

MoTIF: Learning Motion Trajectories with Local Implicit Neural Functions for Continuous Space-Time Video Super-Resolution

Yi-Hsin Chen* Si-Cun Chen* Yi-Hsin Chen Yen-Yu Lin Wen-Hsiao Peng
 National Yang Ming Chiao Tung University, Taiwan
 {yhchen12101, sicun.map1, karta6120}.cs09@nycu.edu.tw
 lin@cs.nycu.edu.tw wpeng@cs.nctu.edu.tw

Abstract

This work addresses continuous space-time video super-resolution (C-STVSR) that aims to up-scale an input video both spatially and temporally by any scaling factors. One key challenge of C-STVSR is to propagate information temporally among the input video frames. To this end, we introduce a space-time local implicit neural function. It has the striking feature of learning forward motion for a continuum of pixels. We motivate the use of forward motion from the perspective of learning individual motion trajectories, as opposed to learning a mixture of motion trajectories with backward motion. To ease motion interpolation, we encode sparsely sampled forward motion extracted from the input video as the contextual input. Along with a reliability-aware splatting and decoding scheme, our framework, termed MoTIF, achieves the state-of-the-art performance on C-STVSR. The source code of MoTIF is available at <https://github.com/sichun233746/MoTIF>.

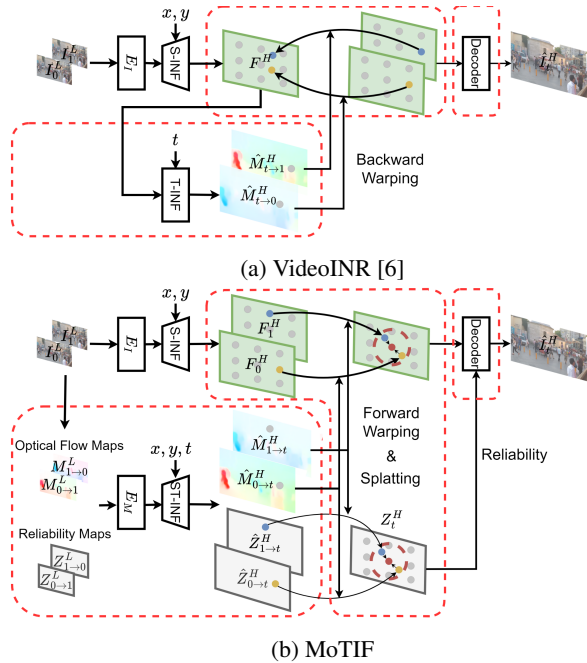


Figure 1: Illustrations of (a) VideoINR [6] and (b) MoTIF. The red dash lines highlight their major differences.

1. Introduction

This work addresses continuous space-time video super-resolution (C-STVSR). The task of C-STVSR is to increase simultaneously the spatial resolution and temporal frame-rate of an input video by any scaling factors with only one single model. It is to be distinguished from fixed-scale space-time video super-resolution (F-STVSR), for which a model is learned to perform space-time super-resolution for only one specific spatiotemporal scale. As compared to F-STVSR, C-STVSR is more flexible and practical in real-world scenarios, which often call for up-scaling low-resolution and low-frame-rate videos of varied spatiotemporal resolutions on heterogeneous video-enabled devices.

C-STVSR remains largely under-explored. One trivial solution to C-STVSR is to perform continuous video frame interpolation [2, 11, 21, 20, 33], followed by interpolat-

ing individual video frames with continuous image super-resolution [5, 34, 14], or the other way around. However, their divide-and-conquer nature of treating C-STVSR as two independent sub-tasks—i.e. temporal interpolation and spatial super-resolution—misses the opportunity to attain the best achievable performance. By leveraging the spatiotemporal information in an end-to-end optimized fashion, some recent works [10, 31, 32, 9] for F-STVSR adopt a one-stage approach, combining the extraction of individual frame features and the temporal aggregation of these features as a unified task. Nonetheless, these F-STVSR methods can hardly be extended straightforwardly to C-STVSR.

Inspired by continuous image super-resolution [5, 34, 14], VideoINR [6] presents an early attempt at C-STVSR. Given any query coordinates (x, y, t) in the continuous spa-

*Both authors contributed equally to this work.

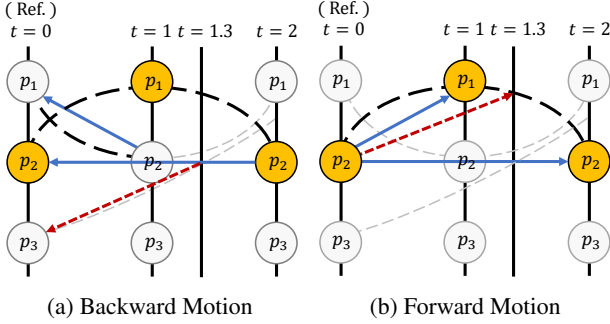


Figure 2: Illustration of backward and forward motion. The circles denote pixels accessible in the input video. The dashed lines display the motion trajectories of pixels in the reference frame at $t = 0$. The blue arrows are backward/forward motion in the form of displacement vectors. The red arrows show the displacement vectors for an arbitrary time instance that are to be predicted from blue arrows.

tiotemporal space, it takes the latent representation of the input video as the contextual information to decode the corresponding RGB value. The process involves learning a spatial implicit neural function (S-INF in Fig. 1 (a)) for super-resolving the frame features, followed by learning another temporal implicit neural function (T-INF in Fig. 1 (a)) to generate motion estimates at time t to backward warp the super-resolved frame features. However, learning implicitly *backward motion* (indicating displacement vectors that identify matching pixels/features in the reference frame) as a function of time is challenging. Essentially, the backward motion at the same spatial coordinates (x, y) yet at different time instances t may capture the motion trajectories of different pixels/features in the reference frame. For example, in Fig. 2 (a), the backward motion vectors of p_2 at $t = 1$ and $t = 2$ are governed by the two distinct motion trajectories that originate from pixels p_1 and p_2 in the reference frame at $t = 0$, respectively. In other words, the backward motion vectors at p_2 , when viewed as a function of time, are a mixture of multiple motion trajectories. This could potentially introduce undesirable randomness and discontinuities in the resulting time function, which must be learned by T-INF in Fig. 1 (a). Furthermore, learning implicitly such a time function based solely on frame features complicates the task.

