \times 1} * \text{LayerNorm}(\hat{h}_{t-1})$, where $K_{3 \times 3}^d$ is the 3×3 depth-wise convolution which double the channel number. Chunking \tilde{h}_{t-1} gives key (K_t) and value (V_t). The calculation of attention map $A_t \in \mathbb{R}^{C \times C}$ follows Eq. (2). The final output O_t is computed as $O_t = K_{1 \times 1} * K_{3 \times 3}^d * K_{1 \times 1} * \mathbf{C}[A_t V_t; f_t]$, where $\mathbf{C}[\cdot; \cdot]$ denotes concatenation.

4.3 Improved Channel Attention (ICA) Module

The design of ICA follows Fig. 5 for empirically verifying findings in Sec. 3.3. **Squeeze and Excite** follows the squeeze-and-excite mechanism in Fig. 5 and helps to predict channels with new information. It squeezes the channels of input feature $X \in \mathbb{R}^{C \times H \times W}$ by factor r . The attention map refers to self-attention on X and of size $\mathbb{R}^{\frac{C}{r} \times \frac{C}{r}}$. Excitation layers expand outputs back to size $\mathbb{R}^{C \times H \times W}$. **Correlation-based channel weighting** weights the output channel based on scalars predicted from the attention map. This design emphasizes channels for predicting more discriminative features in ‘excite’ operation. The attention map $A_r \in \mathbb{R}^{\frac{C}{r} \times \frac{C}{r}}$ is taken for calculating the average and max values along the rows. The average and max values are combined and mapped to weights of size $\mathbb{R}^{\frac{C}{r} \times 1}$ through linear layers and a sigmoid function.

4.4 Implementation Details

Model details. RealViFormer uses SPyNet [29] for flow estimation. After the CAF module, the reconstruction module \mathcal{R} has a three-level encoder-decoder architecture. From level 1 to level 3, there are [2,3,4] transformer blocks with [48,96,192] channels. There are [1,2,4] attention heads in the ICA, all with a squeeze factor of 4. Supplementary B.1 gives the detailed architecture of \mathcal{R} .

Training details. We train using the REDS dataset [27] and follow RealBasicVSR [5] in applying random combinations of blur, noise, JPEG compression, and video compression for synthesizing input degradations. We load 15 frames as an input sequence. The spatial size of inputs is cropped to 64×64 , and the batch size is 16. We use a pre-trained flow estimation model SPyNet, the parameter of which is fixed for the first 5K iterations and tuned with other modules later.

Following RealBasicVSR [5], we perform two-stage training. The first stage trains the model with a Charbonnier loss [19] and SSIM [35] loss for 300K iterations. In the second stage, the model is trained for another 130K iterations with the Charbonnier loss, SSIM loss, perceptual loss [18] and GAN loss [11] together, weighted by 1, 0.001, 1, and 0.005, respectively. The implementations of perceptual loss, GAN loss, and discriminator follow RealBasicVSR [5]. We implement all experiments on 4 Quadro RTX 8000 GPUs with PyTorch. Other details of the training settings are in Supplementary B.2.

4.5 Experimental Results

We compare our model on four datasets, VideoLQ [5], RealVSR [37], REDS4 [27] and UDM10 [38], with RealBasicVSR [5], Real-ESRGAN [34], BSRGAN [40], DBVSR [28], RealVSR [37], DAN [14], and RealSR [17]. VideoLQ and RealVSR are collected from real-world scenarios. VideoLQ is an unpaired dataset. RealVSR has same-size paired low-quality (LQ) and high-quality (HQ) frames. Our model super-resolves the input frames spatially. Although downsampling LQ enables paired data, it alters the original degradation. Thus, we test on the original LQ and do not use the HQ for evaluation. REDS4 and UDM10 have

Table 2: Quantitative comparisons with existing methods with **best** and **second-best** results. Our method achieves the best ILNIQE and NRQM scores over VideoLQ and RealVSR datasets and the best PSNR, SSIM, and LPIPS over synthetic datasets REDS4 and UDM10 with relatively few parameters and the lowest run-time.

