Temporal Modulation Network for Controllable Space-Time Video Super-Resolution

Gang Xu[1]   Jun Xu[2*]   Zhen Li[1]   Liang Wang[3]   Xing Sun[4]   Ming-Ming Cheng[1]

[1] College of Computer Science, Nankai University, Tianjin, China
[2] School of Statistics and Data Science, Nankai University, Tianjin, China
[3] National Lab of Pattern Recognition, Institute of Automation, CAS, Beijing, China
[4] Youtu Lab., Tencent, Shanghai, China

Abstract

Space-time video super-resolution (STVSR) aims to increase the spatial and temporal resolutions of low-resolution and low-frame-rate videos. Recently, deformable convolution based methods have achieved promising STVSR performance, but they could only infer the intermediate frame pre-defined in the training stage. Besides, these methods undervalued the short-term motion cues among adjacent frames. In this paper, we propose a Temporal Modulation Network (TMNet) to interpolate arbitrary intermediate frame(s) with accurate high-resolution reconstruction. Specifically, we propose a Temporal Modulation Block (TMB) to modulate deformable convolution kernels for controllable feature interpolation. To well exploit the temporal information, we propose a Locally-temporal Feature Comparison (LFC) module, along with the Bi-directional Deformable ConvLSTM, to extract short-term and long-term motion cues in videos. Experiments on three benchmark datasets demonstrate that our TMNet outperforms previous STVSR methods. The code is available at `https://github.com/CS-GangXu/TMNet`.

1. Introduction

Nowadays, flat-panel displays using liquid-crystal display (LCD) or light-emitting diode (LED) technologies can broadcast Ultra High Definition Television (UHD TV) videos with 4K (3840×2160) or 8K (7680×4320) full-color pixels, at the frame rate of 120 frames per second (FPS) or 240 FPS [39]. However, currently available videos are commonly in Full High Definition (FHD) with a resolution of 2K (1920×1080) at 30 FPS [45]. To broadcast FHD videos on UHD TVs, it is necessary to increase their
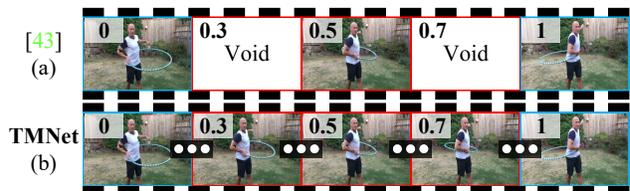
Figure 1: **Flexible STVSR performance by our TMNet**. Given input frames at moments **0** (begin) and **1** (end), [45] could only interpolate pre-defined intermediate frame at moment **0.5** (a), while our TMNet can generate intermediate frames at arbitrary moments (e.g., **0.3**, **0.5**, **0.7**) (b).

space-time resolutions comfortably with the broadcasting standard of UHD TVs. Although it is possible to increase the spatial resolution of videos frame-by-frame via single image super-resolution methods [4, 22], the perceptual quality of the enhanced videos would be deteriorated by temporal distortion [17]. To this end, the space-time video super-resolution (STVSR) methods [31, 45] are developed to simultaneously increase the spatial and temporal resolutions of low-frame-rate and low-resolution videos.

Previous model-based STVSR methods [30–32] rely heavily on precise spatial and temporal registration [38], and would produce inferior reconstruction results when the registration is inaccurate. Besides, they usually require huge computational costs on solving complex optimization problems, resulting in low inference efficiency [21, 25]. Later, deep convolutional neural networks [12, 13, 34, 44] have been widely employed in video restoration tasks such as video super-resolution (VSR) [3, 37], video frame interpretation (VFI) [1, 27, 46], and the more challenging STVSR [17, 45]. A straightforward solution for STVSR is to perform VFI and VSR successively on low-resolution and low-frame-rate videos, to increase their spatial resolutions and frame rates [45]. However, these two-stage methods ignore the inherent correlation between temporal and spatial dimensions. That is, the videos with high-resolution frames contain richer details on moving object(s) and back-

ground, while those in high-frame-rate provide finer pixel alignment between adjacent frames [8]. Therefore, these two-stage STVSR methods would suffer from the temporal inconsistency problem [45] and produce artifacts, e.g., "the attentional blink phenomenon" [36] on STVSR.

To well exploit the correlation between the temporal and spatial dimensions in videos, several one-stage STVSR methods [8, 17, 45] have been proposed to simultaneously perform VFI and VSR reconstruction on low-frame-rate and low-resolution videos. The work of STARnet [8] estimates the motion cues with an additional optical flow branch [5], and performs feature warping of two adjacent frames to interpolate the intermediate frame. But this flow-based method [8] needs to learn an extra branch for optical flow estimation, which consumes expensive costs on computation and memory. To alleviate this problem, Xiang *et al.* [45] employed the deformable convolution backbone [41], and directly performed STVSR on the feature space. Though with promising performance, current STVSR networks could only generate the intermediate frames pre-defined in the network architecture, and thus are restricted to highly-controlled application scenarios with fixed frame-rate videos. However, in many commercial scenarios, such as sports events, it is very common for the user to flexibly adjust the intermediate video frames for better visualization. Thus, it is necessary to develop controllable STVSR methods for smooth motion synthesizing.

To fulfill the versatile requirements of broadcasting scenarios, in this paper, we propose a Temporal Modulation Network (TMNet) to interpolate an arbitrary number of intermediate frames for STVSR, as shown in Figure 1. But current deformable convolution based methods [45] could only generate pre-defined intermediate frame(s). To tackle this problem, we propose a Temporal Modulation Block (TMB) to incorporate motion cues into the feature interpolation of intermediate frames. Specifically, we first estimate the motion between two adjacent frames under the deformable convolution framework [41], and learn controllable interpolation at an arbitrary moment defined by a temporal parameter. In addition, we also propose a Locally-temporal Feature Comparison module to fuse multi-frame features for effective spatial alignment and feature warping, and a globally-temporal feature fusion to explore the long-term variations of the whole video. This two-stage temporal feature fusion scheme accurately interpolates the intermediate frames for STVSR. Extensive experiments on three benchmarks [23, 35, 46] demonstrate that our TMNet is able to interpolate an arbitrary number of intermediate frames, and achieves state-of-the-art performance on STVSR.

The contribution of this work are three-fold:

- **We propose a Temporal Modulation Network (TMNet) to perform controllable interpolation of arbitrary frame-rates** for flexible STVSR performance.