To circumvent the aforementioned issues, we propose learning *forward motion* of pixels in the form of motion trajectories with a space-time implicit neural function (ST-INF in Fig. 1 (b)). Considering each reference frame in the input video as sitting at the origin in time, our ST-INF takes (x, y, t) as input and outputs a displacement vector that specifies where the pixel at the coordinates (x, y) of the reference frame will appear in a synthesized frame at time t . That is, it encodes the motion trajectory of the

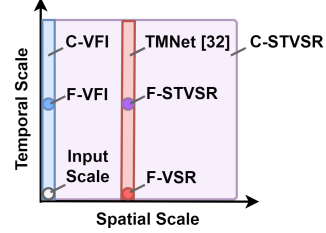


Figure 3: Illustration of fixed-scale video frame interpolation (F-VFI), continuous video frame interpolation (C-VFI), fixed-scale video super-resolution (F-VSR), fixed-scale space-time video super-resolution (F-STVSR), TM-Net [32], and continuous space-time video super-resolution (C-STVSR) in terms of their supported space-time scales.

pixel at (x, y) , e.g. the highlighted motion trajectory of p_2 in the reference frame at $t = 0$ in Fig. 2 (b). Moreover, to facilitate the learning of such a neural function in an explicit way, we supply *forward optical flow maps* estimated between reference frames as the contextual information (i.e. $M_{0 \rightarrow 1}^L, M_{1 \rightarrow 0}^L$ in Fig. 1 (b)). Our space-time neural function is also learned to predict the reliability of every motion trajectory (i.e. $\hat{Z}_{0 \rightarrow t}^H, \hat{Z}_{1 \rightarrow t}^H$ in Fig. 1 (b)), which is essential to ensure the quality of forward warping. Explicit motion modeling allows us to extract rough reliability estimates from the input video for better prediction.

Fig. 1 (b) depicts our end-to-end trainable C-STVSR framework, MoTIF. The main contributions of our work include: (1) we propose a space-time local implicit neural function that predicts *forward* motion and its reliability in a continuous manner; (2) we propose a reliability-aware splatting and decoding scheme that fuses simultaneously information from multiple reference frames; and (3) our MoTIF achieves the state-of-the-art performance on C-STVSR and provides out-of-distribution generalization.

2. Related Work

This section surveys methods for video frame interpolation, video super-resolution, and space-time video super-resolution. Fig. 3 presents a Venn diagram to illustrate how C-STVSR, the focus of our work, is related to these methods in terms of their supported space-time scales. As shown, the fixed-scale methods—e.g. fixed-scale video super-resolution [4], fixed-scale video frame interpolation [1, 13] and F-STVSR [9, 31, 10]—perform only one specific type of space-time interpolation. Their supported space-time scales are visualized as singletons in Fig. 3. In comparison, the continuous-scale methods—such as continuous video frame interpolation [2, 11, 21, 20, 33] and TMNet[32]—are able to cover a continuum of temporal scales. Of these approaches, C-STVSR [6] is the most flexible and challenging one, with its supported space-time scales covering the entire space-time space.

2.1. Video Frame Interpolation

Video frame interpolation [2, 11, 21, 33, 17, 13, 1] aims to increase the frame rate of a video by interpolating between existing reference frames. The key to successful frame interpolation is to predict how the pixels/features of the reference frames progress temporally to the interpolated frame. The flow-based methods [20, 2, 21, 19, 16, 11] rely on optical flow maps to propagate features/pixels from the neighboring reference frames, whereas the kernel-based methods [13, 7, 8] estimate motion implicitly as kernels for motion compensation with deformable convolution. Most flow-based approaches adopt backward warping [11, 16, 23, 22], but more recently, forward warping [2, 20, 21, 19] emerges as an attractive alternative. Forward warping, however, is faced with the challenge that multiple features/pixels in the reference frame may be mapped to the same location in the target frame. To tackle this issue, Niklaus *et al.* [21] introduce softmax splatting, weighting the conflicting features/pixels according to the reliability of their forward motion.

2.2. Video Super-Resolution

Video super-resolution is to increase the spatial resolution of a video. Its central theme is to exploit temporal information from neighboring frames in order to complement a low-resolution video frame in recovering its missing high-frequency details. Early deep learning-based methods [27, 3, 24, 35, 4] rely on optical flows to align the features/pixels of the neighboring frames. However, optical flow estimation can be expensive. As such, Tian *et al.* [29] adopt deformable convolution for temporal alignment. Wang *et al.* [30] extend the idea to perform temporal alignment in a coarse-to-fine manner. These works target fixed-scale video super-resolution.

2.3. Space-Time Video Super-Resolution

Recognizing that both video frame interpolation and video super-resolution involve aggregating temporal information from neighboring frames, Haris *et al.* [10] adopt a unified network to address space-time video super-resolution (STVSR). STVSR is much more challenging than the previous two tasks, as the low-resolution neighboring frames are the only source of information to interpolate a high-resolution video frame. Along this line of research, Xiang *et al.* [31] propose using bidirectional deformable ConvLSTM to mine useful space-time information from the input video in an end-to-end fashion. Based on [31], Xu *et al.* [32] introduce a temporal modulation block, which allows STVSR to be continuous in the temporal scale. By contrast, both [10] and [31] support only F-STVSR.

More recently, Chen *et al.* [6] present the first work on end-to-end learned C-STVSR, allowing both the spatial and temporal scales to be continuous. Inspired by [5], which

learns local implicit neural functions for continuous image super-resolution, their C-STVSR scheme includes a spatial and a temporal implicit neural function. The former generates the pixel features at any given spatial coordinates (x, y) for super-resolution, while the latter predicts the *backward* motion for any spatiotemporal coordinates (x, y, t) to propagate temporally the resulting features to time t . Both neural functions are local; they refer to neighboring latents extracted from the input video as additional contextual information.

3. Proposed Method

Given two low-resolution RGB video frames $I_0^L, I_1^L \in \mathbb{R}^{3 \times H \times W}$ of size $H \times W$, our task is to interpolate a high-resolution video frame $I_t^H \in \mathbb{R}^{3 \times H' \times W'}$ with an arbitrary scale $s = W'/W = H'/H \geq 1$ and at any time $t \in [0, 1]$.

