

Deep Space-Time Video Upsampling Networks

Jaeyeon Kang¹[0000-0003-2172-0026], Younghyun Jo¹[0000-0002-8530-9802],
Seoung Wug Oh¹[0000-0002-8498-0864], Peter Vajda²[0000-0002-2031-4678], and
Seon Joo Kim^{1,2}[0000-0001-8512-216X]

¹Yonsei University, ²Facebook

Abstract. Video super-resolution (VSR) and frame interpolation (FI) are traditional computer vision problems, and the performance have been improving by incorporating deep learning recently. In this paper, we investigate the problem of jointly upsampling videos both in space and time, which is becoming more important with advances in display systems. One solution for this is to run VSR and FI, one by one, independently. This is highly inefficient as heavy deep neural networks (DNN) are involved in each solution. To this end, we propose an end-to-end DNN framework for the space-time video upsampling by efficiently merging VSR and FI into a joint framework. In our framework, a novel weighting scheme is proposed to fuse all input frames effectively without explicit motion compensation for efficient processing of videos. The results show better results both quantitatively and qualitatively, while reducing the computation time ($\times 7$ faster) and the number of parameters (30%) compared to baselines. Our source code is available at <https://github.com/JaeYeonKang/STVUN-Pytorch>.

Keywords: Video Super-Resolution, Video Frame Interpolation, Joint Space-Time Upsampling

1 Introduction

In this paper, we introduce a method of upsampling both the spatial resolution and the frame rate of a video simultaneously. This is an important problem as more high-performance TV displays are being introduced with higher resolution and frame rate, but the video contents have not yet caught up with the capabilities of displays. For example, new UHD displays now come with 4K or even 8K resolution, and the frame rate of 120 fps. On the other hand, most available contents are still HD (1080p) or less in resolution, with the frame rate of 30 fps. Another potential application of this problem is the video replay for sports and security videos. In order to inspect a video in much detail, videos are spatially magnified in slow motion. As shown by these examples, there is definitely a major need for a framework that can convert a given video into a video with higher resolution and frame rate.

Deep neural networks (DNN) have become common solutions for video super-resolution (VSR) and frame interpolation (FI) recently. With DNN, an obvious solution for the joint upsampling problem would be to sequentially run a VSR




Fig. 1: Common pipelines for VSR/FI and the design scheme for our proposed network. By efficiently merging the pipelines for the two tasks with new mechanisms for feature fusion and interactions between modules, we can accurately upsample videos both in space and time in a very efficient fashion. E: extraction, A: alignment, F: fusion, R: reconstruction, I: interpolation.

network followed by a FI network or vice versa. However, running the two algorithms independently is computationally expensive and inefficient, as the state-of-the-art methods for each task employ heavy DNNs. The goal of this paper is to design an efficient DNN for the joint space-time upsampling problem by investigating shareable components between the spatial and the temporal upsampling tasks.

While there are many different DNN architecture for VSR [29,34,23,32] and FI [34,9,2,20], the common design schemes can be summarized as in Fig. 1(a). In VSR, most methods employ four stages – feature extraction, alignment, fusion and reconstruction. For FI, the process can be divided into feature extraction, feature interpolation and reconstruction.

To jointly upsample videos both in space and time, we propose to combine the two tasks in an efficient manner as shown in Fig. 1(b) by sharing the common modules in feature extraction and reconstruction. The modules are designed to interact and learn simultaneously for accurate and efficient reconstruction of jointly upsampled videos.

Furthermore, we propose a novel way to efficiently fuse the features of individual frames for VSR without explicit motion compensation. Most VSR methods rely on aligning many input frames through optical flows [34,23,8,29] or deformable convolutions [32] before fusing the extracted features. As many methods use up to 7 input frames, aligning that many frames takes up a large portion of the computation. To remove the computational burden of motion compensation, we fuse the feature maps without explicit alignment step (Fig. 1(b)). In our feature fusion process, we propose Early Fusion with Spatio-Temporal weights (EFST) module that learns to fuse information by considering spatio-temporal relationship between input frames in an implicit manner. In this module, learnable spatio-temporal weights are computed in order to combine rich information from all frames instead of focusing too much on the target frame.

There are no public datasets available for the joint space-time upsampling, as it is a relatively a new topic. While there are many datasets for VSR and FI separately, they are not ideal for the joint upsampling task. To this end, we

collected a new dataset called the Space-Time Video Test (STVT) dataset that can be used to evaluate joint upsampling methods. This dataset will be publicly available.

In summary, the main contributions of our paper are as follows:

- By efficiently merging two networks of VSR and FI, we propose a novel framework called the Space-Time Video Upsampling Networks (STVUN) for joint space-time video upsampling. With careful design of each module and their interactions, we produce better results while reducing the computation time ($\times 7$ faster) and the number of parameters (30%) compared to sequentially connected state-of-the-art VSR and FI networks.
- We propose Early Fusion with Spatio-Temporal weights (EFST) to fuse input features efficiently without explicit motion compensation for VSR.
- Our framework can deal with more challenging upsampling tasks as it can upsample 4×4 in space and $\times \infty$ in time. In comparison, recent works on joint upsampling have only shown results on doubling the resolution both in space and time ($2 \times 2 \times 2$).
- We collected Space-Time Video Test (STVT) dataset for evaluating the joint space-time upsampling task. This can be very useful for future work in this topic.