This is achieved by our Temporal Modulation Block under the deformable convolution framework.

- **We present a two-stage temporal feature fusion scheme for effective STVSR.** Specifically, we propose a locally-temporal feature comparison module to exploit the short-term motion cues of adjacent frames, and perform globally-temporal feature fusion by exploring the long-term variations over the whole video.

- Experiments on three benchmarks show that **our TMNet is able to perform controllable frame interpolation at arbitrary frame-rate**, and outperforms state-of-the-art STVSR methods.

2. Related Work

Video frame interpolation (VFI) aims to synthesize new intermediate frames between adjacent frames [2, 15, 20]. Early VFI methods mainly resort to optical flow techniques for motion estimation [2, 15, 27]. Jiang *et al.* [15] modeled motion interpretation for arbitrary frame-rate VFI. Niklaus *et al.* [26] warped the input frames with contextual information, and interpolated context-aware intermediate frames. Bao *et al.* employed motion estimation and compensation for VFI in [2], and obtained improved performance by further exploring the depth information [1]. Niklaus *et al.* [27] tackled the conflicts of mapping multiple pixels to the same location in VFI by softmax splatting. However, these optical flow based methods need huge computational costs on motion estimation. Therefore, recently researchers exploited to learn spatially-adaptive convolution kernels [28] or deformable ones for VFI [20].

Video super-resolution (VSR) is the task of increasing the spatial resolutions of low-resolution (LR) videos [16, 37, 41]. Existing VSR methods [16, 37, 41] mainly aggregate spatial information of multiple frames for high-resolution (HR) reconstruction, with the help of optical flow techniques [5]. Jo *et al.* [16] generated dynamic upsampling filters to enhance the LR frames with residual learning [12]. Wang *et al.* [41] proposed the Pyramid, Cascading and Deformable (PCD) module to perform frame alignment, and then fused multiple frames into a single one by spatial and temporal attention. Haris *et al.* [7] designed an iterative refinement framework by integrating the spatial and temporal contexts of multiple frames. Tian *et al.* [37] utilized the learned sampling offsets of deformable convolution kernels to align the supporting frames with the reference ones, which are both used to reconstruct the HR frames.

Space-time video super-resolution (STVSR) aims to increase the spatial and temporal dimensions of the low-frame-rate and low-resolution videos [8, 17, 45]. Shechtman *et al.* [31] tackled the STVSR problem by employing a directional space-time smoothness regularization on the HR video reconstruction problem. Mudenagudi *et al.* [25] for-
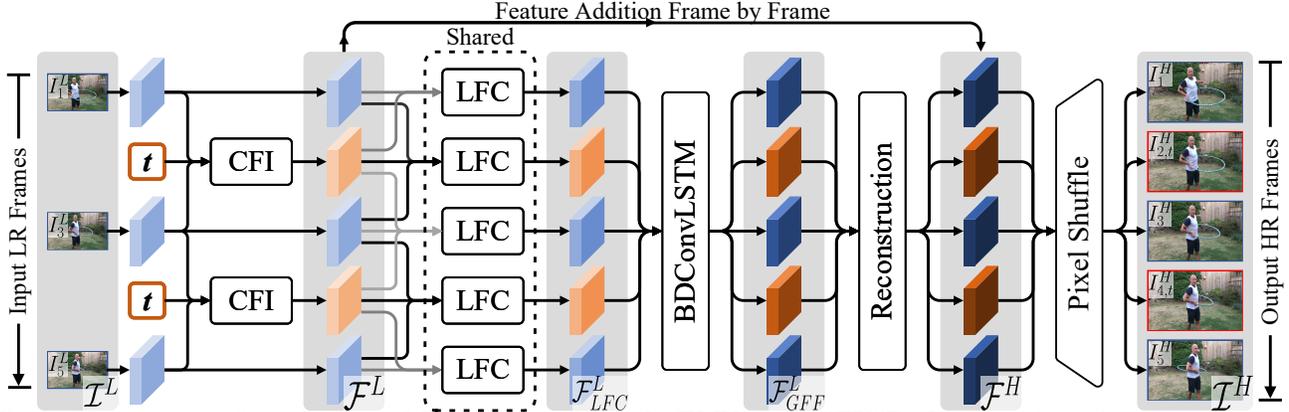
Figure 2: **Overview of our Temporal Modulation Network (TMNet) for STVSR**. Given the input low-frame-rate and low-resolution video $\mathcal{I}^L$, we first extract initial features and perform Controllable Feature Interpolation (CFI, implemented by our Temporal Modulation Block) for the intermediate frame at an arbitrary moment $t \in (0, 1)$. Then, we feed the obtained feature maps $\mathcal{F}^L$ into a two-stage temporal feature fusion scheme. For short-term motion consistency, the feature maps $\mathcal{F}^L$ are refined to $\mathcal{F}^L_{LFC}$ by our Locally-temporal Feature Comparison (LFC) module. To exploit long-term motion cues, the feature maps $\mathcal{F}^L_{LFC}$ are further improved to $\mathcal{F}^L_{GFF}$ by globally-temporal feature fusion (GFF), implemented by Bi-directional Deformable ConvLSTM (BDConvLSTM) [45]. Finally, we employ 40 residual blocks to reconstruct high-resolution feature maps $\mathcal{F}^H$, and two Pixel-Shuffle layers to output the high-frame-rate and high-resolution video $\mathcal{I}^H$.

mulated their STVSR method under the Markov Random Field framework [6]. STARnet [8] leveraged inherent motion relationship between spatial and temporal dimensions with an extra optical flow branch [5], and perform feature warping of two adjacent frames to interpolate the intermediate frame. Xiang *et al*. [45] developed a unified framework to interpolate the multi-frame features via PCD alignment modules [41], the intermediate features by bidirectional deformable ConvLSTM [34], and finally performed STVSR by multi-frame feature fusion. In this work, our goal is to develop a temporally controllable network for powerful and flexible STVSR. Though built upon [45], our TMNet arrives at better performance on benchmark datasets, owing to the proposed locally temporal feature comparison module.

Modulation networks. Recently, researchers proposed to control the restoration intensity of the main network by additional modulation branches [9, 10, 40, 42]. These modulation networks are trained to trade-off the restoration quality and flexibility, which are controlled by hyper-parameters. He *et al*. [9] put feature modulation filters after each convolution layer to modulate the output according to user's preference. Later, He *et al*. [10] expanded this design to multiple dimensions, and modulated the output according to the levels of multiple degradation types. Wang *et al*. [40] learned the features from tuning blocks and residual ones with different objectives, to control the trade-off between noise reduction and detail preservation. In this work, we consider the modulation on temporal dimension, instead of on restoration intensity as in these modulation networks. As far as we know, our work is among the first to implement temporal modulation in the STVSR problem. As will be shown in §4, our TMNet can explore the potential of tem-

poral modulation for controllable STVSR.

