

Scale-Adaptive Feature Aggregation for Efficient Space-Time Video Super-Resolution

Zhewei Huang¹ Ailin Huang¹ Xiaotao Hu^{1,2} Chen Hu¹ Jun Xu^{2,3,*} Shuchang Zhou^{1,*}

¹Megvii Technology ²Nankai University

³Guangdong Provincial Key Laboratory of Big Data Computing, CUHK

hzwer@pku.edu.cn, csjunxu@nankai.edu.cn, shuchang.zhou@gmail.com

<https://github.com/megvii-research/WACV2024-SAFA>

Abstract

The Space-Time Video Super-Resolution (STVSR) task aims to enhance the visual quality of videos, by simultaneously performing video frame interpolation (VFI) and video super-resolution (VSR). However, facing the challenge of the additional temporal dimension and scale inconsistency, most existing STVSR methods are complex and inflexible in dynamically modeling different motion amplitudes. In this work, we find that choosing an appropriate processing scale achieves remarkable benefits in flow-based feature propagation. We propose a novel Scale-Adaptive Feature Aggregation (SAFA) network that adaptively selects sub-networks with different processing scales for individual samples. Experiments on four public STVSR benchmarks demonstrate that SAFA achieves state-of-the-art performance. Our SAFA network outperforms recent state-of-the-art methods such as TMNet [83] and VideoINR [10] by an average improvement of over 0.5dB on PSNR, while requiring less than half the number of parameters and only 1/3 computational costs.

1. Introduction

Constrained by the filming, processing, and distribution pipelines, the majority of videos are commonly stored and displayed at limited resolution and frame rates. Within the sphere of industry, space-time video super-resolution (STVSR) is practical for synthesizing smooth high-definition videos. New applications and software have continuously created demands for improving the performance and efficiency of STVSR models. From a scholarly perspective, the STVSR methods inspire insights into motion modeling for many video processing tasks [62,80]. Formally, given two low-resolution (LR) frames $\{I_0^{LR}, I_1^{LR}\}$ and a target time-step $t \in [0, 1]$, our goal is to synthesize the

*Corresponding authors.

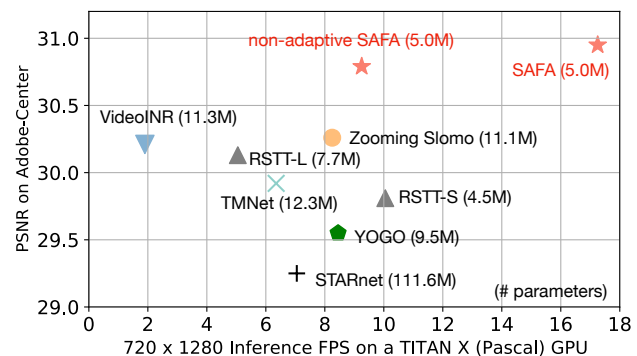


Figure 1. **Performance comparison on the Adobe240-Center dataset [72].** SAFA outperforms the other methods in terms of inference speed and PSNR (dB) metric.

high-resolution (HR) frame I_t^{HR} at moment t . The STVSR task contains two correlated subtasks of video frame interpolation (VFI) and video super-resolution (VSR). Both are among the most studied problems in Computer Vision. A traditional approach to efficiently aggregate information across frames involves estimating the dense displacement field. This is represented as $\mathbf{f}_{0 \rightarrow 1}$ for a pair of frames $\{I_0, I_1\}$, known as the optical flow. This technique is widely used in VSR. Besides, most VFI methods [1, 28, 33, 49, 85] focus on approximating the backward flow fields $\{\mathbf{f}_{t \rightarrow 0}, \mathbf{f}_{t \rightarrow 1}\}$ starting from the intermediate frame I_t . It is different from the former (because I_t is to be predicted) but very related. Successively performing VFI [1, 39] and VSR [5, 78] would overlook their inter-correlation on motion modeling.

Recent one-stage STVSR methods [10,81,83] have made great progress. They model temporal motion [28] by incorporating motion representation operators, e.g., deformable convolution [12] and DConvLSTM [81], into the convolutional neural network (CNN) backbones. Integrating additional components into these STVSR models would in-




Figure 2. $8\times$ VFI using RIFE [28] under different resolution scales of flow computing on the X4K1000FPS dataset [68]. Different video clips have respective suitable processing scales.

crease their complexity and overhead, which is not conducive to fulfilling the flexible requirements of diverse scenarios. Therefore, the complex structures adopted by previous methods might impede the progress of future STVSR methods. Motion estimation components have become more and more complex to manage various scenes and resolutions. Many models are trained on small (*e.g.*, 256×256) patches [86], yet relying on hand-crafted multi-stage and multi-scale strategies [39, 53] to process high-resolution videos [56, 68]. Even with these efforts, current method still perform considerably differently as the processing resolution changes and per video. We show the results of a toy experiment in Figure 2 and put the specific analysis in §3.2.

To mitigate the issue of excessive model complexity, it is imperative to develop efficient and streamlined one-stage STVSR models. As suggested by [5, 58], we decouple various functional components and integrate promising techniques from high-level vision. Since the input frames of STVSR are in low resolution, we focus on aggregating contextual features to supplement the lack of information in pixel-wise alignment [28, 39, 53]. To improve motion estimation and model efficiency, we propose a Scale-Adaptive Flow Estimation (SAFE) block to construct our model, Scale-Adaptive Feature Aggregation (SAFA).

Firstly, the flow estimation blocks recurrently update an intermediate hidden state that encodes motion information based on the contextual feature. Then, we use the dynamic routing technique [14] to adapt SAFA to diverse moving objects and motion scales. Specifically, we introduce a data-dependent scale selector using a Bernoulli distribution, enabling SAFE to well utilize the scalable advantage of motion estimation. The multi-branch selection reduces the computational overhead and helps SAFA fit the data more accurately. The propagation of image information and contextual features complement each other for

better STVSR performance. Experiments on these modules further validate that SAFA is a simple yet effective method that outperforms state-of-the-art STVSR methods. As shown in Figure 1, SAFA is faster and achieves over 0.5dB higher PSNR results than the comparison STVSR methods like VideoINR [10], TMNet [83], and RSTT [16] on the Adobe240 [72] dataset.

In summary, our contributions are three-fold:

- We introduce an innovative SAFE block that estimates motion in an iterative manner with trainable block-wise scale selection.
- With SAFE block, we propose an efficient SAFA method for STVSR. The modular design of SAFA makes it a simple yet efficient solution for STVSR with input-adaptive inference customization.
- Experiments show that SAFA quantitatively and qualitatively outperforms state-of-the-art STVSR methods, while exhibiting faster inference speed and fewer parameter amounts.