2 Related Work

2.1 Video Super-Resolution

After Dong et al. [6] have successfully achieved the high performance by incorporating deep learning into the single image SR task, deep learning approaches have also become prevalent in solving the VSR problem [13, 11, 4, 16, 29, 23].

DUF [10] used dynamic up-sampling filters to improve the resolution while reducing the flickering artifact which is prevalent in VSR task. Their method takes the advantage of the implicit motion computed within the network, and additionally used the learned residual image to enhance the sharpness. RBPN [8] used an iterative refinement framework, which forwards the input frame with other frames at multiple times. They use the idea of back-projection, which computes a residual image for each time step to reduce the error between the target and the output. In EDVR [32], input frames are first aligned with the target frame using the deformable convolution [5]. Aligned frames are then fused using the temporal and spatial attention (TSA) mechanism.

2.2 Video Frame Interpolation

Video frame interpolation can be roughly divided into two categories: kernel-based methods and optical flow-based methods. As an interpolation kernel based approach, Niklaus et al. [21] proposed AdaConv, which produces interpolation kernels to generate intermediate frame. In [22], they extended the method to reduce the computational cost, which is named as SepConv using 1D kernels instead of 2D kernels.

With the introduction of CNN-based optical flow algorithms [7], several frame interpolation algorithms using the optical flow have been developed. Liu et al. [18] produce intermediate frames by the trilinear sampling based on the estimated deep voxel flows called DVF. Xue et al. [35] used the bi-directional flow to warp both input frames using the backward warping function. Jiang et al. [9] obtained the bi-directional flow through the network and then linearly transformed two flows with respect to the time value to generate multiple intermediate frames. Niklaus and Liu [20] used the forward warping and further designed a refinement network in order to fill the holes caused by the forward warping. Liu et al. [17] used cycle consistency loss to enhance synthesized frames to be more reliable as input frames. To deal with the occlusion problem which is a common issue in optical flows, additional depth information was used to refine the optical flows in DAIN [2].

2.3 Space-Time Upsampling

In [26], Shechtman et al. first proposed a space-time super-resolution framework by using multiple low resolution (LR) videos of the same dynamic scene. Different from the frame interpolation methods mentioned above, they explicitly deal with the motion blur to generate sharp interpolated frames. In [24], Shahar et al. extended the work in [26] with a method that only uses a single video to enhance the resolution. Sharma et al. [25] first used a DNN architecture for the joint space-time upsampling. They used the auto-encoder to learn the mapping between LR and high resolution (HR) frames, and the frame interpolation was simply done by the tri-cubic interpolation. Another deep joint upsampling method called FISR [12] was recently introduced, which targets for estimating 4K, 60fps video from 2K, 30fps video. They regularized their joint upsampling network by forwarding multiple chunks of frames into one iteration and set multiple temporal losses at the output of each chunk. Note that, FISR only generates 8 pixels (space $\times 2$, time $\times 2$) per input pixel, while our work aims at more challenging task of generating more pixels (e.g. 64 pixels for space $\times 4$, time $\times 4$).

3 Space-time Video Upsampling Algorithm

Given a sequence of LR frames X_t , our method produces HR frames \hat{Y}_t of inputs as well as the intermediate HR frames \hat{Y}_T in-between the input frames. The term t denotes the input time index, and T indicates the newly created time index. The size of a LR frame is $H \times W \times C$, where H , W , and C are the height, the width, and the number of channels respectively. The output size is $rH \times rW \times C$, with r being the spatial upscaling factor. We can generate N multiple upsampled intermediate frames in-between the two input frames. The problem is very challenging as the algorithm has to generate $r^2 \times (N + 1)$ pixels per pixel in the input frame. For example, we need to generate 64 pixels in the output per input pixel with $r = 4$ and $N = 3$.

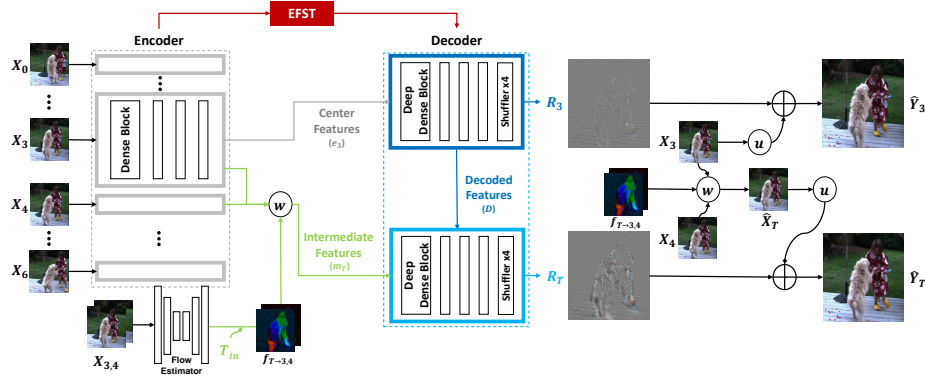


Fig. 2: Overview of our space-time upsampling network. The network is composed of several interacting modules to produce the HR frames of the center frame \hat{Y}_3 and the intermediate frames \hat{Y}_T , where $T \in [3, 4]$. The term w and u indicate backward warping and bilinear upsampling respectively.