3. Proposed Method

In this section, we first overview our Temporal Modulation Network (TMNet) for STVSR in §3.1. Then, we introduce our Temporal Modulation Block for controllable feature interpolation in §3.2. We present temporal feature fusion in §3.3, and high-resolution reconstruction in §3.4. Finally, the training details are given in §3.5.

3.1. Network Overview

As illustrated in Fig. 2, our TMNet consists of three seamless stages: controllable feature interpolation, temporal feature fusion, and high-resolution reconstruction.

Controllable feature interpolation. Given a sequence of low-frame-rate and low-resolution video $\mathcal{I}^L = \{I^L_{2i-1}\}^n_{i=1}$, our TMNet firstly extracts the corresponding initial feature maps $\{F^L_{2i-1}\}^n_{i=1}$ through 5 residual blocks. To perform temporally controllable feature interpolation, we propose a Temporal Modulation Block (TMB) to modulate the deformable convolution kernels with a temporal hyperparameter t. Here, $t \in (0, 1)$ indicates the (arbitrary) moment at which we plan to interpolate a feature map $F^L_{2i,t}$ from the feature maps F^L_{2i-1} and F^L_{2i+1} of two adjacent frames I^L_{2i-1} and I^L_{2i+1}, respectively. Finally, we obtain a feature sequence $\mathcal{F}^L = \{F^L_1, F^L_{2,t}, F^L_3, ..., F^L_{2n-2,t}, F^L_{2n-1}\}$ of high-frame-rate and low-resolution video frames.

Temporal feature fusion. The extracted (or interpolated) feature maps in $\mathcal{F}^L$ are often of low-quality, since they are extracted from individual LR frames (or interpolated by the initial feature maps of adjacent LR frames). Thus, we pro-
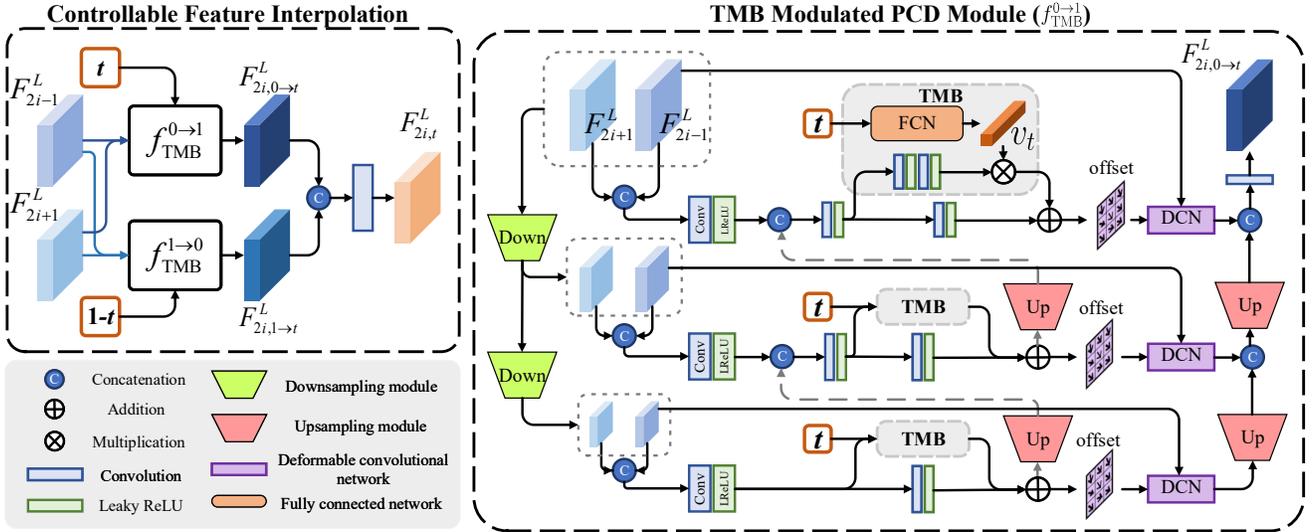
Figure 3: **Proposed Temporal Modulation Block (TMB)** modulated Pyramid, Cascading and Deformable (PCD) module [41] for controllable feature interpolation. $f_{\mathrm{TMB}}^{0\rightarrow1}$ ($f_{\mathrm{TMB}}^{1\rightarrow0}$) is the PCD module modulated by our TMB block to model the forward (backward) motion. Our TMB modulates all the three levels of the PCD module, by transforming the temporal hyper-parameter t into a modulation vector v_t via a fully connected network (FCN) consisted of three convolutional layers.

pose a Locally-temporal Feature Comparison (LFC) module, to refine every feature map in $\mathcal{F}^L$ with the help of the feature maps of adjacent frames. After the local feature refinement, we further improve the feature maps in $\mathcal{F}^L$ by performing globally-temporal feature fusion (GFF). This is implemented by employing a Bi-directional Deformable ConvLSTM (BDConvLSTM) network [45], to consecutively aggregate the useful information from individual feature maps along the temporal direction. Both the LFC and GFF fusion modules well exploit the intra-correlation between spatial and temporal dimensions to improve the quality of feature maps in $\mathcal{F}^L$. In the end, we obtain the sequence of improved feature maps $\mathcal{F}_{GFF}^L$.

High-resolution reconstruction. Here, we feed the sequence of feature maps $\mathcal{F}_{GFF}^L$ into 40 residual blocks to improve their quality along the spatial dimension. Next, we increase the spatial resolution of these improved feature maps via the widely used Pixel-Shuffle layers [33], and output the final high-frame-rate and high-resolution video sequence $\mathcal{I}^H = \{\boldsymbol{I}_1^H, \boldsymbol{I}_{2,t}^H, ..., \boldsymbol{I}_{2n-2,t}^H, \boldsymbol{I}_{2n-1}^H\}$.