2. Related Work

2.1. Optical Flow Estimation

Optical flow estimation has garnered increasing attention in recent research. Pioneering works such as FlowNet [13] and FlowNet2.0 [31] have paved the way by learning to estimate dense optical flow fields from synthetic datasets. Efficient networks like PWC-Net [74] and LiteFlowNet [29] have been constructed by warping features and computing cost volume at various pyramid levels, thereby calculating optical flow in a coarse-to-fine manner [24, 59]. RAFT [76] has achieved a significant breakthrough by iteratively updating a flow field through a recurrent ConvGRU [66]. The iterative framework introduced by RAFT has further evolved into new variants [34, 73]. Transformer-based flow estimation models such as GMFlow [84], FlowFormer [27], and VideoFlow [67] have also shown promising advancements in the field. Moreover, several studies have successfully generalized optical flow techniques for downstream tasks [17, 86]. In this work, we introduce scale-adaptive path selection into task-oriented flow-based propagation for efficient STVSR.

2.2. Video Super-Resolution

In recent years, research on VSR has gradually shifted from sliding-window framework [77, 78, 86] to recurrent framework [5, 15, 26, 32]. The recurrent framework can leverage long-term motion information for better performance. The deformable-model-based alignment scheme [6, 78] is an extension of flow-based alignment and is widely used. Benefiting from in-depth research on long-term feature alignment and propagation, the work of BasicVSR [5,

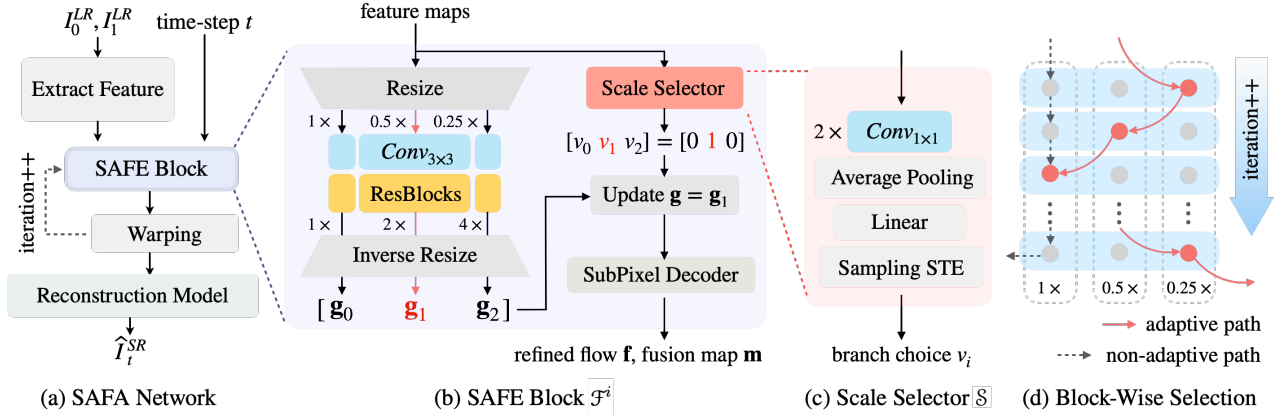


Figure 3. **Architecture of our SAFA.** (a) We design SAFA with low intra-module design complexity. The motion is iteratively estimated using SAFE blocks. (b) A SAFE block consists of several neural layers. We warp the image feature and update the hidden state \mathbf{g} . We use a scale selector \mathcal{S} to choose a suitable processing scale. (c) The scale selector is a small neural network with a differentiable sampling STE [2, 92]. (d) The adaptive selection of the sub-network is block-wise.

[7] has achieved very competitive performance. We refer readers to [7] for detailed technologies of long-term motion modeling. There are also some explorations of reusing optical flow between different frames [90]. These techniques are orthogonal to ours, as we focus more on streaming video processing schemes that do not rely on many distant frames.

2.3. Video Frame Interpolation

Most existing VFI methods can be divided into two main paradigms: kernel-based methods [54, 55] and flow-based methods [28, 33, 39, 49, 57, 88]. They share the idea that the pixels of the generated frames come from nearby regions of the adjacent frames. The kernel-based methods, *e.g.*, AdaConv [54] and SepConv [55], implicitly model the motion by estimating spatially adaptive kernels for different output pixels. On the other hand, flow-based methods explicitly model the motion of objects. Specifically, DVF [49] predicts the intermediate flow and occlusion mask jointly using CNN. SuperSloMo [33] uses two U-Net [61] models to estimate the bi-directional flows and occlusion mask successively. The research on optical flow is gradually deepening. ABME [57] explores asymmetric motion estimation. RIFE [28] and IFRNet [39] propose to build more refined models and learn intermediate flow estimation via knowledge distillation. AdaCoF [42] and EDSC [11] explore the integration of both paradigms under the multi-flow collaboration framework. Some recent works focus on specific scenarios, such as animations [9, 70] and near-duplicate photos [60]. Despite that VFI models are becoming more sophisticated, such as VFIfomer [50] and EMA [87], we notice that the critical path is still the feature extraction and motion estimation. With this clue, our designed model is modular with efficient flow-based feature aggregation.

2.4. Space-Time Video Super-Resolution

Standing on the early pioneer explorations [51, 63, 64], STVSR has achieved new developments in the era of deep learning. FISR [37] is among the pioneering work to unify VFI and VSR networks. As concurrent work, Xiang *et al.* [81] proposed to aggregate temporal information by a unified deformable ConvLSTM [12]. Haris *et al.* [18] proposed the STARnet to leverage mutually informative relationships between temporal and spatial dimensions, and fuse contexts at different resolution scales. To further improve STVSR performance, the work of STVUN [36], MB-Net [91], and YOGO [22] exploit the inter-correlation, interaction, and integration of propagation schemes between VSR and VFI tasks, respectively. Besides, TMNet [83] modulates the deformable convolution kernels for arbitrary-time frame interpolation. CycMu-Net [21] makes full use of spatial-temporal correlations via mutual learning between two processing dimensions. VideoINR [10] utilizes implicit neural representation for the arbitrary scale of temporal and spatial super-resolution. Our work focuses on improving the overall performance of STVSR through explicit and scale-adaptive feature aggregation.