3.1 Network Overview

The overview of our network is shown in Fig. 2. Our network is composed of multiple modules: encoder, feature fusion for spatial upsampling (EFST), flow estimator for frame interpolation, and decoder. Our framework takes 7 LR frames as inputs, for example $X_{[0,1,2,3,4,5,6]}$. Then, it produces the HR frames for the center frame \hat{Y}_3 as well as N HR intermediate frames \hat{Y}_T , where $T \in [3, 4]$.

The encoders that share weights are first used to extract features per frame. The encoded features are fused using EFST for the spatial upsampling, and interpolated using the computed flows for the temporal upsampling. The decoding block that consists of decoders with shared weights produces residual images for the spatial and the temporal upsampling, both of which are added to the bilinearly upsampled images to produce the final output frames.

3.2 Network Details

Encoder Structure of the encoder is shown in Fig. 3. The encoder extracts feature representations for each frame and consists of multiple dense convolution blocks. Each dense block is connected to the corresponding block in decoder through EFST. This allows the decoder to keep considering the temporal relationship of inputs. The encoded features are expressed as e_t^i , where i is the block index and t is the time index of the input frame.

Early Fusion with Spatio-Temporal weights (EFST) In most VSR methods, features from multiple frames are aligned before fusion using explicit motion compensation by optical flows or deformable convolutions. However, aligning multiple frames (6 in most cases including ours) to the center frame is computationally expensive. Therefore, we exclude explicit alignment process by devising

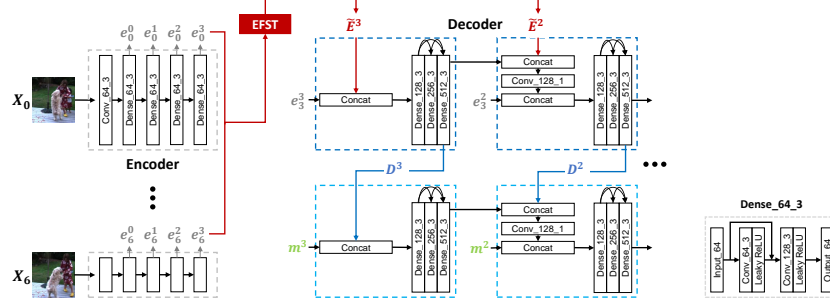


Fig. 3: The structure of the encoder and decoder. Dense_64_3 denotes dense block with 64 input channel dimension and 3×3 kernel size. We use Leaky ReLU with the slope value 0.1.

the EFST module for implicit feature alignment and fusion for spatial upsampling.

To merge input features from the encoders, we first apply early fusion to reduce the computational cost. Early fused features E^i are defined as $\text{Conv}(\text{Concat}[e_0^i, \dots, e_6^i])$ where Conv reduces the channel dimension by the factor of 7 with an 1×1 convolution filter. However, since the early fusion will collapse all temporal information in the first layer, the features of the target (center) frame will be mainly used as mentioned in [27]. Some information in other frames may vanish due to the bottleneck. This is because most information for reconstructing the HR target frame is contained in the input LR center frame.

In order to use valuable information in the features from all the input frames without explicit alignment, we propose the EFST module that computes spatio-temporal weights to compensate E^i . The structure of EFST is shown in Fig. 4. In the early fusion result E^i , most input features are not considered equally since E^i will be computed to mainly focus on the center frame. Therefore, we design a confidence score to effectively fuse informative features from the neighbor frames as well as the center frame. We estimate the confidence score by computing dot-product between E^i and all the e_t^i . We use this confidence score as a temporal attention to find which frames need to be more referred. The confidence score is computed as follows:

$$s_t^i = \text{sigmoid}(\theta(e_t^i) \circ \delta(E^i)), \quad (1)$$

where \circ is dot-product and s is the confidence score. θ and δ are single convolutional layer with filter size 1×1 . s_t^i has the same spatial size as e_t^i and the values of s_t^i are in $[0, 1]$. To pay more attention to the frames with high confidence score, we multiply this value to the original encoded features e_t^i as follows:

$$\bar{e}_t^i = s_t^i \odot e_t^i, \quad (2)$$

where \odot denote element-wise multiplication.