3.1. System Overview

Fig. 4 depicts our proposed MoTIF, which comprises four major components and operates as follows. First, given I_0^L and I_1^L , (1) the encoder E_I converts them into their latent representations $F_0^L, F_1^L, F_{(0,1)}^L \in \mathbb{R}^{C \times H \times W}$, where $F_{(0,1)}^L$ serves as a rough estimate of the feature of the target frame I_t^H . Similar to recent STVSR works [32, 6], we adopt the off-the-shelf video-based encoder from [31], which fuses information from both I_0^L and I_1^L in generating F_0^L, F_1^L and $F_{(0,1)}^L$. Second, (2) the spatial local implicit neural function (S-INF) is queried to super-resolute F_0^L, F_1^L as $F_0^H, F_1^H \in \mathbb{R}^{C \times H' \times W'}$, respectively. Our S-INF follows the design of LIIF [5]. Third, considering I_0^L as sitting at the origin in time, (3) the motion encoder E_M encodes $M_{0 \rightarrow 1}^L \in \mathbb{R}^{2 \times H \times W}$ —namely, the forward optical flow map capturing the forward motion from I_0^L to I_1^L —together with its reliability map $Z_{0 \rightarrow 1}^L \in \mathbb{R}^{3 \times H \times W}$ into $T_0^L \in \mathbb{R}^{C \times H \times W}$. The optical flow estimation is not always perfect; $Z_{0 \rightarrow 1}^L$ indicates how reliable $M_{0 \rightarrow 1}^L$ is across spatial locations (x, y) (Section 3.2). Forth, using T_0^L as the motion latent, (4) our space-time local implicit neural function (ST-INF) renders a high-resolution, forward motion map $\hat{M}_{0 \rightarrow t}^H \in \mathbb{R}^{2 \times H' \times W'}$ and its reliability map $\hat{Z}_{0 \rightarrow t}^H \in \mathbb{R}^{H' \times W'}$ according to the query space-time coordinates (x, y, t) . $\hat{M}_{0 \rightarrow t}^H$ specifies the forward motion of the features in F_0^H and is utilized to forward warp F_0^H to F_t^H (Section 3.2). The same motion encoding, rendering and warping processes are repeated for I_1^L , in aggregating temporally the information from all the reference frames. Lastly, we follow [21] to perform softmax splatting to create F_t^H and Z_t^H , which are further combined with $F_{(0,1)}^H$ to decode the high-resolution video frame \hat{I}_t^H at time t (Section 3.3). Z_t^H indicates how good F_t^H is across spatial locations. It is used to condition the pixel-based decoding of the RGB values from F_t^H and $F_{(0,1)}^H$.

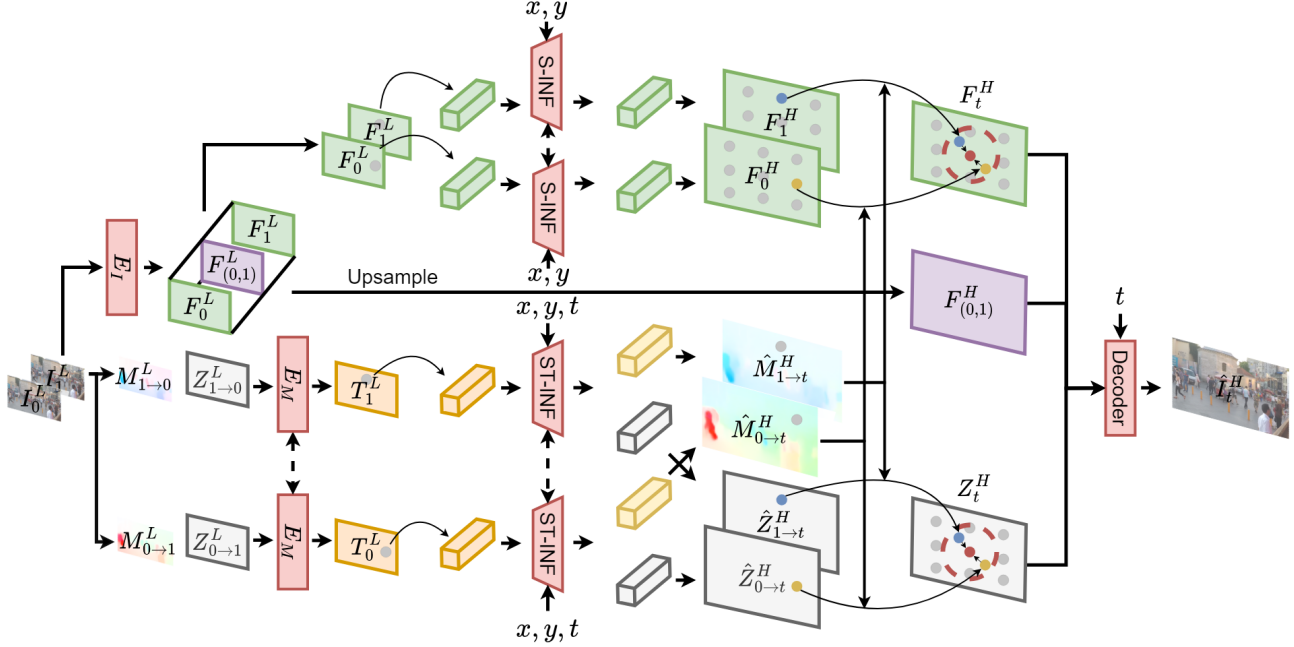


Figure 4: The proposed MoTIF for C-STVSR, where the dash double arrows represent the shared-weight networks.

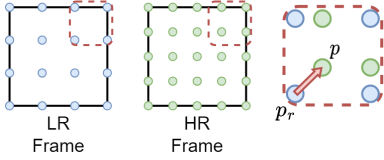


Figure 5: Illustration of low-resolution coordinates (blue dots) and high-resolution coordinates (green dots).

3.2. Space-time Local Implicit Neural Functions

The very core of our C-STVSR scheme is the space-time local implicit neural function (ST-INF) in Fig. 4. Our ST-INF has the striking feature of predicting forward motion rather than backward motion. That is, it specifies how the feature at coordinates $p = (x, y)$ in F_0^H or F_1^H are propagated temporally to any designated time t . The forward motion is represented in the form of displacement vectors along with their reliability values. For example, to get the forward motion $\hat{M}_{0 \rightarrow t}^H(p)$ and its reliability value $\hat{Z}_{0 \rightarrow t}^H(p)$ for propagating the feature $F_0^H(p)$ of F_0^H at $p = (x, y)$, it is queried as follows:

$$\{\hat{Z}_{t_r \rightarrow t}^H(p), \hat{M}_{t_r \rightarrow t}^H(p)\} = f_\theta(v_r, p - p_r, t - t_r), \quad (1)$$

where $v_r = T_0^L(p_r)$ is the motion latent at $p_r = (x_r, y_r)$ that is nearest to the query coordinates $p = (x, y)$, $t_r = 0$ is the temporal location where the reference frame I_0^L sits, and θ is the network parameters. Fig. 5 depicts an example of the geometrical relationship between p and p_r . The sum $p + \hat{M}_{0 \rightarrow t}^H(p)$ gives the landing location of the query feature

$F_0^H(p)$ at time t . In much the same way, $\hat{M}_{1 \rightarrow t}^H(p)$ and $\hat{Z}_{1 \rightarrow t}^H(p)$ for propagating the feature $F_1^H(p)$ can be obtained by having in Eq. (1) $v_r = T_1^L(p_r)$ and $t_r = 1$, i.e. the temporal location of I_1^L .