3. Proposed Method

Here, we first overview SAFA for STVSR in §3.1. Then, we describe the scale inconsistency issue and scale-adaptive flow estimation in §3.2. Finally, we provide the training details in §3.3 for reproducible research.

3.1. Overall Network

Following previous settings [81, 83], we first use bicubic interpolation to upsample the input LR frames $\{I_0^{LR}, I_1^{LR}\}$, by a scaling factor of $4\times$. Here, we still keep the index

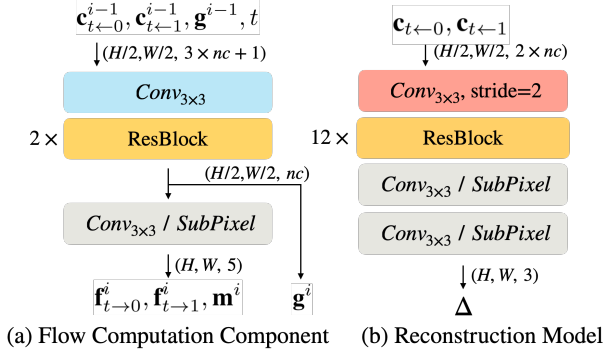


Figure 4. **Architecture of a single-path SAFE block \mathcal{F}^i (a) and reconstruction model \mathcal{R} (b).** Multiple inputs of each module are combined with concatenate operator. The hidden layers use PReLU [19] as nonlinear activation. nc is a hyper-parameter that controls the complexity of SAFA.

“LR” on the upsampled frames as $\{I_0^{LR}, I_1^{LR}\}$ for descriptive convenience. Thus, in SAFA, the input and output frames are of the same size. As shown in Figure 3, the proposed SAFA basically contains three parts: a spatial feature extractor f_θ , our proposed SAFE network \mathcal{F} for temporal feature aggregation, and a reconstruction module \mathcal{R} .

SAFA network firstly extracts the feature maps \mathbf{c}_0 and \mathbf{c}_1 of the two adjacent frames I_0^{LR} and I_1^{LR} , respectively, by the feature extractor f_θ . That is, $f_\theta(I_0^{LR}) = \mathbf{c}_0$, $f_\theta(I_1^{LR}) = \mathbf{c}_1$. The extracted feature maps will be employed to calculate the intermediate flow field and further propagated to fuse the latent feature of the frame at a target intermediate moment t . In the era of deep learning, there are many milestone works extracting useful representations from images [20, 41, 43, 69]. Inspired by the successful FPN networks [8, 44], we employ the ResNet-18 [20] as our feature extractor f_θ and integrate the feature maps output by the input stem, stage-1, and stage-2. To seamlessly integrate these feature maps of varying sizes, we use a 1×1 convolutional layer and bilinear upsampling to adjust the channel numbers and spatial sizes, respectively, illustrated in **Appendix**.

With the extracted feature maps $\{\mathbf{c}_0, \mathbf{c}_1\}$, we then perform temporal feature aggregation by SAFE blocks. The SAFE network \mathcal{F} estimates the bidirectional intermediate flow fields $\mathbf{f}_{t \rightarrow 0}$ and $\mathbf{f}_{t \rightarrow 1}$, as well as a fusion map $\mathbf{m} \in [0, 1]$ to model object occlusion [33, 49]. That is,

$$(\mathbf{f}_{t \rightarrow 0}, \mathbf{f}_{t \rightarrow 1}, \mathbf{m}) = \mathcal{F}(\mathbf{c}_0, \mathbf{c}_1, t). \quad (1)$$

To synthesize accurate target frame at moment t , we further propagate the LR frames $\{I_0^{LR}, I_1^{LR}\}$ and the feature maps $\{\mathbf{c}_0, \mathbf{c}_1\}$ to those at the target intermediate moment t . This is implemented by pixel-wise backward warping \overleftarrow{W} with the flow fields $\{\mathbf{f}_{t \rightarrow 0}, \mathbf{f}_{t \rightarrow 1}\}$, as follows:

$$\hat{I}_{t \leftarrow 0} = \overleftarrow{W}(I_0^{LR}, \mathbf{f}_{t \rightarrow 0}), \quad \hat{I}_{t \leftarrow 1} = \overleftarrow{W}(I_1^{LR}, \mathbf{f}_{t \rightarrow 1}), \quad (2)$$

$$\mathbf{c}_{t \leftarrow 0} = \overleftarrow{W}(\mathbf{c}_0, \mathbf{f}_{t \rightarrow 0}), \quad \mathbf{c}_{t \leftarrow 1} = \overleftarrow{W}(\mathbf{c}_1, \mathbf{f}_{t \rightarrow 1}), \quad (3)$$

After obtaining the propagated frames $\{\hat{I}_{t \leftarrow 0}, \hat{I}_{t \leftarrow 1}\}$ and feature maps $\{\mathbf{c}_{t \leftarrow 0}, \mathbf{c}_{t \leftarrow 1}\}$, we estimate a residual refinement Δ of the SR frame \hat{I}_t^{SR} by the reconstruction module:

$$\Delta = \mathcal{R}(\mathbf{c}_{t \leftarrow 0}, \mathbf{c}_{t \leftarrow 1}). \quad (4)$$

Finally, the reconstructed frame \hat{I}_t^{SR} is obtained by

$$\hat{I}_t^{SR} = [\mathbf{m} \odot \hat{I}_{t \leftarrow 0} + (1 - \mathbf{m}) \odot \hat{I}_{t \leftarrow 1}] + \Delta, \quad (5)$$

where \odot means pixel-wise multiplication.

The obtained target frame \hat{I}_t^{SR} is consisted of two parts: the pixel-wise fusion of the predicted intermediate frames $\{\hat{I}_{t \leftarrow 0}, \hat{I}_{t \leftarrow 1}\}$ propagated from the input LR images [49, 93] and the residual refinement Δ . \mathbf{f} , \mathbf{m} and Δ are of the same sizes with the HR frame. We have $0 \leq \mathbf{m} \leq 1$ and $-1 \leq \Delta \leq 1$, since \mathbf{m} and Δ are obtained by the Sigmoid and Tanh functions, respectively. The feature maps \mathbf{c}_0 and \mathbf{c}_1 are at $1/2$ resolution of the input frames. When warping the feature maps $\{\mathbf{c}_0, \mathbf{c}_1\}$, we need to resize the flow fields bilinearly to match their resolution. The reconstruction module \mathcal{R} is mainly consisted of ResBlocks [20] and SubPixel [65] operator, as shown in Figure 4.