All temporally weighted encoded features \bar{e}_t^i are then concatenated and forwarded to pyramid designed convolutional layers to further consider spatio-

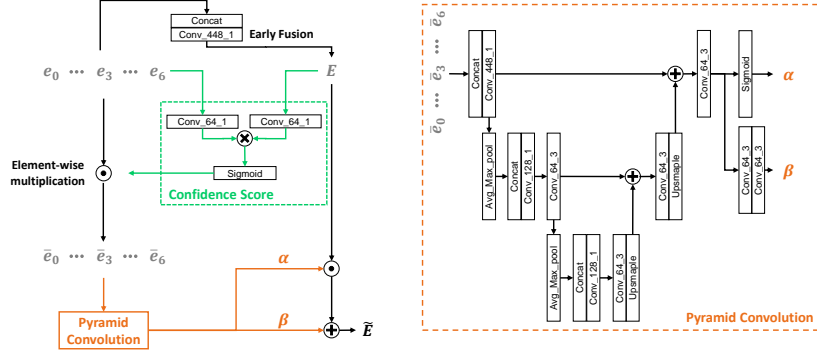


Fig. 4: Early Fusion with Spatio-Temporal weights (EFST) module. Avg_Max_pool means pooling separately with average pooling and max pooling. For temporally attending more frames, the confidence score is computed from early fusion result E^i and e_t^i . Then, the spatio-temporal weights α and β are computed and applied. For simplicity, we omit the superscript i , and see text for details.

temporal information. Pyramid convolution can effectively enlarge the receptive field with just few convolution layers. Afterwards, we generate learnable spatio-temporal weights α, β . It is a tensor with same size of E^i . It transform the initial early fusion result to learn the alignment in an implicit way. The final fused features is computed as follows:

$$\tilde{E}^i = \alpha \odot E^i + \beta. \quad (3)$$

Our EFST module is similar to Fusion with Temporal and Spatial Attention called TSA in [32]. TSA measures similarity distance between aligned frames and target frame to temporally weight more on well-aligned frames, since misalignment can severely interfere with learning. In comparison, we use confidence score as a way to involve more features from more input frames, which eventually works as a joint alignment and fusion process without explicit alignment.

Intermediate Feature Interpolation The feature interpolation process is shown in green in Fig. 2. Features m_T^i of an intermediate frame are synthesized by warping the features of the input frames (X_3, X_4) using the optical flow estimated by the flow estimator. To reduce the computational cost, we warp the encoded features of the two inputs to produce the target intermediate features instead of first creating the intermediate frame and encoding it again. Note that explicit alignment is only used for intermediate frames but not for merging all input features (EFST). We first use the optical flow computed by the PWC-Net [28]. Then, we use the following formulation [9] to estimate the flow between the input frames and the intermediate frame:

$$\begin{aligned} f_{T \rightarrow 3} &= -(1 - T_{in})T_{in}f_{3 \rightarrow 4} + T_{in}^2f_{4 \rightarrow 3}, \\ f_{T \rightarrow 4} &= (1 - T_{in})^2f_{3 \rightarrow 4} - T_{in}(1 - T_{in})f_{4 \rightarrow 3}, \end{aligned} \quad (4)$$

where f indicates the optical flow and T_{in} is a relative scalar value (e.g. when we want to get $\hat{Y}_{3.5}$, then T_{in} is set to 0.5). Note that T_{in} is given as an input to the network to provide the time index of the intermediate frame.

Finally, the features of the intermediate frame are computed as follows:

$$m_T^i = \frac{w(e_3^i, f_{T \rightarrow 3}) + w(e_4^i, f_{T \rightarrow 4})}{2}, \quad (5)$$

where w stands for the backward warping.

At the same time, we generate LR intermediate frames X_T for the subsequent process of bilinear upsampling with

$$\hat{X}_T = \frac{w(X_3, f_{T \rightarrow 3}) + w(X_4, f_{T \rightarrow 4})}{2}, \quad (6)$$

and for finetuning PWC-net with ground-truth intermediate frames during training whole network, we set loss function as:

$$\mathcal{L}_{\mathcal{M}} = \sum_T \|\hat{X}_T - X_T\|_1. \quad (7)$$

Decoder The decoder reconstructs target HR residual image and it consists of multiple dense convolution blocks. The same number of blocks is used to connect with each block of the encoder. We design a more deeper dense block which is shown in Fig. 3, since more layers and connections could boost performance [14,33]. To generate the residual image of the target frame, features from the last layer of the last block are convolved with a filter having $C \cdot r \cdot r$ output channels, where the output is then reshaped to the size of $rH \times rW \times C$ through the pixel shuffler [27] with the scale factor of 4.

For space upsampling, features of target (center) frame e_3 and fused features from EFST \bar{E} are used as inputs. Output residual image R_3 is added to the upsampled target frame to generate final HR output as follows:

$$\hat{Y}_3 = u(X_3) + R_3, \quad (8)$$

where $u(\cdot)$ is the bilinear upsampling function. The loss function for space upsampling is defined as:

$$\mathcal{L}_{\mathcal{S}} = \|\hat{Y}_3 - Y_3\|_1. \quad (9)$$

Since the decoder consists of more deeply stacked convolution layers, it creates more refined spatio-temporal information from the EFST features. Therefore, during the space upsampling task, the decoded features D^i are generated at the end of each dense block and forwarded to the space-time upsampling task to supplement with more rich information. For space-time upsampling, the intermediate features m_T are passed to another decoder that shares weights. Here, different from space upsampling, decoded features D^i are fed instead of feeding \bar{E} .