	RealSR	DAN	RealVSR	DBVSR	BSRGAN	Real-ESRGAN	RealBasicVSR	Ours
Params (M)	16.7	<u>4.3</u>	2.7	25.5	16.7	16.7	6.3	5.3
Runtime (ms)	180	250	772	-	180	196	<u>73</u>	49
VideoLQ	ILNIQE↓	26.63	28.28	31.94	27.85	27.97	<u>25.98</u>	25.94
	NRQM↑	6.054	3.742	3.460	3.851	6.156	<u>6.306</u>	6.338
RealVSR	ILNIQE↓	32.81	32.29	34.39	-	32.65	<u>30.37</u>	28.61
	NRQM↑	5.610	3.523	3.795	-	6.152	<u>6.582</u>	6.588
REDS4	PSNR↑	22.02	22.67	18.30	22.35	22.94	<u>23.09</u>	23.34
	SSIM↑	0.5097	0.5571	0.4900	0.5530	0.5750	<u>0.6076</u>	0.6079
	LPIPS↓	0.5991	0.6315	0.7240	0.6211	0.3766	<u>0.2991</u>	0.2877
UDM10	PSNR↑	25.37	25.90	23.35	25.08	<u>25.97</u>	24.96	26.42
	SSIM↑	0.6658	0.7229	0.7115	0.7112	<u>0.7568</u>	0.7432	0.7609
	LPIPS↓	0.4811	0.4781	0.4761	0.4756	0.3388	<u>0.3209</u>	0.3063

ground-truth images, and we synthesize low-quality inputs through the same pipeline during training. All tested recurrent-based VSR models load the half sequence for videos in RealVSR and the whole sequence for others each time. The quantitative and qualitative results are discussed below.

Quantitative Results. For evaluation without reference, we apply the ILNIQE [41], the improved version of NIQE [26], and NRQM [24] metrics based on each output sequence’s first, middle, and last frames in RGB format [5]. These metrics appear less biased towards oversharpened features; details are given in Supplementary B.3. For evaluation on REDS4 and UDM10, we report more reliable metrics, PSNR, SSIM, and LPIPS [42]. We collect released models of all compared methods and generate sequences. As shown in Tab. 2, RealViformer performs better than other methods with smaller parameters than the most competitive RealBasicVSR [5] and the shortest runtime.

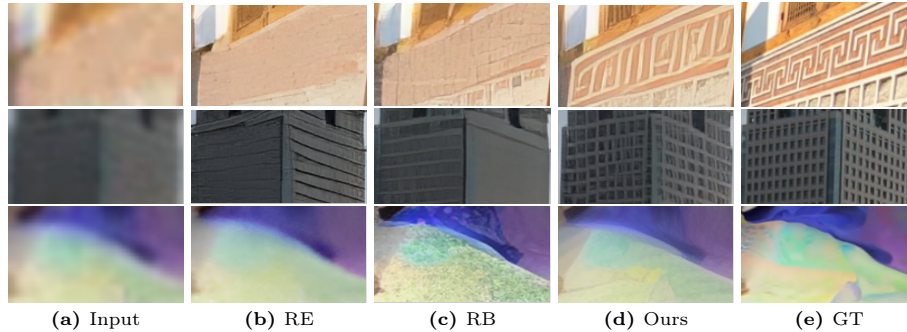


Fig. 8: Qualitative comparisons on synthetic datasets. Our method produces clearer than RealESRGAN (RE) [34] and RealBasicVSR (RB) [5] for very hard inputs.




Fig. 9: Qualitative comparisons on real-world datasets. Our method produces less high-frequency artifacts and overshoot effects than RealESRGAN [34] and RealBasicVSR [5].

Qualitative Results. We show qualitative comparisons on synthetic (see Fig. 8) and real-world (see Fig. 9) datasets. Compared to the listed methods, RealViformer generates clear structures with much fewer high-frequency artifacts. More visual comparisons are in Supplementary B.5.

User Study. We also performed a user study and asked 30 evaluators on Amazon MTurk to score reconstructions for 85 frames sampled from datasets VideoLQ [5], RealVSR [37], REDS4 [27] and UDM10 [38]. Each worker saw five HR results of the same frame and rated them based on the visual quality, from 1 (the worst) to 5 (the best); as shown in Fig. 10, our method surpasses BSRGAN [40], Real-ESRGAN [34], RealSR [17], and RealBasicVSR [5].