3.2. Scale-Adaptive Flow Estimation

The versatile motion amplitudes and object sizes bring great challenges into flow estimation [53, 68, 76]. In high-level vision tasks like image classification, researchers often enlarge the size of input images to obtain more accurate prediction results [75]. However, the effective receptive fields of flow-based video synthesizing models [28] may not fully cover the objects of interests and their **large** motions in HR frames. We not only want the spatial information provided by the HR frame, but also want the receptive field of the model to cover the large motion of the object. This inspires us to design dynamic inference paths in the model to handle different scenarios.

Scale inconsistency issue. We clarify the scale inconsistency issue by a toy experiment. We use a pre-trained VFI model, RIFE [28], to perform $8 \times$ time scale VFI on the X4K1000FPS [68] benchmark. The details of the evaluation settings are consistent with those in [68]. We use RIFE [28] to infer the flow fields at different scales and then interpolate the intermediate frames. As shown in Figure 2, each video clip gains improvement on PSNR with different inference scales, even if RIFE [28] has equipped the coarse-to-fine flow estimation [25, 74]. By default, RIFE selects the $1/4 \times$ scale and achieves an average PSNR of 30.58dB. It further gains an improvement of about 0.4dB if we manually pick the inference scales with best results for each video clip. This demonstrates that it is useful to select an adaptive inference scale for each video clip. However, it

is laborious to manual inspection and selection of the inference scale clip by clip.

Intermediate flow estimation. We follow the trial-and-error manner for flow estimation [30,31,74,76]. That is, the model estimates an optical flow field and uses the current field to warp the spatial information of the original input frame or feature map, and so on. We encode the time-step t as a separate channel [28,39], and feed it into the flow estimation module along with the image features. We employ a bidirectional network to project the features $\{c_0, c_1\}$ from both directions to that at the moment t , under an iterative trial-and-error manner. Formally, we denote the i^{th} SAFE block as \mathcal{F}^i , as shown in Figure 4. They share the same model structure, but with independent parameters:

$$\mathcal{F} = \{\mathcal{F}^1, \mathcal{F}^2, \dots, \mathcal{F}^K\}. \quad (6)$$

Here, we set $K = 6$ in SAFA.

As shown in Figure 3, we iteratively update the hidden state \mathbf{g} that encodes motion. In each iteration, we use the current estimated flow field to propagate extracted feature c to obtain the bidirectional feature maps $\{c_{t \leftarrow 0}^{i-1}, c_{t \leftarrow 1}^{i-1}\}$ of the intermediate frame at a target moment t . Then these feature maps are input to the next SAFE block. The calculation process of K iterations is as follows:

$$\mathbf{g}^0 = \text{Conv}_{3 \times 3}(\text{Concat}(c_0, c_1)), \quad (7)$$

$$\mathbf{g}^i = \mathcal{F}^i(c_{t \leftarrow 0}^{i-1}, c_{t \leftarrow 1}^{i-1}, \mathbf{g}^{i-1}, t), \quad i = 1, 2, \dots, K. \quad (8)$$

Finally, we get iteratively refined flow fields $\{f_{t \rightarrow 0}, f_{t \rightarrow 1}\}$ and an occlusion mask map \mathbf{m} , which will be used to propagate spatial information of input frames and their feature maps, as described in §3.1.

Input-adaptive scale selection. To endow SAFA with the ability of dynamic scale selection, we propose a SAFE block to update the hidden states $\mathbf{g} = \mathcal{F}(x)$ on different scales, and select the one with the best performance, as shown in Figure 3. The computational costs for calculating the state $\mathbf{g}^{1/s}$ with $1/s$ resolution are $1/s^2$ of that in the original resolution. This provides SAFA an appropriate path to predict the intermediate frame at a target moment t from two given frames.

We introduce the STE sampling [2,23,92] for multi-branch selection. The neural network of scale selector \mathcal{S} estimates the selection probability P_i of B branches, where $i = 0, 1, \dots, B-1$ and $0 \leq P_i \leq 1$. We then choose a branch based on P_i :

$$\text{Forward} : p_i \sim \text{Bernoulli}(P_i), \quad p_i \in \{0, 1\}, \quad (9)$$

$$v_0 = p_0, \quad v_1 = (1 - p_0) * p_1, \quad (10)$$

$$v_2 = (1 - p_0) * (1 - p_1) * p_2, \quad \dots, \quad (11)$$

$$\mathbf{g} = \sum_{i=0}^{B-1} v_i * \mathbf{g}_i, \quad (12)$$

$$\text{Backward} : \frac{\partial o}{\partial P_i} = \frac{\partial o}{\partial p_i}, \quad (13)$$

where we calculate \mathbf{g} in multiple branches gated by the branch choice v_i , o is the objective function. Essentially, the above formula performs multiple dual-branch selections to construct a computable multi-branch selection.

In practice, we use three branches of different scales $\{\mathbf{g}^1, \mathbf{g}^{1/2}, \mathbf{g}^{1/4}\}$, *i.e.*, $B = 3$. During training, our multi-branch design increases the training overhead by 30% compared to that with a single branch \mathbf{g}^1 . In the inference stage, only one branch needs to be calculated. The parameters of the different branches here are shared. On the Adobe240 benchmark [72], the runtime of SAFA is 60% of non-adaptive SAFA. The scale selector \mathcal{S} occupies only 2% of computation overhead by SAFA model.

3.3. Implementation Details

SAFA model is optimized from scratch (random initialization) by Adam [38]. We train SAFA end-to-end using L_1 loss function with 32×32 down-sampled patches for 600,000 iterations. The batch size in training is 24. We gradually reduce the learning rate from 2×10^{-4} to 0 using cosine annealing scheme [81,83] during the training process. We use a sliding window with a length of 9 to select frames from videos. In one training sample, we select the first and last frames of the sliding window as the inputs and the intermediate frames are taken as the ground truths. The corresponding time-step t is from $\{0, \frac{1}{8}, \frac{2}{8}, \dots, 1\}$. Note that we randomly select one intermediate frame for STVSR prediction during the training stage. We randomly augment the training sample by horizontal and vertical flipping, rotating (90° , 180° and 270°), temporal order reversing. The training is performed on four Pascal TITAN X GPUs, which takes about 50 hours. For comparison, VideoINR [10] and TMNet [83] require more than one week for training.