In Eq. (1), both $p = (x, y)$ and t can take any values. Together they can refer to any space-time coordinates. Therefore, f_θ is able to generate forward motion in a continuous manner to warp F_0^H, F_1^H of any spatial resolution to any time instance $t \in [0, 1]$. However, in essence, f_θ is a local function that predicts forward motion in the vicinity of the reference space-time coordinates p_r, t_r by referring to the local motion latent v_r .

Learning motion trajectories. Learning forward motion can be interpreted as learning motion trajectories along the temporal axis. To see this, in Eq. (1), we fix $p = (x, y)$ at some coordinates, e.g. p_2 in Fig. 2 (b), take $t_r = 0$, and view f_θ as a function of time t . With this setting, the forward motion predicted by f_θ specifies a displacement vector indicating where p_2 should appear at time t . Collectively, the displacement vectors evaluated at different time instances t 's define the motion trajectory of p_2 . Generally, this motion trajectory is a smooth function of time and is relatively easier to approximate. While it is completely feasible to change the output semantics of f_θ to learn backward motion, the resulting time function can be discontinuous. The reason is illustrated in Fig. 2 (a), where fixing the query coordinates p at p_2 , f_θ returns at every time instance t a backward displacement vector identifying the location of the matching pixel/feature in the reference frame

at $t_r = 0$. In this case, the displacement vectors evaluated for the same p_2 yet at different time instances t 's may correspond to the distinct motion trajectories of different matching pixels. This suggests that f_θ has to model a less smooth function of time. Section 5.1 presents an ablation study to justify the use of forward motion.

Learning motion latents. Predicting the forward motion of a pixel (or a feature vector) at any given $p = (x, y)$ and for any t is a non-trivial task. We formulate the problem as learning a f_θ that interpolates between forward motion sampled sparsely in both the spatial and temporal dimensions. This is achieved by providing f_θ with the motion latent that encodes the sparsely sampled forward motion as the contextual input. Take Eq. (1) as an example, where f_θ is queried to predict the forward motion of $F_0^H(p)$ for time t . The prediction is conditioned on the nearest motion latent $T_0^L(p_r)$, which captures the forward motion $M_{0 \rightarrow 1}^L$ estimated from I_0^L to I_1^L in the vicinity of p_r . In this work, we adopt Raft-lite [28] to estimate the forward optical flow map $M_{0 \rightarrow 1}^L$. Recognizing that the flow estimation is often not perfect, we follow [19] to quantify the reliability of the resulting flow map $M_{0 \rightarrow 1}^L$ based on three metrics, including (1) the intensity warping error, (2) the flow warping error, and (3) the local variances of the flow map. Further details of these metrics are provided in the supplementary document. The reliability evaluation with each of these metrics yields a real-valued map of size the same as $M_{0 \rightarrow 1}^L$, reflecting the reliability of $M_{0 \rightarrow 1}^L$ across spatial locations. These maps are concatenated channel-wisely to form $Z_{0 \rightarrow 1}^L$, which is encoded jointly with $M_{0 \rightarrow 1}^L$ by the motion encoder E_M as T_0^L . Section 5.1 shows that $Z_{0 \rightarrow 1}^L$ benefits f_θ considerably in interpolating forward motion.

3.3. Multi-Frame Forward Warping

To come up with a prediction of F_t^H for decoding a high-resolution video frame \hat{I}_t^H at time t , we aggregate temporally F_0^H, F_1^H , each of which represents the high-resolution feature of a reference frame (Fig. 4). Inspired by [21], we adopt softmax splatting to resolve the potential issue that multiple features from F_0^H, F_1^H or both may be forward warped to the same location in F_t^H . Considering that our task is to interpolate and super-resolute a new frame from the ground up, we perform softmax splatting after F_0^H, F_1^H have both been forward warped to time t . Our approach differs from [21], which targets video frame interpolation and applies softmax splatting separately to individual reference frames for late fusion. In symbols, we have

$$F_t^H(p) = \sum_{i=0}^1 \sum_q \frac{b(u) \cdot \exp(\alpha \cdot \hat{Z}_{i \rightarrow t}^H(q)) \cdot F_i^H(q)}{\sum_{i=0}^1 \sum_q b(u) \cdot \exp(\alpha \cdot \hat{Z}_{i \rightarrow t}^H(q))}, \quad (2)$$

where the feature $F_t^H(p)$ of F_t^H at p is formulated as a weighted sum of all the reference features $F_0^H(q), F_1^H(q)$, with the weighting determined by the distance $u = p - (q + \hat{M}_{i \rightarrow t}^H(q))$, the bilinear kernel $b(u) = \max(0, 1 - |u_x|) \cdot \max(0, 1 - |u_y|)$, as well as the reliability $\hat{Z}_{i \rightarrow t}^H(q)$ of the forward motion at q . $\alpha = -20$ is the temperature of the softmax operation. Since the bilinear kernel has a finite support, only those $F_0^H(q), F_1^H(q)$ warped to the neighborhood of p will actually contribute to the evaluation of $F_t^H(p)$.

Additionally, we generate a map Z_t^H to indicate how good F_t^H is across spatial locations. Intuitively, if $F_t^H(p)$ is synthesized from those $F_0^H(q), F_1^H(q)$ whose forward motion is unreliable, the quality of $F_t^H(p)$ should be downgraded. $Z_t^H(p)$ serves as a conditioning factor for decoding the RGB values at p , and is obtained by

$$Z_t^H(p) = \max_{i=0,1} \max_q b(u) \cdot \exp(\alpha \cdot \hat{Z}_{i \rightarrow t}^H(q)), \quad (3)$$

which takes the maximum value among the (unnormalized) contributing weights from $F_0^H(q), F_1^H(q)$. When none of these contributing $F_0^H(q), F_1^H(q)$ has reliable forward motion, the quality of $F_t^H(p)$ is regarded as poor.

To synthesize a high-resolution video frame \hat{I}_t^H , we implement a pixel-wise decoder that incorporates a multi-layer perceptron. It decodes the RGB values at p by taking as inputs $F_t^H(p), F_{(0,1)}^H(p), Z_t^H(p)$, and t (the rightmost part of Fig. 4).

3.4. Training Objective

We train our MoTIF end-to-end with the following objective:

$$\mathcal{L} = \mathcal{L}_{char}(\hat{I}_t^H, I_t^H) + \beta \sum_{i=0}^1 \mathcal{L}_{char}(\hat{M}_{i \rightarrow t}^H, M_{i \rightarrow t}^H), \quad (4)$$

where $\mathcal{L}_{char}(\hat{x}, x) = \sqrt{\|\hat{x} - x\|^2 + \epsilon^2}$ is the Charbonnier loss [12] and β is a hyper-parameter. ϵ, β are set empirically to 10^{-3} and 0.01, respectively. Our objective requires both the decoded frame \hat{I}_t^H and the predicted forward motion $\hat{M}_{i \rightarrow t}^H$ to approximate their respective ground-truths.