Then, HR residual image R_T of the intermediate frame is generated and the final HR intermediate frame is computed as follows:

$$\hat{Y}_T = u(\hat{X}_T) + R_T, \quad (10)$$

where \hat{X}_T is from Eq. (6). Our loss function to train space-time upsampling is:

$$\mathcal{L}_{\mathcal{F}} = \sum_T \|\hat{Y}_T - Y_T\|_1, \quad (11)$$

where T can be any values in $[3, 4]$. Note that we can generate arbitrary number of intermediate frames using Eq. (4).

3.3 Training

Vimeo septuplets dataset [35] is usually used to train VSR and FI tasks. But the length of video frames in Vimeo dataset is too short for our task. It consists of 7 frames per clip, but we need at least 8 frames for training. Therefore, we collect training videos of 240fps from YouTube. This training dataset consists of various scenes with global camera motions and local object motions. In total, the dataset contains about 1800 video clips and 220K frames. To make LR frames, HR frames are first smoothed with a Gaussian filter and then subsampled with respect to the scaling factor $r = 4$. For the data augmentation, we randomly flip left-right and rotate 90/180 degrees. We also reverse the order of the sequence to enlarge the training dataset. The whole training and test is processed in RGB channels.

It is difficult to train all the networks in our framework simultaneously from scratch, as there are many interactions between the components. We first pre-train the encoder and the spatial decoder by minimizing $\mathcal{L}_{\mathcal{S}}$ (VSR part only). For this pretraining, we use 7 frames $Y_{[0,1,2,3,4,5,6]}$ in the training dataset and 128×128 patches are cropped. We use the Adam optimizer for 300K iterations with the mini-batch size of 32. The learning rate is initialized to 0.0001 and decreased by a factor of 2 every 100K iterations.

After pretraining the VSR part, we train the whole network using the following total loss function:

$$\mathcal{L} = \lambda_{\mathcal{M}}\mathcal{L}_{\mathcal{M}} + \lambda_{\mathcal{S}}\mathcal{L}_{\mathcal{S}} + \lambda_{\mathcal{F}}\mathcal{L}_{\mathcal{F}}, \quad (12)$$

where $\lambda_{\mathcal{M}}$, $\lambda_{\mathcal{S}}$, and $\lambda_{\mathcal{F}}$ are the weight parameters. In our experiment, we empirically set $\lambda_{\mathcal{M}} = 1$, $\lambda_{\mathcal{S}} = 1$, and $\lambda_{\mathcal{F}} = 1$ for the best results. For the joint training, 256×256 patches are used rather than 128×128 in order to deal with large motions. Intermediate frames in-between Y_3, Y_4 (e.g. $Y_{3.5}$) as well as the 7 frames are used for training VSR and FI part together. We train the whole network for 400K iterations and the initial learning rate is set to 0.00005. The same learning rate decay is used.

4 Experiments

In this section, we provide both quantitative and qualitative evaluations of our algorithm.

Testsets While there are some datasets for VSR and FI separately, they are not ideal for the joint space-time upsampling task. For example, the Vid4 testset [15] for VSR have a lot of details, but the motion between the frames is too small. This limits the assessment of FI performance. MPI Sintel testset [3] is synthetic dataset which dose not have much detail to assess VSR performance. REDS-VTSR dataset [19] is used for VSR and FI separately, but it contains unnatural camera movements. In addition, the Vimeo [34], Middlebury [1] and FISR [12] testset are not available, since at least 15 frames are required for each scene to evaluate the performance.

To this end, we use Vid4, MPI Sintel and REDS-VTSR [19] for testing the generalization our performance. In addition, we create Space-Time Video Test (STVT) dataset that consists of 12 dynamic scenes with both natural motions and spatial details for the joint upsampling evaluation. Each scene has at least 50 frames, and we will make STVT dataset publicly available to promote more research in this topic.

Baselines We make two baseline methods ($V \rightarrow F$ and $F \rightarrow V$) that combine VSR and FI, which run consequently. V and F indicate VSR and FI respectively. For example, $F \rightarrow V$ indicates running FI first and then VSR. For V and F , we use EDVR [32] and DAIN [2] respectively, which are the state-of-the-art methods with publicly available codes. As the bias of the dataset affect the evaluation performance [30], for fair comparison, we try to finetune the baseline methods with our YouTube training dataset. However, since their weights are already highly finetuned, we find that the performance is rather reduced when we jointly train both networks at the same time (0.15dB is reduced for Vid4 testset). Therefore, we fix their weights to produce the results. We also compare our method with FISR [12], the only deep learning based work that we can compared to at this moment.