4. Experiments

4.1. Datasets and Evaluation Metrics

For STVSR methods, a typical training or testing sample of video frame tuples is $\{I_0^{LR}, I_t^{HR}, I_1^{LR}, t\}$. Here, the input LR frames I_0^{LR} and I_1^{LR} are bicubically down-sampled by $0.25 \times$ from the corresponding HR frames I_0^{HR} and I_1^{HR} , respectively. Following VideoINR [10], we train the comparison methods on the Adobe240 training dataset [72], which contains 100 720p videos and each video has about 3,000 frames. Our evaluation is on the following datasets:

Vid4 [46] is a popular test set containing 171 frames of 480p. On this dataset, we conduct STVSR experiments for single-frame interpolation. Following VideoINR [10], the metrics reported on Vid4 differ from some previous literatures [81,83] due to using a smaller sliding window.

Table 1. **Quantitative comparison** on the Vid4 [45], GoPro [52], and Adobe240 [72] datasets. We omit some results of methods that can only synthesize frames at fixed times $t = 0.5$. Some of the previous methods are reported by VideoINR [10].

VFI Method	VSR Method	Vid4		GoPro-Center		GoPro-Average		Adobe-Center		Adobe-Average		# Param (M)
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
DAIN [1]	EDVR [78]	23.48	0.6547	28.01	0.8239	26.37	0.7964	27.06	0.7895	26.01	0.7703	44.7
DAIN [1]	BasicVSR [5]	23.43	0.6514	28.00	0.8227	26.46	0.7966	27.07	0.7890	26.23	0.7725	30.3
IFRNet [39]	EDVR [78]	23.68	0.6515	28.49	0.8379	26.41	0.7980	27.31	0.7981	26.12	0.7710	25.7
IFRNet [39]	BasicVSR [5]	23.76	0.6603	28.55	0.8392	26.58	0.8012	27.44	0.8005	26.35	0.7769	11.3
Zooming SloMo [81]		25.72	0.7717	30.69	0.8847	-	-	30.26	0.8821	-	-	11.1
TMNet [83]		<u>25.96</u>	<u>0.7803</u>	30.14	0.8692	28.83	0.8514	29.41	0.8524	28.30	0.8354	12.3
RSTT-L [16]		25.94	0.7801	30.37	0.8762	-	-	30.13	0.8754	-	-	7.7
VideoINR-fixed [10]		25.78	0.7730	30.73	0.8850	-	-	30.21	0.8805	-	-	11.3
VideoINR [10]		25.61	0.7709	30.26	0.8792	29.41	0.8669	29.92	0.8746	29.27	0.8651	11.3
non-adaptive SAFA		25.82	0.7759	<u>30.97</u>	<u>0.8871</u>	<u>30.01</u>	<u>0.8719</u>	<u>30.79</u>	<u>0.8854</u>	<u>29.92</u>	<u>0.8736</u>	5.0
SAFA		25.98	0.7807	31.28	0.8894	30.22	0.8761	30.97	0.8878	30.13	0.8782	5.0

Adobe240 [72] contains 17 videos in its test subset. We take the first frame out of every eight frames (*i.e.*, 1st, 9th, 17th frame,...) to produce the input LR frames, and use them to interpolate the intermediate frames.

GoPro [52] contains 11 720p and 240FPS videos. These videos are mostly street scenes captured by action cameras. We pre-process these videos in the same way as the Adobe240 [72] test set. We do not use the training set [52].

X4K1000FPS is a high frame rate 4K dataset [68] containing 15 street scenes for testing.

For Adobe240 [72] and GoPro [52] datasets, we separately evaluate the average metrics of the center (*i.e.*, 1st, 5th, 9th, ...) frames and all output frames, which are denoted as *-Center* and *-Average*, respectively. We use the metrics of Peak Signal-to-Noise Ratio (PSNR) and structural similarity (SSIM) [79] for quantitative evaluation. Following previous works [10, 81, 83], the generated images will be evaluated on Y channel of YCbCr space. All the methods are tested on a Pascal TITAN X GPU. To report the runtime, we calculate the average process time for 100 runs after a warm-up process.

4.2. Comparisons to State-of-the-Arts

Comparison methods. We compare SAFA with six two-stage “VFI+VSR” methods and four one-stage state-of-the-art STVSR methods [10, 16, 81, 83]. We reproduce the other methods following VideoINR [10]. For two-stage methods, we employ DAIN [1] and IFRNet [39] for VFI, and employ Bicubic Interpolation, EDVR [78], and BasicVSR [5] for VSR, both trained on the Adobe240 training set [72]. The one-stage VideoINR [10], RSTT [16], and Zooming SloMo [81] are trained in similar settings with SAFA. With an additional training stage on Vimeo90K [86], TMNet [83] is fine-tuned on the Adobe240 [72] training set. As suggested in [10], VideoINR has a variant fixing the interpolation time as $t = 0.5$, denoted as VideoINR-fixed.

Quantitative results. As shown in Table 1, the two-stage

Table 2. **Quantitative comparison for different time scales** on the GoPro dataset [52].

Time Scale	TMNet [83]	VideoINR [10]	SAFA
6×	30.49dB	30.78dB	31.67dB
8×	28.83dB	29.41dB	30.22dB
12×	26.38dB	27.32dB	27.97dB
16×	24.72dB	25.81dB	26.32dB

methods consuming more parameters perform worse than the one-stage methods on STVSR. This is consistent with previous literature [71, 81, 83] that demonstrated the benefits of one-stage model design for STVSR methods.

For one-stage methods, SAFA enables arbitrary time-step interpolation and achieves clearly better performance than the comparison methods on Gopro [52] and Adobe240 [72]. Due to small inter-frame object motion, SAFA is only on par with TMNet [83] on Vid4 [45].

The original *-Average* evaluation regime can be taken as 8× time scale evaluation. To study the ability of different methods on modeling time-step t , we also conduct 6×, 12×, and 16× experiments on GoPro [52]. As shown in Table 2, SAFA consistently achieves substantial advantages on all regimes over TMNet [83] and VideoINR [10].

Visual effects. In Figures 5 and 6, we compare SAFA with TMNet [83] and VideoINR [10]. SAFA recovers better the details of different scenes, especially on the frames at $t = 0.5$. The reconstruction of frames at $t = 0.5$ is relatively difficult because it is far from both input frames. We provide some video results in *Supplementary Materials*.

Runtime. We test all comparison methods on the same platform. As shown in Figure 7, SAFA has about 0.25× inference time when compared to TMNet [83]. The most efficient two-stage method, *i.e.*, IFRNet [39] + BasicVSR [5], is considerably slower than SAFA.