3.5. Comparison with Prior Works

Both our MoTIF and VideoINR [6] use implicit neural functions to tackle space-time video super-resolution (STVSR). As illustrated in Fig. 1, our MoTIF differs from VideoINR [6] in three significant ways:

First, for the C-STVSR task, our MoTIF uses *forward* motion rather than *backward* motion. This aspect in its own right has a significant impact on the quality of the generated videos (Section 5.1 and Table 4).

Second, our MoTIF models motion *explicitly* rather than *implicitly*. This allows our ST-INF to directly learn to interpolate between motion trajectories derived from a pre-trained optical flow estimation model. The supplementary

document provides additional results, showing that our MoTIF can work well with well-behaved, off-the-shelf optical flow estimation networks. Using explicit motion also allows us to evaluate the reliability information $Z_{0 \rightarrow 1}^L, Z_{1 \rightarrow 0}^L$ based on the input video for better predicting $\hat{Z}_{0 \rightarrow t}^H, \hat{Z}_{1 \rightarrow t}^H$ (see Fig. 4).

Third, our MoTIF introduces the reliability-aware splatting and decoding schemes, which are not seen in VideoINR [6]. Their benefits are studied in Section 5.1 and Table 6.

Different from [21], our reliability-aware splatting adopts early fusion of reference frames by forward warping all the reference features to the target frame according to Eq. (2). In contrast, [21] applies softmax splatting to each individual reference frame, followed by late fusing the results with a synthesis network. Our reliability-aware decoding scheme, which incorporates the reliability information for decoding (Eq. (3)) is not seen in [21].

4. Experiments

To our best knowledge, VideoINR [6] is the only prior work that addresses specifically C-STVSR. We thus follow its training and test protocols, unless otherwise specified. VideoINR [6] is also included as the major baseline method.

Training Datasets. We train our model on Adobe240 dataset [26], which contains 133 720P hand-held videos. Of these videos, 100 are used for training, 16 for validation, and 17 for test. In each video, we take 9 consecutive frames to form a training sample, where the 1st and 9th frames are bicubic down-sampled and used as the low-resolution, low-frame-rate input.

Evaluation. We compare the competing methods on Vid4 [15], Adobe240 [26], and Gopro [18] datasets. Unless otherwise specified, the spatial scaling factor defaults to 4. On Vid4, the temporal scaling factor is fixed at 2 to test single-frame interpolation. On Adobe240-*average* and Gopro-*average*, the temporal scaling factor is set to 8 for multi-frame interpolation. Under the same setting, we also report results on Adobe240-*center* and Gopro-*center* for only the 1st, 4th and 9th frames (namely, single-frame interpolation).

Baselines and Quality Metrics. The baseline methods include (1) two-stage F-STVSR methods, namely video frame interpolation (SuperSloMo [11], QVI [33], DAIN [2]) plus video super-resolution (Bicubic Interpolation, EDVR [30], BasicVSR [4]); (2) one-stage F-STVSR methods (Zooming SloMo [31]); (3) two-stage C-STVSR methods, namely continuous video frame interpolation (SuperSloMo [11], DAIN [2]) plus continuous image super-resolution (LIIF [5]); (4) one-stage C-STVSR methods (VideoINR [6]); and (5) TMNet [32]. The quality metrics

are Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) on the Y channel.

Implementation and Training Details. We adopt the same two-stage training strategy as VideoINR [6]. The spatial scaling factor is set to 4 for the first 450,000 iterations, and is chosen uniformly from [1, 4] in the following 150,000 iterations. The training batch size is 24; within each batch, every input frame is down-sampled spatially by the same factor and cropped to 32×32 . For training stability, we use the ground-truth forward motion in place of the predicted forward motion with a certain probability, the value of which is attenuated from 1 to 0 in the first 150,000 iterations. We adopt Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$ and cosine annealing to decay the learning rate from 10^{-4} to 10^{-7} for every 150,000 iterations. For data augmentation, we perform random rotation and horizontal-flipping. Both our S-INF and ST-INF (Fig. 4) are implemented with 3-layer SIRENs [25], the hidden dimensions of which are 64, 64, and 256. More network details are in the supplementary document.

4.1. Comparison with State-of-the-art Methods

Table 1 presents qualitative results, comparing the competing methods on the F-STVSR task. Both VideoINR [6] and our MoTIF are trained for C-STVSR, whereas the other methods are trained for F-STVSR and their results are excerpted from [6]. Notably, VideoINR-*fixed* is trained specifically for single-frame interpolation. From Table 1, several observations can be made. (1) Our MoTIF outperforms VideoINR [6] in all the test cases. It also outperforms VideoINR-*fixed*, although not trained for F-STVSR. (2) While both VideoINR [6] and our MoTIF adopt the same E_I encoder from Zooming SloMo [31], VideoINR [6] performs worse than Zooming SloMo [31] under the single-frame interpolation on Vid4, GoPro-*Center*, and Adobe-*Center*; on the contrary, our MoTIF is superior to Zooming SloMo [31]. (3) On Vid4, our MoTIF performs slightly worse than TMNet [32]. This may be because TMNet [32] is trained on Vimeo-90K dataset [35], which shares similar characteristics to Vid4. (4) All the one-stage methods (our MoTIF, [32, 6, 31]) performs better than the two-stage methods (video frame interpolation plus video super-resolution) due to end-to-end optimization. (5) Our MoTIF (with Raft-lite [28] included) has a similar model size to VideoINR [6]. Section 5 further shows that MoTIF (including Raft) has comparable or even lower GMACs than VideoINR, and thus similar or higher FPS.

Table 3 further presents results on the C-STVSR task, with most of the spatiotemporal scaling factors not seen during training. Except TMNet [32], which supports continuous temporal scaling but only 4x spatial scaling, all the methods are able to achieve C-STVSR. Again, our MoTIF achieves the best performance in all the test cases, con-

Table 1: Performance comparison on the F-STVSR task. **Red**, **blue**, and **green** indicate the best, the second best, and the third best performance, respectively. Quality metrics: PSNR/SSIM.