4.1 Comparisons

For the evaluation, we extract odd numbered frames in the testset and set them as ground-truth frames. Only the even numbered frames are used to generate the space-time upsampled results. We first compare our method with the two baseline methods. We set $T_{in} = 1/2$ for generating the HR intermediate frame. Table 1 shows the quantitative results of different approaches for $\times 4$ space and $\times 2$ time. In every testset, $F \rightarrow V$ consistently shows the worst performance, because FI works better on HR input frames due to sufficient details. On the other hand, in the case of $V \rightarrow F$, FI can access sufficient details from VSR, thus it can generate sharper results. However, the improvement in the resolution increases the amount of computation for FI ($\times 4$ slower).

Our results show better performance for all datasets as shown in Table 1. The performance difference in Vid4, Sintel and REDS-VTSR testset is not that big because those testsets are not constructed for this particular tasks and not optimal for evaluating the joint upsampling task. The performance gap between our method and the baselines become larger with the STVT dataset, which is specifically designed for the joint upsampling.

Table 1: Quantitative evaluation of the joint space-time upsampling on multiple testsets. We compare our method with the two baseline approaches by measuring the PSNR and SSIM. We set $T_{in} = 1/2$ for comparison. We also write down the number of parameters and the running time for each method. The running time is measured when generating the results with the resolution 960×540 . The best is shown in bold.

		$F \rightarrow V$	$V \rightarrow F$	Ours
Dataset	Vid4	25.22/0.7506	26.39/0.8163	26.49/0.8231
	Sintel	26.99/0.7986	27.56/ 0.8185	27.58 /0.8134
	REDS-VTSR	23.70/0.6541	23.63/0.6533	23.78/0.6601
	STVT	26.43/0.8435	26.96/0.8619	27.23/0.8644
#Params		44.7M	44.7M	30.9M
Running Time		0.52s	2.14s	0.30s

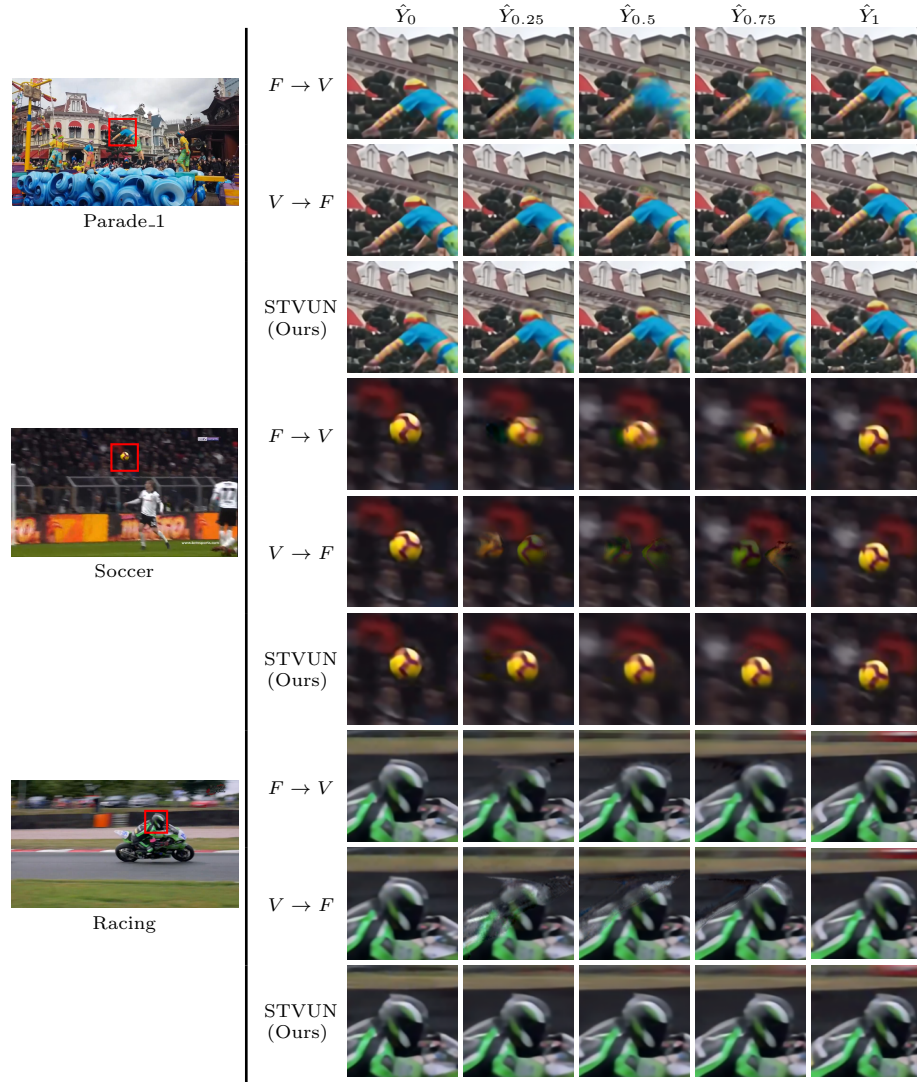
Table 2: Comparison with FISR [12]. We train our model with the upsampling factor space $\times 2$ and time $\times 2$ which is the same as FISR.