3.2. Controllable Feature Interpolation

Given a sequence of low-frame-rate and low-resolution video frames $\mathcal{I}^L = \{\boldsymbol{I}_{2i-1}^L\}_{i=1}^n$, we first extract the corresponding initial features $\mathcal{F}^L = \{\boldsymbol{F}_{2i-1}^L\}_{i=1}^n$ via five residual blocks. Each residual block contains a sequence of "Conv-ReLU-Conv" operations with a skip connection. For any two adjacent frames $\boldsymbol{I}_{2i-1}^L$ and $\boldsymbol{I}_{2i+1}^L$ ($i \in \{1, ..., n-1\}$), our goal here is to interpolate the feature of intermediate frame at an arbitrary moment $t \in (0, 1)$. To this end, we need to estimate the motion cues from $\boldsymbol{I}_{2i-1}^L$ to the intermediate frame (forward) and that from $\boldsymbol{I}_{2i+1}^L$ to the intermedi-

ate frame (backward). Previous STVSR methods [41, 45] utilized the Pyramid, Cascading and Deformable (PCD) module to estimate the offset between $\boldsymbol{F}_{2i-1}^L$ and $\boldsymbol{F}_{2i+1}^L$ as the motion cues, to align and interpolate the features of the intermediate frame under the deformable convolutional framework [48]. However, the vanilla PCD module could only estimate the motion to a predefined moment, which is fixed in both training and inference stages.

To overcome this limitation, we propose a Temporal Modulation Block (TMB) to modulate the learned offset between $\boldsymbol{F}_{2i-1}^L$ and $\boldsymbol{F}_{2i+1}^L$. The modulation is controlled by a hyper-parameter $t \in (0, 1)$, indicating an arbitrary moment that we plan to interpolate a new frame. This enables our TMNet to control the feature interpolation process upon the initial feature maps $\boldsymbol{F}_{2i-1}^L$ and $\boldsymbol{F}_{2i+1}^L$ of two adjacent frames $\boldsymbol{I}_{2i-1}^L$ and $\boldsymbol{I}_{2i+1}^L$ in the input video. The PCD module modulated by our TMB block can estimate the forward and backward motions and interpolate the feature map $\boldsymbol{F}_{2i,t}^L$ of a new frame at the arbitrary moment $t \in (0, 1)$.

Denote by $f_{\mathrm{TMB}}^{0\rightarrow1}$ and $f_{\mathrm{TMB}}^{1\rightarrow0}$ the PCD modules modulated by our TMB block, to model the forward and backward motions, respectively. Here, we perform modulated feature interpolation from the forward and backward directions:

$$
\begin{aligned}
\boldsymbol{F}_{2i,0\rightarrow t}^L &= f_{\mathrm{TMB}}^{0\rightarrow1}(\boldsymbol{F}_{2i-1}^L, \boldsymbol{F}_{2i+1}^L, t), \\
\boldsymbol{F}_{2i,1\rightarrow t}^L &= f_{\mathrm{TMB}}^{1\rightarrow0}(\boldsymbol{F}_{2i-1}^L, \boldsymbol{F}_{2i+1}^L, 1-t),
\end{aligned} \tag{1}
$$

where $\boldsymbol{F}_{2i,0\rightarrow t}^L$ and $\boldsymbol{F}_{2i,1\rightarrow t}^L$ are the interpolated features aligned from the feature maps $\boldsymbol{F}_{2i-1}^L$ and $\boldsymbol{F}_{2i+1}^L$ of the adjacent frames. Note that the two TMB-modulated PCD modules share the same network structure but have different weights. Here we only take the $f_{\mathrm{TMB}}^{0\rightarrow1}$ as an example to explain how the PCD modules modulated by our TMB work

on modeling the forward motion. The PCD module $f_{\text{TMB}}^{1\to0}$ modeling the backward motion can be similarly explained.

As shown in Fig. 3, the PCD module has three levels to estimate the motion in different scales. To realize flexible modulation on the temporal dimension, we embed our TMB block into each level of the vanilla PCD module independently, to modulate the offset before the deformable convolutional network (DCN). The benefits of adding our TMB block to all three levels of the PCD module will be verified in §4. To adaptively modulate the offset by our TMB at different levels of PCD, we use three convolutional layers to map the temporal hyper-parameter t onto a modulation vector $\boldsymbol{v}_t$ of size $1 \times 1 \times 64$. To well exploit the motion cues for precise modulation, we feed the features in each vanilla PCD level into two convolutional layers to enlarge their receptive fields. Then, the generated feature is multiplied with the modulation vector $\boldsymbol{v}_t$ along the channel dimension, to produce the TMB-modulated features. For robustness, we add the TMB-modulated features with the corresponding pre-modulated features before DCN.

Once obtaining the modulated feature maps $\boldsymbol{F}_{2i,0\to t}^{L}$ and $\boldsymbol{F}_{2i,1\to t}^{L}$, we interpolate the intermediate feature $\boldsymbol{F}_{2i,t}^{L}$ via channel-wise concatenation "$[\cdot, \cdot]$" followed by a 1×1 convolution layer $f_{1\times1}$ as:

$$\boldsymbol{F}_{2i,t}^{L} = f_{1\times1}([\boldsymbol{F}_{2i,0\to t}^{L}, \boldsymbol{F}_{2i,1\to t}^{L}]). \tag{2}$$

Now, we obtain the features of the interpolated sequence $\mathcal{F}^{L} = \{\boldsymbol{F}_{1}^{L}, \boldsymbol{F}_{2,t}^{L}, \boldsymbol{F}_{3}^{L}, ..., \boldsymbol{F}_{2n-2,t}^{L}, \boldsymbol{F}_{2n-1}^{L}\}$ for the high-frame-rate and low-resolution video. Next, we perform feature fusion along the temporal dimension.

3.3. Temporal Feature Fusion

Here, the initial features are extracted (or interpolated) from individual (or adjacent) frames. There is considerable leeway to improve their quality. But we also feed the initial features into the Pixel-Shuffle part of our TMNet.