VFI Method	VSR Method	Vid4	GoPro-Center	GoPro-Average	Adobe-Center	Adobe-Average	Parameters (Millions)
SuperSloMo [11]	Bicubic	22.42 / 0.5645	27.04 / 0.7937	26.06 / 0.7720	26.09 / 0.7435	25.29 / 0.7279	19.8
SuperSloMo [11]	EDVR[30]	23.01 / 0.6136	28.24 / 0.8322	26.30 / 0.7960	27.25 / 0.7972	25.95 / 0.7682	19.8+20.7
SuperSloMo [11]	BasicVSR [4]	23.17 / 0.6159	28.23 / 0.8308	26.36 / 0.7977	27.28 / 0.7961	25.94 / 0.7679	19.8+6.3
QVI [33]	Bicubic [30]	22.11 / 0.5498	26.50 / 0.7791	25.41 / 0.7554	25.57 / 0.7324	24.72 / 0.7114	29.2
QVI [33]	EDVR [30]	23.60 / 0.6471	27.43 / 0.8081	25.55 / 0.7739	26.40 / 0.7692	25.09 / 0.7406	29.2+20.7
QVI [33]	BasicVSR [4]	23.15 / 0.6428	27.44 / 0.8070	26.27 / 0.7955	26.43 / 0.7682	25.20 / 0.7421	29.2+6.3
DAIN [2]	Bicubic	22.57 / 0.5732	26.92 / 0.7911	26.11 / 0.7740	26.01 / 0.7461	25.40 / 0.7321	24.0
DAIN [2]	EDVR [30]	23.48 / 0.6547	28.01 / 0.8239	26.37 / 0.7964	27.06 / 0.7895	26.01 / 0.7703	24.0+20.7
DAIN [2]	BasicVSR [4]	23.43 / 0.6514	28.00 / 0.8227	26.46 / 0.7966	27.07 / 0.7890	26.23 / 0.7725	24.0+6.3
Zooming SlowMo [31]		25.72 / 0.7717	30.69 / 0.8847	- / -	30.26 / 0.8821	- / -	11.10
TMNet [32]		25.96 / 0.7803	30.14 / 0.8692	28.83 / 0.8514	29.41 / 0.8524	28.30 / 0.8354	12.26
VideoINR-fixed [6]		25.78 / 0.7730	30.73 / 0.8850	- / -	30.21 / 0.8805	- / -	11.31
VideoINR [6]		25.61 / 0.7709	30.26 / 0.8792	29.41 / 0.8669	29.92 / 0.8746	29.27 / 0.8651	11.31
Ours		25.79 / 0.7745	31.04 / 0.8877	30.04 / 0.8773	30.63 / 0.8839	29.82 / 0.8750	12.55

Table 2: PSNR/SSIM performance comparison on the C-STVSR task (on Gopro). **Bold** indicates the best performance.

Temporal Scale	Spatial Scale	SuperSloMo [11] + LIIF [5]	DAIN [2] + LIIF [5]	TMNet [32]	VideoINR [6]	Ours
×6	×4	26.70 / 0.7988	26.71 / 0.7998	30.49 / 0.8861	30.78 / 0.8954	31.56 / 0.9064
×6	×6	23.47 / 0.6931	23.36 / 0.6902	-	25.56 / 0.7671	29.36 / 0.8505
×6	×12	21.92 / 0.6495	22.01 / 0.6499	-	24.02 / 0.6900	25.81 / 0.7330
×12	×4	25.07 / 0.7491	25.14 / 0.7497	26.38 / 0.7931	27.32 / 0.8141	27.77 / 0.8230
×12	×6	22.91 / 0.6783	22.92 / 0.6785	-	24.68 / 0.7358	26.78 / 0.7908
×12	×12	21.61 / 0.6457	21.78 / 0.6473	-	23.70 / 0.6830	24.72 / 0.7108
×16	×4	24.42 / 0.7296	24.20 / 0.7244	24.72 / 0.7526	25.81 / 0.7739	25.98 / 0.7758
×16	×6	23.28 / 0.6883	22.80 / 0.6722	-	23.86 / 0.7123	25.34 / 0.7527
×16	×12	21.80 / 0.6481	22.22 / 0.6420	-	22.88 / 0.6659	23.88 / 0.6923
×6	×1	-	-	-	32.34 / 0.9545	34.77 / 0.9696
×1	×4	-	-	33.02 / 0.9206	32.26 / 0.9198	33.84 / 0.9328

firming its better generalization to unseen scaling factors. Notably, on the video frame interpolation task (i.e. temporal scale = 6 and spatial scale = 1), MoTIF outperforms VideoINR [6] by 2.5dB in terms of PSNR. This underlines the merit of using forward motion for better modeling.

In Fig. 6, our MoTIF shows consistently better subjective quality than VideoINR [6]. More results are provided in the supplementary document.

5. Complexity Comparison

Table 3 characterizes the complexity of the competing methods [31, 32, 6]. We follow [31] to report frames per second (FPS) on Vid4 [15] dataset; that is, FPS is evaluated to be the ratio of the total number of output frames to the total runtime for processing the entire dataset. We also report the corresponding multiply-accumulate (MAC) operations per frame. These numbers are evaluated on one Tesla V100. From Table 3, our MoTIF has comparable FPS and GMACs to the other baseline methods on the lower-scale tasks (i.e. temporal scale = 2 and spatial scale = 4) while showing higher FPS and lower GMACs than the base-

line methods on the tasks with higher temporal and spatial scales.

Note that when we increase the temporal scale while fixing the spatial scale, the FPS increases and the MAC per frame decreases. The same observation holds true for all the competing methods. This is because higher temporal scales invoke less frequent feature extraction to generate $F_0^L, F_1^L, F_{(0,1)}^L$. For example, a temporal scale of 16 (respectively, 2) implies that the feature extraction process is invoked only once every 16 frames (respectively, 2 frames). Given that the total number of frames to be processed is fixed, more frequent feature extraction leads to lower FPS and higher GMACs per frame. It is also seen that the complexity advantage of our MoTIF over VideoINR becomes more obvious when the temporal scale becomes higher and the spatial scale remains fixed. This is mainly because VideoINR backward warps the latent representations $F_0^H, F_1^H, F_{(0,1)}^H \in \mathbb{R}^{C \times H' \times W'}$ simultaneously and this operation is done twice. In comparison, our MoTIF forward warps the latent $F_0^H, F_1^H \in \mathbb{R}^{C \times H' \times W'}$ individually and this operation is done only once.