Figure 5. **Qualitative comparisons on the GoPro dataset [52].** (a) VideoINR [10], (b) SAFA, (c) ground truth. We perform $8\times$ time scale interpolation and $4\times$ super-resolution. Three of the eight results are shown. The PSNR annotations are calculated on the full 720p frame. We crop 120×300 patches for visualization.

4.3. Ablation Study

We perform extensive ablation studies shown in Table 3.

(a) Feature extractor. The feature extractor in SAFA is built on out-of-the-box ResNet-18 [20]. We denote the default setting as $\{R18: \text{stem}, s1, s2\}$, indicating that SAFA fuses the features output by the stem and the first two stages. Using only deep features **(a1, a2)** brings severe performance degradation. The reason is that more down-sampling layers lead to location confusion [48] and make feature alignment difficult. Fusion of multi-level features improves the performance of SAFA **(a1 vs. a3)**. However, advanced models [82] **(a4, a5)** that perform better in high-level tasks may not be suitable for the STVSR task. We found the bottleneck block of ResNet-50 can lead to a performance drop in STVSR. More encoding layers may require additional techniques to preserve spatial information.

(b) Information aggregation. SAFA **(b1)** computes optical flow for explicit propagation of both image information and features, while implicit propagation [5,81] usually works on features. The variants **(b2, b3)** using only image or feature propagation suffer from clear performance drop.

(c) Setting of flow estimation. In the modular design of SAFA, we can perform speed-accuracy trade-off by adjusting the number of flow estimation iterations **(c1-c4)**.

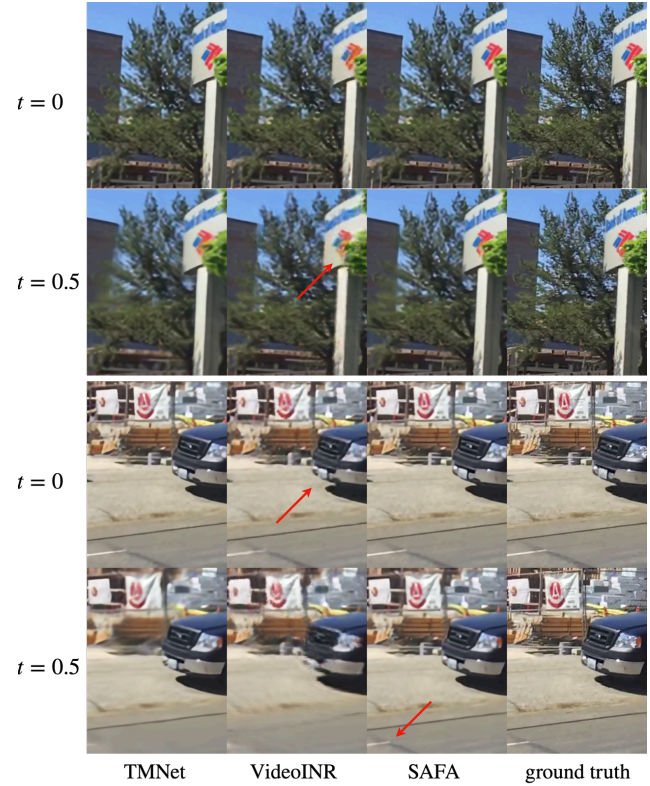


Figure 6. **Qualitative comparisons on the Adobe240 dataset [72].** We perform $8\times$ time scale interpolation and $4\times$ super-resolution. Two of the eight results are used for presentation. We crop 320×240 patches for visualization.

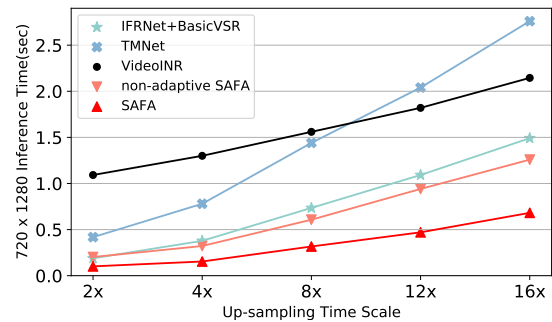


Figure 7. **Inference time on different up-sampling time scales.** SAFA is efficient and has high accuracy.

Similar to the finding in the supervised optical flow estimation [76], the iterative estimation indeed improves the STVSR performance of SAFA. However, sharing weights between different SAFE blocks **(c5)** brings about 0.2dB performance drop. Multi-branch selection contributes 0.31/0.18dB improvements and increases the speed **(c4 vs. c6, Figure 7)**. The weights between different scales can also be shared, which reduces model parameters and makes SAFA more effective **(c4 vs. c7)**. Directly optimizing the

Table 3. **Ablation studies** on the Gopro-Center [52] and Adobe240-Center [72] datasets. The default settings of SAFA are marked using gray backgrounds.

Setting	GoPro PSNR	Adobe240 PSNR	# Param (M)
<i>Feature Extractor f_θ</i>			
a1: {R18: s2}	31.09	30.77	5.0
a2: {R18: s3}	31.01	30.68	7.0
a3: {R18: stem, s1, s2}	31.28	30.97	5.0
a4: {R50: stem, s1, s2}	31.23	30.94	6.0
a5: {ResNeXt101: stem, s1, s2}	31.17	30.87	7.0
<i>Information Aggregation</i>			
b1: image + feature fusion	31.28	30.97	5.0
b2: image fusion	30.57	30.25	3.3
b3: feature fusion	31.15	30.77	5.0
<i>Setting of SAFE block</i>			
c1: $K = 1$ iteration	30.73	30.61	2.5
c2: $K = 2$ iterations	31.04	30.77	3.0
c3: $K = 4$ iterations	31.21	30.90	4.0
c4: $K = 6$ iterations	31.28	30.97	5.0
c5: + cross-block weight sharing	31.07	30.81	2.5
c6: non-adaptive flow estimation	30.97	30.79	5.0
c7: w/o cross-scale weight sharing	31.25	30.91	6.8
c8: w/o scale selector	30.89	30.71	5.0
<i>Initialization of Feature Extractor</i>			
d1: from scratch	31.28	30.97	5.0
d2: supervised pre-trained [20]	31.34	31.02	5.0
d3: DINO pre-trained [4]	31.38	31.12	5.0
<i>Peak Learning Rate β, number of channels nc</i>			
e1: $\beta = 2 \times 10^{-4}$, $nc = 80$	31.28	30.97	5.0
e2: $nc = 60$	31.17	30.85	2.8
e3: $\beta = 1 \times 10^{-3}$	31.22	30.91	5.0
e4: $\beta = 4 \times 10^{-5}$	30.94	30.62	5.0

Table 4. **Statistics of scale selection** for different input resolution and performance impact.