Locally-temporal feature comparison. It is essential to maintain short-term temporal consistency for each current frame. For this purpose, we propose a Locally-temporal Feature Comparison (LFC) module to exploit the complementary information (e.g., motion cues) from adjacent frames. As illustrated in Fig. 4, to refine the feature map $\boldsymbol{F}_{2i,t}^{L}$ of current frame from adjacent feature maps $\boldsymbol{F}_{2i-1}^{L}$ and $\boldsymbol{F}_{2i+1}^{L}$, we concatenate current frame ($\boldsymbol{F}_{2i,t}^{L}$) and adjacent frames ($\boldsymbol{F}_{2i-1}^{L}$, $\boldsymbol{F}_{2i+1}^{L}$), and employ two convolutional layers to learn the offset in the deformable convolutional framework [48]. Note that we learn two offsets to describe the motion cues in the forward (from $\boldsymbol{I}_{2i-1}^{L}$ to current frame) and the backward (from $\boldsymbol{I}_{2i+1}^{L}$ to current frame) directions. Then, the learned offset from forward (or backward) direction is used to align the feature map $\boldsymbol{F}_{2i-1}^{L}$ of previous (or $\boldsymbol{F}_{2i+1}^{L}$ of next) frame with that of the current frame, via one deformable convolutional layer. After the alignment,
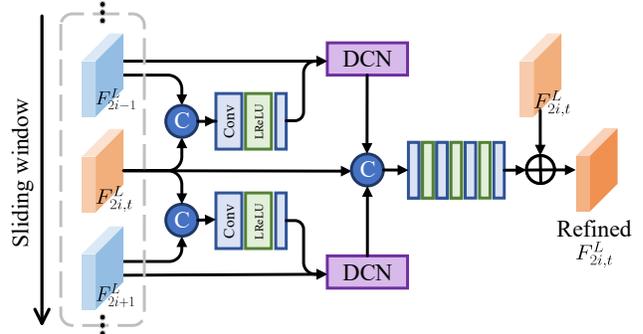

Figure 4: **Proposed Locally-temporal Feature Comparison** (LFC) module refines the interpolated feature $\boldsymbol{F}_{2i,t}^{L}$ by exploiting short-term motion cues among adjacent frames.

we concatenate the aligned feature maps of two adjacent frames with that of the current frame, and perform feature comparison via four 1×1 convolutional layers and an addition operation. For the first (or last) frame, the previous (or next) adjacent frame is just itself. Now we get a refined feature sequence $\mathcal{F}_{LFC}^{L}$.

Globally-temporal feature fusion. The feature sequence refined by our LFC module is able to maintain short-term consistency in the interpolated video. But it would fail on large or fast motions, since LFC lacks the capability of modeling the motions over the whole video. To tackle this problem, we propose to exploit the long-term information in videos by globally-temporal feature fusion. Inspired by [45], we feed the feature sequence $\mathcal{F}_{LFC}^{L}$ generated by our LFC into the BDConvLSTM network, and obtain the features $\mathcal{F}_{GFF}^{L}$ in long-term temporal consistency.

As will be illustrated in the experimental section, our short-term LFC module and the long-term BDConvLSTM indeed boost the performance of our TMNet on STVSR.

3.4. High-Resolution Reconstruction

Until now, the intra-correlation of temporal and spatial dimensions is well explored to obtain the high-quality feature sequence $\mathcal{F}_{GFF}^{L}$ of the whole video. Then, we perform spatial refinement for the feature maps via 40 residual blocks, and get the refined feature maps $\mathcal{F}^{H}$. Then we add the features $\mathcal{F}^{H}$ with the corresponding initial feature maps in $\mathcal{F}^{L}$, and obtain the reconstructed feature maps $\mathcal{F}_{final}^{H}$. Finally, we feed the reconstructed feature maps $\mathcal{F}_{final}^{H}$ into two Pixel-Shuffle layers, followed by a sequence of "Conv-LeakyReLU-Conv" operations, to output the reconstructed HR video frames $\mathcal{I}^{H} = \{\boldsymbol{I}_{1}^{H}, \boldsymbol{I}_{2,t}^{H}, ..., \boldsymbol{I}_{2n-2,t}^{H}, \boldsymbol{I}_{2n-1}^{H}\}$.

3.5. Training Details

Implementation details. We employ the Adam optimizer [18] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ to optimize our TMNet with the Charbonnier loss function [19], as suggested in [45]. The learning rate is initialized as 4×10^{-4}, and is decayed to 1×10^{-7} with a cosine annealing [24] for

Table 1: **Comparison of PSNR, SSIM [43], speed (in `fps`), and parameters (in `million`) by different STVSR methods** on Vid4 [35], Vimeo-`Fast`, Vimeo-`Medium`, Vimeo-`Slow` [46]. "↑" means that larger is better. The speed is evaluated on Vid4 [35]. The best, second best and third best results are highlighted in red, blue and **bold**, respectively.

Method	Vid4 [35]		Vimeo-Fast		Vimeo-Medium		Vimeo-Slow		Speed	Parameters
VFI+(V)SR / STVSR	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	fps↑	million↓
SuperSloMo [15] + Bicubic	22.84	0.5772	31.88	0.8793	29.94	0.8477	28.37	0.8102	-	**19.8**
SuperSloMo [15] + RCAN [47]	23.80	0.6397	34.52	0.9076	32.50	0.8884	30.69	0.8624	2.49	19.8+16.0
SuperSloMo [15] + RBPN [7]	23.76	0.6362	34.73	0.9108	32.79	0.8930	30.48	0.8584	2.06	19.8+12.7
SuperSloMo [15] + EDVR [41]	24.40	0.6706	35.05	0.9136	33.85	0.8967	30.99	0.8673	6.85	19.8+20.7
SepConv [28] + Bicubic	23.51	0.6273	32.27	0.8890	30.61	0.8633	29.04	0.8290	-	21.7
SepConv [28] + RCAN [47]	24.92	0.7236	34.97	0.9195	33.59	0.9125	32.13	0.8967	2.42	21.7+16.0
SepConv [28] + RBPN [7]	26.08	0.7751	35.07	0.9238	34.09	0.9229	32.77	0.9090	2.01	21.7+12.7
SepConv [28] + EDVR [41]	25.93	0.7792	35.23	0.9252	34.22	0.9240	32.96	0.9112	6.36	21.7+20.7
DAIN [1] + Bicubic	23.55	0.6268	32.41	0.8910	30.67	0.8636	29.06	0.8289	-	24.0
DAIN [1] + RCAN [47]	25.03	0.7261	35.27	0.9242	33.82	0.9146	32.26	0.8974	2.23	24.0+16.0
DAIN [1] + RBPN [7]	25.96	0.7784	35.55	0.9300	34.45	0.9262	32.92	0.9097	1.88	24.0+12.7
DAIN [1] + EDVR [41]	**26.12**	0.7836	35.81	0.9323	34.66	0.9281	**33.11**	0.9119	5.20	24.0+20.7
STARnet [8]	26.06	0.8046	**36.19**	0.9368	34.86	0.9356	33.10	0.9164	**14.08**	111.61
Zooming Slow-Mo [45]	26.31	0.7976	36.81	0.9415	35.41	0.9361	33.36	0.9138	16.50	11.10
TMNet (Ours)	26.43	0.8016	37.04	0.9435	35.60	0.9380	33.51	0.9159	14.69	12.26

every 150,000 iterations. We initialize the parameters of our TMNet by Kaiming initialization [11] without pre-trained weights. The batch size is 24. Our TMNet, implemented in PyTorch [29] and Jittor [14], is trained in a total of 600,000 iterations on four RTX 2080Ti GPUs, which takes about 8.71 days (209.04 hours). For each input video clip, we randomly crop it into a sequence of downsampled patches of size 32×32. For data argumentation, we horizontal-flip each frame, and randomly rotate it with $90°$, $180°$, or $270°$.