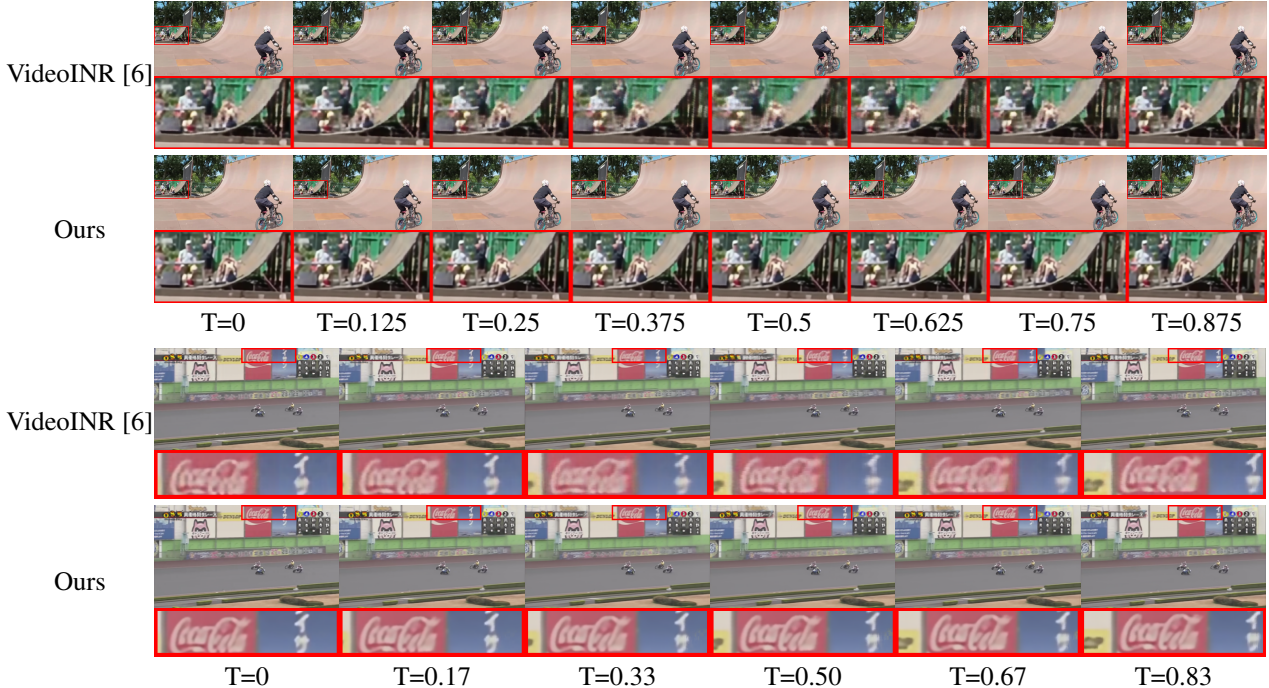


Figure 6: Subjective quality comparison. The temporal scaling factor of the upper example is 8 (in-distribution), whereas that of the lower example is 6 (out-of-distribution). Zoom in for better visualization.

Table 3: Complexity comparison on the C-STVSR task. **Red** indicates the best performance. Complexity metrics: FPS (\uparrow) / GMACs (\downarrow) per frame. The FPS and GMACs are evaluated based on processing the entire Vid4 [15] dataset on one Tesla V100.

Temporal Scale	Spatial Scale	ZSM [31]	TMNet[32]	VideoINR[6]	Ours
$\times 2$	$\times 4$	25.52 / 42.27	24.75 / 45.36	16.05 / 51.18	15.15 / 52.11
$\times 4$	$\times 4$	-	23.68 / 46.30	21.69 / 36.59	20.70 / 33.90
$\times 8$	$\times 4$	-	22.84 / 46.57	26.17 / 26.65	27.71 / 21.49
$\times 16$	$\times 4$	-	21.54 / 48.15	29.06 / 21.50	42.50 / 14.69
$\times 2$	$\times 6$	-	-	11.76 / 69.51	13.59 / 70.82
$\times 4$	$\times 6$	-	-	15.13 / 54.93	17.51 / 47.91
$\times 8$	$\times 6$	-	-	17.63 / 44.85	21.20 / 32.54
$\times 16$	$\times 6$	-	-	18.14 / 40.18	25.20 / 24.46
$\times 2$	$\times 8$	-	-	9.24 / 95.19	10.72 / 97.44
$\times 4$	$\times 8$	-	-	10.40 / 80.61	13.84 / 68.95
$\times 8$	$\times 8$	-	-	11.93 / 70.34	16.30 / 49.22
$\times 16$	$\times 8$	-	-	14.09 / 66.32	23.93 / 35.49

5.1. Ablation Experiments

Backward vs. Forward Motion. Table 4 presents results for an ablation experiment that replaces forward motion with backward motion in our MoTIF. This replacement includes the following changes: (1) learning ST-INF to predicting backward motion $\hat{M}_{t \rightarrow 0}^H, \hat{M}_{t \rightarrow 1}^H$ and their reliability maps $\hat{Z}_{t \rightarrow 0}^H, \hat{Z}_{t \rightarrow 1}^H$, (2) applying backward warping with $\hat{M}_{t \rightarrow 0}^H, \hat{M}_{t \rightarrow 1}^H$ to each reference feature F_0^H, F_1^H , and (3) synthesizing a high-resolution video frame \hat{I}_t^H by taking as inputs the two backward warped reference features, their warped reliability maps, $F_{(0,1)}^H$, and t . From Table 4, using

backward motion instead of forward motion in MoTIF results in a considerable PSNR drop (0.4-1dB) across the test cases. Our supplementary document provides additional Fourier analyses to compare forward and backward motion.

Implicit vs. Explicit Motion Modeling. Table 5 presents ablation results based on predicting the high-resolution forward motion without using a pre-trained optical flow estimation network, i.e. the implicit motion modeling. In this case, F_0^L, F_1^L (see Fig. 4) are used as inputs to our ST-INF. Compared with the explicit method, the implicit method has 0.1-0.4dB PSNR loss. Notably, even with the implicit method, MoTIF outperforms VideoINR by 0.1-0.4dB in

Table 4: Backward vs. forward motion in MoTIF. Quality metrics: PSNR/SSIM. **Bold** indicates the best performance.

	Backward	Forward
Vid4	25.35 / 0.7696	25.79 / 0.7745
GoPro-Center	29.98 / 0.8765	31.04 / 0.8877
GoPro-Average	29.38 / 0.8693	30.04 / 0.8773
Adobe-Center	29.73 / 0.8723	30.63 / 0.8839
Adobe-Average	29.14 / 0.8658	29.82 / 0.8750

Table 5: Explicit vs. implicit motion modeling in MoTIF. Quality metrics: PSNR/SSIM. **Bold** indicates the best performance.

	VideoINR	MoTIF (Implicit)	MoTIF (Explicit)
Vid4	25.61 / 0.7709	25.71 / 0.7721	25.79 / 0.7745
GoPro-Center	30.26 / 0.8792	30.58 / 0.8856	31.04 / 0.8877
GoPro-Average	29.41 / 0.8669	29.81 / 0.8744	30.04 / 0.8773
Adobe-Center	29.92 / 0.8746	30.24 / 0.8796	30.63 / 0.8839
Adobe-Average	29.27 / 0.8651	29.59 / 0.8719	29.82 / 0.8750

PSNR, suggesting that the other components of MoTIF are essential.