		FISR [12]	Ours
Dataset	Vid4	26.93/0.8534	30.60/0.9369
	Sintel	27.17/0.8115	28.36/0.8329
	REDS-VTSR	23.89/0.6601	23.66/0.6550
	STVT	26.49/0.8514	28.01/0.8895
#Params		62.3M	30.9M
Running Time		1.10s	0.98s

Table 1 also shows the number of parameters and the computation time of different methods. In this experiment, our total parameters and computational times does include PWC-Net [28]. We run the methods on Nvidia Geforce Titan X and measure the time taken to generate one 960×540 jointly upsampled frame. The number of parameters is reduced by more than 30% compared to the baseline methods, and the speed is 7 times faster than $V \rightarrow F$ and 1.7 times faster than $F \rightarrow V$. Although ours is lighter than the baseline methods, it exceeds the performance of baseline methods, indicating that our network is designed efficiently.

Additionally, we compare our model with FISR [12] in Table 2. As the upsampling factor of FISR is space $\times 2$ and time $\times 2$, we train our network with the same settings. Note that only the number of output channels of the last convolutional layer in the decoder is changed. As can be seen in Table 2, our method outperforms FISR by a large margin except for REDS-VTSR. In addition, our method runs faster than FISR with fewer parameters.

Fig. 5 visually compares our method with the two baseline methods and FISR. In Fig. 5(a), we generate multiple frame ($T_{in} = 0.25, 0.5, 0.75$) in-between two input frames. To better illustrate the results, we enlarge the corresponding red areas. As STVT dataset has a large motion, two baseline methods have



(a) Comparison with baseline methods on STVT dataset.



(b) Comparison with FISR on REDS-VTSR testset.

Fig. 5: Visual comparisons of the space-time upsampling results. In (a), we generate a total of 5 frames that consist of 2 space upsampling and 3 intermediate frames. *Parade_1*, *Soccer*, *Racing* scenes are used in our *STVT* dataset. In (b), we generate one intermediate frame. *002* and *007* in REDS-VTSR testset are used.

Table 3: Ablation studies on the EFST and our network structure. STVT dataset is used for comparison.

	w/o EFST	w/o D	w/ A&F	Ours
PSNR/SSIM	27.06/0.8613	27.15/0.8615	27.20/ 0.8652	27.23 /0.8644
#Params	30.5M	30.9M	32.4M	30.9M
Running Time	0.27s	0.30s	0.75s	0.30s

difficulty in handling the large motion. In the soccer scene, $F \rightarrow V$ shows more pleasing result than $V \rightarrow F$ because it is easier to estimate the motion at smaller input size. Except for the large motion scene, $V \rightarrow F$ is clearer than $F \rightarrow V$. Overall, our method is more accurate in estimating the motion and shows less artifacts. In Fig. 5(b), we generate one intermediate frame ($T_{in} = 0.5$) for comparison with FISR. The results in FISR show ghost artifact due to wrong motion estimation, but ours restore sharper edge details. However, due to the unnatural movement of REDS-VTSR testset, most center frames are not in the middle of the front and rear frames. So, only for this testset, the blurry results of FISR reduce average pixel error than ours. We recommend watching our demo video in the supplementary material to see the difference more clearly.

Beside the STVUN, our network can be used for VSR. As our main objective is the space-time upsampling, the experiments on VSR will be shown in the supplementary material.

4.2 Ablation Studies

We conduct ablation studies to investigate the contribution of EFST and our network design. Table 3 summarizes the ablation results. First, we test our model without EFST (w/o EFST), which means only the early fusion is used to fuse input features. This test demonstrates the effectiveness of EFST as it shows that our final model improves the performance without the large difference in running time.

To show the effectiveness of using decoded features D^i for space-time upsampling, we test our model when EFST features is used instead (w/o D). The performance gain shows learned features from space upsampling enhance the space-time upsampling results, indicating the decoder learn more rich information from EFST features.

We also evaluate our model with the explicit alignment (w/ A&F). The overall structure is the same as our proposed method except for the alignment and fusion parts. We use two modules in EDVR [31] – Pyramid, Cascading and Deformable Convolution (PCD) for the alignment and TSA for the feature fusion. As the explicit alignment process is a computational burden, it increases the running time by about 2.5 times. But the performance gap is minor, demonstrating EFST can effectively fuse features without explicit motion compensation.

Fig. 6 shows the visualization of the confidence score to analyze how confidence score changes with learning in EFST. In the early stages of training, confi-




Fig. 6: We visualize the confidence score s^i . Green box numbers are the time index, and the blue box numbers are the number of iterations. High confidence score is shown in red, and this means to be more referred and dark blue is vice versa. In the last row, we overlap results of the input and color map of 320k iterations. Zoom in to see better visualization.

dence scores are ambiguous to determine where to concentrate more. Therefore, the overall scores are high and shown in dark red. As the learning progresses, the confidence score gets the ability to determine the important parts among all inputs. High confidence scores remain for the regions which are helpful for reconstructing the center frame. On the other hand, occluded regions such as under the trunk lid in frame 0 have low confidence score because they are unnecessary for reconstructing center frame. It demonstrates that our confidence score effectively fuses features from all frames without explicit alignment.