Input Resolution	4K	2K	540p	270p
Ratio of scale=1/4	0.71	0.58	0.49	0.41
Ratio of scale=1/2	0.17	0.18	0.15	0.16
Ratio of scale=1	0.12	0.24	0.36	0.43
Fix scale=1 PSNR	27.05	25.50	23.55	21.62
SAFA PSNR	32.28	29.02	24.90	22.11
VideoINR [10] PSNR	29.70	26.64	23.09	20.56

scale selector (**c8**) rather than associating it with the inputs [47] brings a performance drop, which confirms the importance of input-adaptive scale selector.

We study how SAFE block is benefited from selecting model branches. We perform $8 \times$ time scale and $4 \times$ space scale STVSR on the X4K1000FPS dataset [68]. We bilinearly down-sample the original frames to get video clips of different resolution. As shown in Table 4, SAFA tends to select larger scales when processing videos in higher resolution. When fixing scale=1, the receptive field may not cover large motions, bringing performance degradation.

(d) Benefited from pre-trained methods. Based on out-of-the-box CNN architectures, the feature extractor of SAFA can benefit from the research on other tasks. Due to the huge difference between low-level and high-level vision tasks, using the paradigm of fine-tuning pre-trained models in low-level tasks is still challenging [40]. SAFA can utilize pre-trained models to improve its STVSR performance, especially using unsupervised pre-trained parameter weights (e.g., DINO [4]) (**d3**). For a fair comparison with the comparison methods, we initialize from scratch the parameters of the feature extractor in SAFA. We reported these results to show potential future improvements for inspiring future work.

(e) Choice of hyper-parameters. We set the number of channels $nc = 80$ in SAFA. Reducing model size could increase speed, but also decrease performance (**e2**). Choosing a larger or smaller peak learning rate degrades the performance (**e3,e4**).

5. Conclusion

We developed a powerful network for STVSR. Specifically, we proposed a Scale-Adaptive Flow Estimation (SAFE) block to perform scale selection for accurate motion modeling. We further introduced an iterative estimation scheme into SAFA to perform scale-adaptive feature propagation and fusion for efficient STVSR performance. Experimental results demonstrated that SAFA outperforms predecessors comprehensively. Extensive ablation studies confirm the efficacy of our approach. We envision SAFA as a solid foundation for future research. Further direction includes exploring the potential of SAFA on the real-world STVSR task and more advanced feature extractor.

Limitation. Limited by the studied topic, our work may not cover some scenarios. Firstly, following the setting of VideoINR [10], SAFA focuses on using two input images. For multi-frame inputs, SAFA could benefit from more implicit feature propagation [5, 35, 81]. SAFA is suitable for streaming systems as it does not depend on distant future frames. Secondly, our experiments are done with PSNR and SSIM as quantitative metrics to objectively measure the capability of the models [3]. SAFA can be readily changed to use the perceptual losses to cater the perception. Thirdly, SAFA may not tackle the diverse degradation that exists in real-world videos currently. To this end, it is necessary to train SAFA using additional training data and specialized training pipelines designed for real-world STVSR [89].

Acknowledgements. Jun Xu is partially supported by the National Natural Science Foundation of China (No. 62002176 and 62176068), and the Open Research Fund (No. B10120210117-OF03) from the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen.

References

- [1] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 6
- [2] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 3, 5
- [3] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 8
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 8
- [5] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 6, 7, 8
- [6] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Understanding deformable alignment in video super-resolution. In *AAAI Conference on Artificial Intelligence*, 2021. 2
- [7] Kelvin CK Chan, Shangchen Zhou, Xiangyu Xu, and Chen Change Loy. Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [8] Qiang Chen, Yingming Wang, Tong Yang, Xiangyu Zhang, Jian Cheng, and Jian Sun. You only look one-level feature. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 4
- [9] Shuhong Chen and Matthias Zwicker. Improving the perceptual quality of 2d animation interpolation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 3
- [10] Zeyuan Chen, Yinbo Chen, Jingwen Liu, Xingqian Xu, Vidit Goel, Zhangyang Wang, Humphrey Shi, and Xiaolong Wang. Videoinr: Learning video implicit neural representation for continuous space-time super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 5, 6, 7, 8
- [11] Xianhang Cheng and Zhenzhong Chen. Multiple video frame interpolation via enhanced deformable separable convolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(10):7029–7045, 2022. 3
- [12] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 3
- [13] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. Flownet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 2
- [14] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *The Journal of Machine Learning Research*, 20(1):1997–2017, 2019. 2
- [15] Dario Fuoli, Shuhang Gu, and Radu Timofte. Efficient video super-resolution through recurrent latent space propagation. In *2019 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. IEEE, 2019. 2
- [16] Zhicheng Geng, Luming Liang, Tianyu Ding, and Ilya Zharkov. Rstt: Real-time spatial temporal transformer for space-time video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 6
- [17] Yunhui Han, Kunming Luo, Ao Luo, Jiangyu Liu, Haoqiang Fan, Guiming Luo, and Shuaicheng Liu. Realflow: Embased realistic optical flow dataset generation from videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [18] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Space-time-aware multi-resolution video enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 4
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4, 7, 8
- [21] Mengshun Hu, Kui Jiang, Liang Liao, Jing Xiao, Junjun Jiang, and Zheng Wang. Spatial-temporal space hand-in-hand: Spatial-temporal video super-resolution via cycle-projected mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [22] Mengshun Hu, Kui Jiang, Zhixiang Nie, and Zheng Wang. You only align once: Bidirectional interaction for spatial-temporal video super-resolution. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. 3
- [23] Xiaotao Hu, Zhewei Huang, Ailin Huang, Jun Xu, and Shuchang Zhou. A dynamic multi-scale voxel flow network for video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 5
- [24] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [25] Yinlin Hu, Rui Song, and Yunsong Li. Efficient coarse-to-fine patchmatch for large displacement optical flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4