Feature Warping and Reliability Maps. Table 6 presents ablation results to understand the contributions of different components in MoTIF. Four variants of MoTIF are investigated, including (a) using only $F_{(0,1)}^H$ for decoding, (b) using both $F_{(0,1)}^H$ and F_t^H for decoding, (c) using $F_{(0,1)}^H, F_t^H$ for decoding while encoding the reliability maps $Z_{0 \rightarrow 1}^L, Z_{1 \rightarrow 0}^L$ of forward motion into the motion latents, and (d) the full model (i.e. variant (c) plus Z_t^H). From Table 6, the considerable PSNR gain of (b) over (a) indicates that our ST-INF is effective in interpolating forward motion for propagating the reference features. The incremental improvement from (b) to (c) suggests that the reliability maps $Z_{0 \rightarrow 1}^L, Z_{1 \rightarrow 0}^L$ help to improve the quality of the interpolated forward motion. Last but not least, the additional use of Z_t^H in the decoding process ((d) vs. (c)) does allow the decoder to better combine $F_{(0,1)}^H$ and F_t^H .

Tri-linear Motion and More Reference Frames. Table 7 investigates the benefits of our space-time implicit neural function (ST-INF) by comparing its performance with tri-linear motion interpolation and by showing its applicability to more reference frames. From Table 7, we observe that when ST-INF is replaced with tri-linear motion interpolation—i.e., $\hat{M}_{0 \rightarrow t}^H, \hat{M}_{1 \rightarrow t}^H$ are interpolated tri-linearly from $M_{0 \rightarrow 1}^L, M_{1 \rightarrow 0}^L$, respectively—the performance drops by 0.1-0.2dB in PSNR. Although ST-INF provides seemingly moderate gain, its advantage becomes obvious when the number of reference frames goes beyond two. In this case, ST-INF can benefit from encoding more forward motion into the motion latents. Conceptually, this amounts to taking more forward motion samples, which help to construct accurate motion trajectories. In the ablation experiment, two more reference frames I_{-1}^L, I_2^L

Table 6: Ablation experiment on individual components. (d) is the proposed MoTIF. Quality metric: PSNR.

Settings	(a)	(b)	(c)	(d)
$F_{(0,1)}^H$	V	V	V	V
F_t^H		V	V	V
Z_0^L, Z_1^L			V	V
Z_t^H				V
Vid4	22.38	25.26	25.61	25.79
GoPro-Center	26.68	30.54	30.97	31.04
GoPro-Average	26.44	29.72	29.97	30.08
Adobe-Center	25.82	30.04	30.49	30.63
Adobe-Average	25.61	29.40	29.68	29.82

Table 7: Ablation experiment on tri-linear motion interpolation and multiple reference frames. Quality metrics: PSNR/SSIM.

	Tri-linear Motion	Ours (2 ref.)	Ours (4 ref.)
Vid4	25.57 / 0.7728	25.79 / 0.7745	26.32 / 0.7864
GoPro-Center	30.89 / 0.8860	30.96 / 0.8868	31.44 / 0.9003
GoPro-Average	29.93 / 0.8759	30.08 / 0.8780	30.77 / 0.8948
Adobe-Center	30.42 / 0.8818	30.63 / 0.8839	31.03 / 0.8919
Adobe-Average	29.64 / 0.8727	29.82 / 0.8750	30.37 / 0.8849

are made available; we thus encode jointly information from $\{M_{0 \rightarrow i}^L, Z_{0 \rightarrow i}^L\}_{i=-1,1,2}$ as T_0^L , and information from $\{M_{1 \rightarrow i}^L, Z_{1 \rightarrow i}^L\}_{i=-1,0,2}$ as T_1^L . As a consequence, the PSNR improves by 0.5-0.7dB. We expect the gain to be even higher if we generate more motion latents to propagate information from these extra reference frames (instead of only I_0^L, I_1^L). The simple tri-linear motion interpolation cannot benefit similarly from having more reference frames.

6. Conclusion

This paper introduces a C-STVSR framework, known as MoTIF. It features a space-time implicit neural function for encoding forward motion, and a reliability-aware splatting and decoding scheme for fusing spatiotemporal information from multiple reference frames. We show that learning forward motion amounts to learning individual motion trajectories rather than a mixture of motion trajectories as with learning backward motion. In addition, for better aggregating temporal information via forward warping, performing splatting and decoding based on the reliability of forward motion is crucial. With all these techniques combined, MoTIF demonstrates superior quantitative and qualitative performance to the state-of-the-art methods for C-STVSR.

Acknowledgement

This work is supported by National Science and Technology Council, Taiwan under Grants NSTC 111-2634-F-A49-010- and MOST 110-2221-E-A49-065-MY3, and National Center for High-performance Computing.

References

- [1] 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 Transactions on Pattern Analysis and Machine Intelligence*, pages 933–948, 2021.
- [2] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [4] Kelvin C.K. 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4947–4956, 2021.
- [5] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8628–8638, 2021.
- [6] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2047–2057, 2022.
- [7] Xianhang Cheng and Zhenzhong Chen. Video frame interpolation via deformable separable convolution. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):10607–10614, 2020.
- [8] Tianyu Ding, Luming Liang, Zhihui Zhu, and Ilya Zharkov. Cdfi: Compression-driven network design for frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8001–8011, 2021.
- [9] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17441–17451, 2022.
- [10] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Space-time-aware multi-resolution video enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [11] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [12] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5835–5843, 2017.
- [13] Hyeongmin Lee, Taehoon Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. Adacof: Adaptive collaboration of flows for video frame interpolation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [14] Jaewon Lee and Kyong Hwan Jin. Local texture estimator for implicit representation function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [15] Ce Liu and Deqing Sun. A bayesian approach to adaptive video super resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, page 209–216, 2011.
- [16] Z. Liu, R. A. Yeh, X. Tang, Y. Liu, and A. Agarwala. Video frame synthesis using deep voxel flow. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4473–4481, 2017.
- [17] Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, and Christopher Schroers. Phasenet for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] 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.
- [19] Simon Niklaus, Ping Hu, and Jiawen Chen. Splatting-based synthesis for video frame interpolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2017.
- [20] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [21] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [22] Junheum Park, Keunsoo Ko, Chul Lee, and Chang-Su Kim. Bmbc: Bilateral motion estimation with bilateral cost volume for video interpolation. In *European Conference on Computer Vision (ECCV)*, 2020.
- [23] Junheum Park, Chul Lee, and Chang-Su Kim. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14539–14548, 2021.
- [24] Mehdi S. M. Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [25] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit

- neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020.
- [26] Shuochen 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.
- [27] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [28] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European Conference on Computer Vision (ECCV)*, 2020.
- [29] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally-deformable alignment network for video super-resolution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [30] Xintao Wang, Kelvin C.K. Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [31] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [32] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6388–6397, 2021.
- [33] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. *Quadratic Video Interpolation*. 2019.
- [34] Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Ultrasr: Spatial encoding is a missing key for implicit image function-based arbitrary-scale super-resolution. *arXiv preprint arXiv:2103.12716*, 2021.
- [35] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 127(8):1106–1125, 2019.