Network training. When directly trained with the proposed TMB block, our TMNet suffers from clear performance drops on STVSR, as shown in our experiments. One possible reason is that our TMNet can not accurately estimate the motion cues to interpolate an intermediate frame at the arbitrary moment $t \in (0, 1)$ since our TMB does not know the moment before the training modulated feature. To resolve this problem, we propose to train our TMNet by a two-step strategy: `Step 1`, we train our main TMNet without the proposed TMB block; `Step 2`, we only train our TMB block while fixing the trained main network.

In `Step 1`, we train our TMNet on the Vimeo-90K dataset [46], which will be introduce in §4.1. This dataset consists of 7-frame video clips. For each clip, the 1-st, 3-rd, 5-th, and 7-th LR frames are input into our TMNet as low-frame-rate and low-resolution video. We set $t = 0.5$ to get rid of the TMB block from our TMNet, and learn to generate the 7-frame high-resolution and high-frame-rate video. This enables our TMNet to fairly compare with previous STVSR methods [8, 35, 41, 45]. For supervision, we calculate the loss function over the corresponding 7-frame HR video clip in the Vimeo-90K dataset [46].

In `Step 2`, we fix the learned weights of our main network, and only train our TMB block for temporal mod-

ulation. The training is performed on the Adobe240fps dataset [35], which is in high-frame-rate and suitable for training our TMB block. We also split it into groups of 7-frame video clips. For each clip, the 1-st and 7-th HR frames are downsampled as the inputs of our TMNet. We set the temporal hyper-parameter $t \in \{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\}$ to interpolate 5 intermediate frames. This step costs 35.26 minutes.

4. Experiments

4.1. Experimental Setup

Dataset. We use Vimeo-90K `septuplet` dataset [46] as the training set. It contains 91,701 video sequences, extracted from 39K video clips selected from Vimeo-90K. Each sequence contains 7 continuous frames of resolution 448×256. The Vid4 [23] and Vimeo-90K `test` set are used as evaluation datasets. As suggested in [45], we split the Vimeo-90K `septuplet test` set into three subsets of `Fast` motion, `Medium` motion, and `Slow` motion, which include 1225, 4977, and 1613 video clips, respectively. We also remove 5 video clips from the original `Medium` motion set and 3 clips from the `Slow` motion set, which contain only all-black backgrounds.

To make our TMNet feasible for controllable feature interpolation, we train our TMB block individually on the Adobe240fps dataset [35]. It has 133 videos (in 720P) taken with hand-held cameras, and is randomly split into the `train`, `val`, and `test` subsets with 100, 16, and 17 videos, respectively. For each video, we split it into groups of 7-frame video clips. We feed the 1-st and 7-th frames in each clip into our TMNet to generate 5 intermediate frames.

We downsample the HR frames to create the LR frames by bicubic interpolation, with a factor of 4.

Figure 5: **Qualitative and quantitative results of different methods on STVSR**. The test video clips are from the Adobe240fps [35] (1-st row), Vimeo-Fast [46] (2-nd row, left) and Vimeo-Slow [46] (2-nd row, right) datasets.

Evaluation metric. We employ the widely used Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) [43] to evaluate different methods on the STVSR task. The PSNR and SSIM metrics are calculated on the Y channel of the YCbCr color space, as favored by previous VSR [7, 41] and STVSR [45] methods.

4.2. Comparison to State-of-the-arts

Comparison methods. We compare our TMNet with state-of-the-art two-stage and one-stage STVSR methods. For the two-stage STVSR methods, we perform video frame interpolation (VFI) by SuperSloMo [15], DAIN [1] or SepConv [28], and perform video super-resolution (VSR) by Bicubic Interpolation (BI), RCAN [47], RBPN [7] or EDVR [41]. For one-stage STVSR methods, we compare our TMNet with the recently developed Zooming SlowMo [45] and STARnet [8]. To fairly compare with these competitors, we set $t = 0.5$ in our TMNet to generate the frame at the middle moment of any two adjacent frames. That is, the 1-st, 3-rd, 5-th, and 7-th LR frames of each clip in Vimeo-90K are fed into our TMNet to reconstruct the 7 HR frames. All these methods are trained on the Vimeo-90K septuplet dataset [46], evaluated on the Vimeo-90K test set [46] and the Vid4 [35] dataset.

Objective results. We list the quantitative comparison results in Table 1. As suggested in [45], we omit the baseline models with Bicubic Interpolation when comparing the speed. One can see that our TMNet outperforms the Zooming SlowMo [45] by 0.12dB, 0.23dB, 0.19dB, and 0.15dB on the Vid4, Vimeo-Fast, Vimeo-Medium, and Vimeo-Slow datasets in terms of PSNR. On SSIM [43], our TMNet achieves better results than the competitors in most cases, but is only slightly inferior to STARnet [8] on Vid4 [23] and Vimeo-Slow. However, our TMNet needs only one-ninth of the parameters in STARnet. On speed, one-stage methods [8, 45] run much faster than two-stage ones [1, 7, 15, 28, 41]. Our TMNet runs at 14.69fps, and is only slower than Zooming Slow-Mo [45]. All these results

validate the effectiveness of our TMNet on STVSR.