Fig. 10: User study results from 30 evaluators on 85 frames. Our method achieves the best among all five methods regarding mean opinion scores (MOS).

Table 3: Ablations of CAF and ICA modules. Channel-attention baseline (ch-baseline) performs better than the spatial-attention baseline (sp-baseline). The CAF module improves the performances for both datasets, and ICA further improves the performances to the state-of-the-art.

Method	CAF	ICA	VideoLQ NRQM↑	UDM10 LPIPS↓
Sp-baseline	-	-	6.061	0.3482
Ch-baseline	✗	✗	6.181	0.3085
RealViformer ⁻	✓	✗	6.196	0.2933
RealViformer	✓	✓	6.338	0.2877

Ablation Studies. We conduct ablations to validate the advantage of channel attention and the efficacy of the CAF and ICA. We build a spatial-attention baseline (sp-baseline) by replacing the reconstruction module of BasicVSR [3] with SwinIR [21]. Channel-attention baseline (ch-baseline) is built with the same




Fig. 11: (a) Visual comparison between RealViformer and its ablations. Red circles highlight the improved details. (b) Radial Power Spectrum (RPS) of model predictions. Using ICA improves the power of high-frequency components (blue region).

overall architecture as RealViformer but replaces CAF with simple concatenation and substitutes ICA with the channel attention block in Restormer [39]. RealViformer⁻ applies CAF on ch-baseline. All models are trained with the same settings in Sec. 4.4. We report NRQM scores for the VideoLQ [5] dataset and LPIPS scores for the UDM10 [38] dataset. As shown in Tab. 3, using the original channel attention module, ch-baseline already yields better performances than sp-baseline. The CAF and ICA modules further improve RealViformer to state-of-the-art performance with channel correlation of propagated information decreasing from 0.436 to 0.422. Fig. 11a visually compares RealViformer with its ablations. Applying CAF reduces artifacts, while ICA provides further improvements. Fig. 11b supplements the Radial Power Spectrum of model predictions. The blur region shows ICA increases the power in the high-frequency region.

5 Conclusion

This paper proposes a real-world VSR model, RealViformer, based on findings from investigating channel and spatial attention in a real-world setting. Explorations reveal that channel attention is less sensitive to the artifacts in query and better serves as a temporal aggregation module to limit model-produced artifacts in hidden states. Additionally, we observe the higher covariance of channel attention outputs and propose the Improved Channel Attention (ICA) Module with a squeeze-and-excite and a covariance-based rescaling mechanism. Based on our findings, we build RealViformer, a channel-attention-based recurrent model for real-world VSR. We propose the CAF module to limit artifact propagation and use the ICA module to achieve better reconstructions. RealViformer performs state-of-the-art on two real-world video datasets with fewer parameters and shorter runtime. On the other hand, we value our findings w.r.t. comparisons between channel and spatial attention and exploration of covariance in channel attention as inspiration for further real-world VSR research.

Acknowledgement This research is supported by the National Research Foundation, Singapore under its NRF Fellowship for AI (NRF-NRFFAI1-2019-0001).

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

1. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. arXiv preprint arXiv:2105.04906 (2021)
2. Cao, J., Li, Y., Zhang, K., Van Gool, L.: Video super-resolution transformer. arXiv preprint arXiv:2106.06847 (2021)
3. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: Basicvsr: The search for essential components in video super-resolution and beyond. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4947–4956 (2021)
4. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5972–5981 (2022)
5. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Investigating tradeoffs in real-world video super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5962–5971 (2022)
6. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12299–12310 (2021)
7. Chen, X., Wang, X., Zhou, J., Qiao, Y., Dong, C.: Activating more pixels in image super-resolution transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22367–22377 (2023)
8. Cogswell, M., Ahmed, F., Girshick, R., Zitnick, L., Batra, D.: Reducing overfitting in deep networks by decorrelating representations. arXiv preprint arXiv:1511.06068 (2015)
9. Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-order attention network for single image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11065–11074 (2019)
10. Fuoli, D., Gu, S., Timofte, R.: Efficient video super-resolution through recurrent latent space propagation. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). pp. 3476–3485. IEEE (2019)
11. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM **63**(11), 139–144 (2020)
12. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)
13. Hua, T., Wang, W., Xue, Z., Ren, S., Wang, Y., Zhao, H.: On feature decorrelation in self-supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9598–9608 (2021)
14. Huang, Y., Li, S., Wang, L., Tan, T., et al.: Unfolding the alternating optimization for blind super resolution. Advances in Neural Information Processing Systems **33**, 5632–5643 (2020)