- [26] Yan Huang, Wei Wang, and Liang Wang. Video super-resolution via bidirectional recurrent convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(4):1015–1028, 2017. 2
- [27] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [28] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Real-time intermediate flow estimation for video frame interpolation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3, 4, 5
- [29] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Lite-flownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [30] Junhwa Hur and Stefan Roth. Iterative residual refinement for joint optical flow and occlusion estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5
- [31] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 5
- [32] Takashi Isobe, Songjiang Li, Xu Jia, Shanxin Yuan, Gregory Slabaugh, Chunjing Xu, Ya-Li Li, Shengjin Wang, and Qi Tian. Video super-resolution with temporal group attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [33] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 3, 4
- [34] Shihao Jiang, Yao Lu, Hongdong Li, and Richard Hartley. Learning optical flow from a few matches. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [35] Tarun Kalluri, Deepak Pathak, Manmohan Chandraker, and Du Tran. Flavr: Flow-agnostic video representations for fast frame interpolation. *arXiv preprint arXiv:2012.08512*, 2020. 8
- [36] Jaeyeon Kang, Younghyun Jo, Seoung Wug Oh, Peter Vajda, and Seon Joo Kim. Deep space-time video upsampling networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 3
- [37] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Fisr: Deep joint frame interpolation and super-resolution with a multi-scale temporal loss. In *AAAI Conference on Artificial Intelligence*, 2020. 3
- [38] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [39] Lingtong Kong, Boyuan Jiang, Donghao Luo, Wenqing Chu, Xiaoming Huang, Ying Tai, Chengjie Wang, and Jie Yang. Ifrnet: Intermediate feature refine network for efficient frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 5, 6
- [40] Klemen Kotar, Gabriel Ilharco, Ludwig Schmidt, Kiana Ehsani, and Roozbeh Mottaghi. Contrasting contrastive self-supervised representation learning pipelines. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 8
- [41] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 4
- [42] Hyeongmin Lee, Taeoh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoung Lee. Adacof: Adaptive collaboration of flows for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [43] Jingyun Liang, Jie Zhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 4
- [44] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 4
- [45] Ce Liu and Deqing Sun. A bayesian approach to adaptive video super resolution. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2011. 6
- [46] Ce Liu and Deqing Sun. On bayesian adaptive video super resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(2):346–360, 2013. 5
- [47] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 8
- [48] Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, and Jason Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 7
- [49] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 3, 4
- [50] Liying Lu, Ruizheng Wu, Huaijia Lin, Jiangbo Lu, and Jiaya Jia. Video frame interpolation with transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [51] Uma Mudenagudi, Subhashis Banerjee, and Prem Kumar Kalra. Space-time super-resolution using graph-cut optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(5):995–1008, 2010. 3
- [52] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene

- deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6, 7, 8
- [53] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 4
- [54] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive convolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [55] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 3
- [56] Junheum Park, Chul Lee, and Chang-Su Kim. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [57] Junheum Park, Chul Lee, and Chang-Su Kim. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 3
- [58] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [59] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [60] Fitsum Reda, Janne Kontkanen, Eric Tabellion, Deqing Sun, Caroline Pantofaru, and Brian Curless. Film: Frame interpolation for large motion. In *Proceedings of the European conference on computer vision (ECCV)*, 2022. 3
- [61] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention (MICCAI)*, 2015. 3
- [62] Laura Sevilla-Lara, Yiyi Liao, Fatma Güney, Varun Jampani, Andreas Geiger, and Michael J Black. On the integration of optical flow and action recognition. In *Pattern Recognition: 40th German Conference, GCPR 2018, Stuttgart, Germany, October 9-12, 2018, Proceedings 40*, pages 281–297. Springer, 2019. 1
- [63] Oded Shahar, Alon Faktor, and Michal Irani. *Space-time super-resolution from a single video*. IEEE, 2011. 3
- [64] Eli Shechtman, Yaron Caspi, and Michal Irani. Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 27(4):531–545, 2005. 3
- [65] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [66] Xingjian Shi, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 2
- [67] Xiaoyu Shi, Zhaoyang Huang, Weikang Bian, Dasong Li, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Videoflow: Exploiting temporal cues for multi-frame optical flow estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023. 2
- [68] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. Xvfi: extreme video frame interpolation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2, 4, 6, 8
- [69] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 4
- [70] Li Siyao, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris Metaxas, Chen Change Loy, and Ziwei Liu. Deep animation video interpolation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [71] Sanghyun Son, Suyoung Lee, Seungjun Nah, Radu Timofte, and Kyoung Mu Lee. Ntire 2021 challenge on video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2021. 6
- [72] Shuo Chen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 5, 6, 7, 8
- [73] Xiuchao Sui, Shaohua Li, Xue Geng, Yan Wu, Xinxing Xu, Yong Liu, Rick Goh, and Hongyuan Zhu. Craft: Cross-attentional flow transformer for robust optical flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [74] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 4, 5
- [75] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the IEEE International Conference on Machine Learning (ICML)*, 2019. 4
- [76] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 4, 5, 7
- [77] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally-deformable alignment network for video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [78] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced

- deformable convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019. 1, 2, 6
- [79] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [80] Yue Wu, Qiang Wen, and Qifeng Chen. Optimizing video prediction via video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [81] Xiaoyu Xiang, Yapeng Tian, Yulun Zhang, Yun Fu, Jan P. Allebach, and Chenliang Xu. Zooming slow-mo: Fast and accurate one-stage space-time video super-resolution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 3, 5, 6, 7, 8
- [82] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [83] Gang Xu, Jun Xu, Zhen Li, Liang Wang, Xing Sun, and Ming-Ming Cheng. Temporal modulation network for controllable space-time video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3, 5, 6
- [84] Haoifei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [85] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1
- [86] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. In *International Journal of Computer Vision (IJCV)*, 2019. 2, 6
- [87] Guozhen Zhang, Yuhan Zhu, Haonan Wang, Youxin Chen, Gangshan Wu, and Limin Wang. Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [88] Haoxian Zhang, Yang Zhao, and Ronggang Wang. A flexible recurrent residual pyramid network for video frame interpolation. In *Proceedings of the European conference on computer vision (ECCV)*, 2020. 3
- [89] Kai Zhang, Shuhang Gu, and Radu Timofte. Ntire 2020 challenge on perceptual extreme super-resolution: Methods and results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020. 8
- [90] Yuantong Zhang, Huairui Wang, Han Zhu, and Zhenzhong Chen. Optical flow reusing for high-efficiency space-time video super resolution. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 2022. 3
- [91] Chengcheng Zhou, Zongqing Lu, Linge Li, Qiangyu Yan, and Jing-Hao Xue. How video super-resolution and frame interpolation mutually benefit. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. 3
- [92] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016. 3, 5
- [93] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 4