Visualization. In the 1-st row of Figure 5, we present the 5 intermediate frames (Frame 1550 to Frame 1554) interpolated by our TMNet on the sequence "0056" from the Adobe240fps test set [35], given the Frame 1549 and Frame 1555 as inputs. It can be seen that our TMNet is able to perform flexible frame interpolation for STVSR. In the 2-nd row of Figure 5, we show the reconstructed frames by different STVSR methods from Vimeo-Fast and and Vimeo-Slow datasets [46] generated by the competing methods respectively. We observe that our TMNet, with the proposed LFC module, can restore more clearly the structures and textures than the competitors. For example, on the Clip "0001" in the sequence "0070" of Vimeo-Slow datasets, our TMNet reconstructs clearly the texture pattern on the bag. In summary, our TMNet demonstrates flexible and powerful STVSR ability quantitatively and qualitatively. More visual comparison on the Vid4 [23], Vimeo-90K test set [35], and Adobe240fps [46] datasets are provided in the *Supplementary File*, because of page limitation.

4.3. Ablation Study

Here, we conduct detailed examinations of our TMNet on STVSR. Specifically, we assess 1) the importance of our Temporal Modulation Block (TMB) for controllable feature interpolation; 2) different strategies that our TMB block modulates the PCD module; 3) how to design our TMB block; 4) how our Locally-temporal Feature Comparison (LFC) module contribute to the temporal feature fusion in our TMNet; 5) the combination of high-quality feature maps $\mathcal{F}^H$ and initial feature maps $\mathcal{F}^L$ for STVSR.

1. Does our TMB block contribute to controllable feature interpolation? To answer this question, we compare our TMNet with previous STVSR methods [8, 45] on generating intermediate frames from two adjacent frames. Due to limited space, we provide the comparison of visual results on Adobe240fps test set [35] in the *Supplementary File*. We observe that our TMNet with TMB block indeed

exhibits temporally controllable STVSR performance.

2. How different strategies that our TMB block modulate the PCD module influence our TMNet on STVSR? The PCD module [41] has a three-level pyramid structure: the 1-st level L1; the 2-nd level L2 is downsampled from the features in L1 by convolution filters at a stride of 2; similarly, the 3-rd level L3 is downsampled from L2 by a stride of 2. In our TMNet, the proposed TMB modulates all the three levels of the PCD module. But our TMB can also modulate only one level (L1, L2, or L3) of PCD, resulting three variants of our TMNet called TMB-L1, TMB-L2, and TMB-L3. These variants are trained on the Adobe240fps train set [35] and evaluate them on test set. As shown in Table 2, the three variants perform in descending order, indicating that the 1-st level of PCD is more important for temporal modulation. By modulating all three levels of PCD, our TMNet outperforms the three variants on STVSR, by better exploiting the motion cues of videos.

Table 2: **PSNR results on Adobe240fps `test` set by different strategies that our TMB modulates the PCD.**

Variant	TMB-L1	TMB-L2	TMB-L3	TMNet
PSNR (dB)	26.92	26.82	26.60	**26.95**

3. How to design our TMB block? The goal of our TMB is to transform the hyper-parameter t into a modulation vector v_t comfortable with the PCD module. A trivial design of our TMB is a linear convolutional layer. We call it TMB-Linear. We train our TMB and the TMB-Linear on Adobe240fps train set [35] and evaluate them on test set. The PSNR results are listd in Table 3, in which the TMB-Linear is 0.02dB lower than our TMB with three nonlinear convolutional layers. This shows that nonlinear transformation is only a little better than the linear one.

Table 3: **PSNR results on Adobe240fps `test` set by our TMB with linear or nonlinear design.**

Variant	TMB-Linear	TMB
PSNR (dB)	26.93	26.95

4. How important is the proposed LFC module to our TMNet? Our TMNet performs a two-stage temporal feature fusion: first local fusion via LFC and then global fusion via GFF. Thus, our TMNet can be called "LFC→GFF". Inverting the order, i.e., GFF→LFC, makes our TMNet collapse during the training. The main reason is that performing GFF before LFC brings noisy long-term information, confusing the learning of deformable convolution in LFC. Thus, we do not evaluate this variant. To study how our LFC contributes to the two-stage fusion in our TMNet, we remove LPC from our TMNet and call this variant "GFF". Besides, the features from LFC and GFF can be concatenated and fused by a convolutional layer, resulting in a variant "LPC+GFF". We train our TMNet and its variants on the Vimeo-90K septuplet dataset and evaluate them on

the Vid4 [23], Vimeo-Fast, Vimeo-Medium, and Vimeo-Slow datasets. The PSNR results are listed in Table 4. One can see that our TMNet (LPC→GFF) achieves the best results on all cases, and outperforms GFF by 0.07dB on Vid4, 0.17dB on Vimeo-Fast, 0.15dB on Vimeo-Medium, and 0.11dB on Vimeo-Slow. This indicates that our LFC module is essential to the success of our TMNet on STVSR, by exploiting short-term motion cues among adjacent frames.

Table 4: **Comparison of PSNR (dB) results by different variants of our TMNet** on STVSR datasets.

Variant	GFF	LFC+GFF	LFC→GFF
Vid4 [23]	26.36	26.35	**26.43**
Vimeo-Fast	36.87	36.90	**37.04**
Vimeo-Medium	35.45	35.47	**35.60**
Vimeo-Slow	33.40	33.43	**33.51**

5. The benefits of combining the high-quality feature maps $\mathcal{F}^H$ and the initial feature maps $\mathcal{F}^L$ for STVSR. In our TMNet, we combine the high-quality features $\mathcal{F}^H$ with the initial features $\mathcal{F}^L$ before the Pixel-Shuffle layers for final STVSR. Since the initial features $\mathcal{F}^L$ largely influence our LFC module, we remove them both from our TMNet and obtain a variant "Baseline". Then we add $\mathcal{F}^L$ to the "Baseline", and obtain a variant model "+$\mathcal{F}^L$". We train our TMNet and the two variants on the Vimeo-90K septuplet dataset [46], and evaluate them on Vimeo-90K test and Vid4 [23] datasets. As shown in Table 5, the variant "+$\mathcal{F}^L$" clearly exceeds the "Baseline". This validates that combining high-quality features $\mathcal{F}^H$ with initial ones $\mathcal{F}^L$ is helpful to our TMNet on STVSR.

Table 5: **Comparison of PSNR (dB) by our TMNet and its variants on different STVSR datasets**.