15. Isobe, T., Jia, X., Gu, S., Li, S., Wang, S., Tian, Q.: Video super-resolution with recurrent structure-detail network. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16. pp. 645–660. Springer (2020)
16. Jeelani, M., Cheema, N., Illgner-Fehns, K., Slusallek, P., Jaiswal, S., et al.: Expanding synthetic real-world degradations for blind video super resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1199–1208 (2023)
17. Ji, X., Cao, Y., Tai, Y., Wang, C., Li, J., Huang, F.: Real-world super-resolution via kernel estimation and noise injection. In: proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 466–467 (2020)
18. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14. pp. 694–711. Springer (2016)
19. Lai, W.S., Huang, J.B., Ahuja, N., Yang, M.H.: Fast and accurate image super-resolution with deep laplacian pyramid networks. *IEEE transactions on pattern analysis and machine intelligence* **41**(11), 2599–2613 (2018)
20. Liang, J., Cao, J., Fan, Y., Zhang, K., Ranjan, R., Li, Y., Timofte, R., Van Gool, L.: Vrt: A video restoration transformer. *arXiv preprint arXiv:2201.12288* (2022)
21. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1833–1844 (2021)
22. Liang, J., Fan, Y., Xiang, X., Ranjan, R., Ilg, E., Green, S., Cao, J., Zhang, K., Timofte, R., Gool, L.V.: Recurrent video restoration transformer with guided deformable attention. *Advances in Neural Information Processing Systems* **35**, 378–393 (2022)
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: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
24. Ma, C., Yang, C.Y., Yang, X., Yang, M.H.: Learning a no-reference quality metric for single-image super-resolution. *Computer Vision and Image Understanding* **158**, 1–16 (2017)
25. Mei, Y., Fan, Y., Zhou, Y.: Image super-resolution with non-local sparse attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3517–3526 (2021)
26. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. *IEEE Signal processing letters* **20**(3), 209–212 (2012)
27. Nah, S., Baik, S., Hong, S., Moon, G., Son, S., Timofte, R., Mu Lee, K.: Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
28. Pan, J., Bai, H., Dong, J., Zhang, J., Tang, J.: Deep blind video super-resolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4811–4820 (2021)
29. Ranjan, A., Black, M.J.: Optical flow estimation using a spatial pyramid network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4161–4170 (2017)
30. Shi, S., Gu, J., Xie, L., Wang, X., Yang, Y., Dong, C.: Rethinking alignment in video super-resolution transformers. *arXiv preprint arXiv:2207.08494* (2022)

31. Song, Y., Wang, M., Yang, Z., Xian, X., Shi, Y.: Negvsr: Augmenting negatives for generalized noise modeling in real-world video super-resolution. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 10705–10713 (2024)
32. Wang, H., Chen, X., Ni, B., Liu, Y., Liu, J.: Omni aggregation networks for lightweight image super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22378–22387 (2023)
33. Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: Edvr: Video restoration with enhanced deformable convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
34. 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)
35. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
36. Xie, L., Wang, X., Shi, S., Gu, J., Dong, C., Shan, Y.: Mitigating artifacts in real-world video super-resolution models. *arXiv preprint arXiv:2212.07339* (2022)
37. Yang, X., Xiang, W., Zeng, H., Zhang, L.: Real-world video super-resolution: A benchmark dataset and a decomposition based learning scheme. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4781–4790 (2021)
38. Yi, P., Wang, Z., Jiang, K., Jiang, J., Ma, J.: Progressive fusion video super-resolution network via exploiting non-local spatio-temporal correlations. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3106–3115 (2019)
39. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5728–5739 (2022)
40. 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)
41. Zhang, L., Zhang, L., Bovik, A.C.: A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing* **24**(8), 2579–2591 (2015)
42. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
43. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: Proceedings of the European conference on computer vision (ECCV). pp. 286–301 (2018)