Variant	Baseline	+$\mathcal{F}^L$	TMNet
Vid4 [23]	26.33	26.36	**26.43**
Vimeo-Fast	36.75	36.87	**37.04**
Vimeo-Medium	35.35	35.45	**35.60**
Vimeo-Slow	33.28	33.40	**33.51**

5. Conclusion

In this work, we proposed a Temporal Modulation Network (TMNet) to flexibly interpolate intermediate frames for space-time video super-resolution (STVSR). Specifically, we introduced a Temporal Modulation Block to modulate the learning of the deformable convolution framework for controllable feature interpolation. To well exploit motion cues, we performed short-term and long-term temporal feature fusion consisting of our proposed Locally-temporal Feature Comparison (LFC) module and a Bi-directional Deformable ConvLSTM, respectively. Experiments on three benchmarks demonstrated the flexibility of our TMNet on interpolating intermediate frames, quantitative and qualitative advantages of our TMNet over previous methods, and effectiveness of our LFC module, for STVSR.

References

[1] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3703–3712, 2019. 1, 2, 6, 7

[2] Wenbo Bao, Wei-Sheng Lai, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Memc-net: Motion estimation and motion compensation driven neural network for video interpolation and enhancement. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. 2

[3] Jose Caballero, Christian Ledig, Andrew Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4778–4787, 2017. 1

[4] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(2):295–307, 2015. 1

[5] 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 *Int. Conf. Comput. Vis.*, December 2015. 2, 3

[6] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 6(6):721–741, 1984. 3

[7] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Recurrent back-projection network for video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3897–3906, 2019. 2, 6, 7

[8] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Space-time-aware multi-resolution video enhancement. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 3, 6, 7

[9] Jingwen He, Chao Dong, and Yu Qiao. Modulating image restoration with continual levels via adaptive feature modification layers. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11056–11064, 2019. 3

[10] Jingwen He, Chao Dong, and Yu Qiao. Multi-dimension modulation for image restoration with dynamic controllable residual learning. *arXiv preprint arXiv:1912.05293*, 2019. 3

[11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Int. Conf. Comput. Vis.*, pages 1026–1034, 2015. 6

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 1, 2

[13] Sepp Hochreiter and Jurgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 1

[14] Shi-Min Hu, Dun Liang, Guo-Ye Yang, Guo-Wei Yang, and Wen-Yang Zhou. Jittor: a novel deep learning framework with meta-operators and unified graph execution. *Science China Information Sciences*, 63, 222103, 2020. 6

[15] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9000–9008, 2018. 2, 6, 7

[16] Younghyun Jo, Seoung Wug Oh, Jaeyeon Kang, and Seon Joo Kim. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3224–3232, 2018. 2

[17] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Fisr: Deep joint frame interpolation and super-resolution with a multi-scale temporal loss. In *Association for the Advancement of Artificial Intelligence*, pages 11278–11286, 2020. 1, 2

[18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Int. Conf. Learn. Represent.*, 2015. 5

[19] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 624–632, 2017. 5

[20] Hyeongmin Lee, Taeoh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. Adacof: Adaptive collaboration of flows for video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5316–5325, 2020. 2

[21] Tao Li, Xiaohai He, Qizhi Teng, Zhengyong Wang, and Chao Ren. Space-time super-resolution with patch group cuts prior. *Signal Processing: Image Communication*, 30:147–165, 2015. 1

[22] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 136–144, 2017. 1

[23] Ce Liu and Deqing Sun. A bayesian approach to adaptive video super resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 209–216. IEEE, 2011. 2, 6, 7, 8

[24] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5

[25] Uma Mudenagudi, Subhashis Banerjee, and Prem Kumar Kalra. Space-time super-resolution using graph-cut optimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):995–1008, 2010. 1, 2

[26] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1701–1710, 2018. 2

[27] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5437–5446, 2020. 1, 2

[28] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *Int. Conf. Comput. Vis.*, pages 261–270, 2017. 2, 6, 7

[29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,

Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Adv. Neural Inform. Process. Syst.*, 2019. 6

[30] Oded Shahar, Alon Faktor, and Michal Irani. Space-time super-resolution from a single video. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2011. 1

[31] Eli Shechtman, Yaron Caspi, and Michal Irani. Increasing space-time resolution in video. In *Eur. Conf. Comput. Vis.*, pages 753–768. Springer, 2002. 1, 2

[32] Eli Shechtman, Yaron Caspi, and Michal Irani. Space-time super-resolution. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):531–545, 2005. 1

[33] Wenzhe Shi, Jose Caballero, Ferenc Huszar, 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1874–1883, 2016. 4

[34] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Adv. Neural Inform. Process. Syst.*, pages 802–810, 2015. 1, 3

[35] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1279–1288, 2017. 2, 6, 7, 8

[36] Matthew F Tang, Lucy Ford, Ehsan Arabzadeh, James T Enns, Troy AW Visser, and Jason B Mattingley. Neural dynamics of the attentional blink revealed by encoding orientation selectivity during rapid visual presentation. *Nature Communications*, 11(1):1–14, 2020. 2

[37] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally-deformable alignment network for video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3360–3369, 2020. 1, 2

[38] Roger Y. Tsai and ?Thomas S. Huang. Multiframe image restoration and registration. In *Advances in Computer Vision and Image Processing*, pages 317–339, 1984. 1

[39] ETSI TS 101 154 V2.3.1. *Digital Video Broadcasting (DVB); Specification for the use of Video and Audio Coding in Broadcasting Applications based on the MPEG-2 Transport Stream*. ETSI, Feb. 2017. 1

[40] Wei Wang, Ruiming Guo, Yapeng Tian, and Wenming Yang. Cfsnet: Toward a controllable feature space for image restoration. In *Int. Conf. Comput. Vis.*, pages 4140–4149, 2019. 3

[41] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 0–0, 2019. 2, 3, 4, 6, 7, 8

[42] Xintao Wang, Ke Yu, Chao Dong, Xiaoou Tang, and Chen Change Loy. Deep network interpolation for continuous imagery effect transition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1692–1701, 2019. 3

[43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibil-ity to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004. 6, 7

[44] Yu-Huan Wu, Shang-Hua Gao, Jie Mei, Jun Xu, Deng-Ping Fan, Rong-Guo Zhang, and Ming-Ming Cheng. Jcs: An explainable covid-19 diagnosis system by joint classification and segmentation. *IEEE Transactions on Image Processing*, 30:3113–3126, 2021. 1

[45] 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 Conf. Comput. Vis. Pattern Recog.*, pages 3370–3379, June 2020. 1, 2, 3, 4, 5, 6, 7

[46] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *Int. J. Comput. Vis.*, 127(8):1106–1125, 2019. 1, 2, 6, 7, 8

[47] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Eur. Conf. Comput. Vis.*, pages 286–301, 2018. 6, 7

[48] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9308–9316, 2019. 4, 5