5 Conclusion

In this paper, we present a deep Space-Time Video Upsampling Networks (STVUN) for joint space-time video upsampling by merging VSR and FI network efficiently. This task has many practical applications, yet a challenging task as the network has to perform two tasks in an efficient manner. In addition, we propose Early Fusion with Spatio-Temporal weights (EFST) modules that learns to fuse information by considering spatio-temporal relationship without any explicit alignment. Our network can generate visually pleasing results with reduced computational time ($\times 7$) and number of parameters (30%) compared to sequentially connected VSR and FI networks. Our method also outperforms a previous space-time upsampling task by a large margin.

Acknowledgements This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2014-0-00059, Development of Predictive Visual Intelligence Technology)

References

1. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. *International Journal of Computer Vision* **92**(1), 1–31 (2011) 10
2. Bao, W., Lai, W.S., Ma, C., Zhang, X., Gao, Z., Yang, M.H.: Depth-aware video frame interpolation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3703–3712 (2019) 2, 4, 10
3. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: A. Fitzgibbon et al. (Eds.) (ed.) *European Conf. on Computer Vision (ECCV)*. pp. 611–625. Part IV, LNCS 7577, Springer-Verlag (Oct 2012) 10
4. Caballero, J., Ledig, C., Aitken, A., Acosta, A., Totz, J., Wang, Z., Shi, W.: Real-time video super-resolution with spatio-temporal networks and motion compensation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017) 3
5. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 764–773 (2017) 3
6. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: *European Conference on Computer Vision*. pp. 184–199. Springer (2014) 3
7. Fischer, P., Dosovitskiy, A., Ilg, E., Häusser, P., Hazırbaş, C., Golkov, V., Van der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. *arXiv preprint arXiv:1504.06852* (2015) 4
8. Haris, M., Shakhnarovich, G., Ukita, N.: Recurrent back-projection network for video super-resolution. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3897–3906 (2019) 2, 3
9. Jiang, H., Sun, D., Jampani, V., Yang, M.H., Learned-Miller, E., Kautz, J.: Super sloMo: High quality estimation of multiple intermediate frames for video interpolation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 9000–9008 (2018) 2, 4, 7
10. Jo, Y., Wug Oh, S., Kang, J., Joo Kim, S.: Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3224–3232 (2018) 3
11. Kappeler, A., Yoo, S., Dai, Q., Katsaggelos, A.K.: Video super-resolution with convolutional neural networks. *IEEE Transactions on Computational Imaging* **2**(2), 109–122 (2016) 3
12. Kim, S.Y., Oh, J., Kim, M.: Fsr: Deep joint frame interpolation and super-resolution with a multi-scale temporal loss. In: *AAAI*. pp. 11278–11286 (2020) 4, 10, 11
13. Liao, R., Tao, X., Li, R., Ma, Z., Jia, J.: Video super-resolution via deep draft-ensemble learning. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 531–539 (2015) 3
14. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 136–144 (2017) 8
15. Liu, C., Sun, D.: On bayesian adaptive video super resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(2), 346–360 (2014) 10

16. Liu, D., Wang, Z., Fan, Y., Liu, X., Wang, Z., Chang, S., Huang, T.: Robust video super-resolution with learned temporal dynamics. In: *Proceedings of the IEEE International Conference on Computer Vision* (2017) [3](#)
17. Liu, Y.L., Liao, Y.T., Lin, Y.Y., Chuang, Y.Y.: Deep video frame interpolation using cyclic frame generation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 8794–8802 (2019) [4](#)
18. Liu, Z., Yeh, R.A., Tang, X., Liu, Y., Agarwala, A.: Video frame synthesis using deep voxel flow. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4463–4471 (2017) [4](#)
19. Nah, S., Baik, S., Hong, S., Moon, G., Son, S., Timofte, R., Mu Lee, K.: Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pp. 0–0 (2019) [10](#)
20. Niklaus, S., Liu, F.: Context-aware synthesis for video frame interpolation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1701–1710 (2018) [2](#), [4](#)
21. Niklaus, S., Mai, L., Liu, F.: Video frame interpolation via adaptive convolution. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 670–679 (2017) [3](#)
22. Niklaus, S., Mai, L., Liu, F.: Video frame interpolation via adaptive separable convolution. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 261–270 (2017) [3](#)
23. Sajjadi, M.S., Vemulapalli, R., Brown, M.: Frame-recurrent video super-resolution. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6626–6634 (2018) [2](#), [3](#)
24. Shahrar, O., Faktor, A., Irani, M.: Space-time super-resolution from a single video. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3353–3360 (2011) [4](#)
25. Sharma, M., Chaudhury, S., Lall, B.: Space-time super-resolution using deep learning based framework. In: *International Conference on Pattern Recognition and Machine Intelligence*. pp. 582–590. Springer (2017) [4](#)
26. Shechtman, E., Caspi, Y., Irani, M.: Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (4), 531–545 (2005) [4](#)
27. Shi, W., Caballero, J., Huszar, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1874–1883 (2016) [6](#), [8](#)
28. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8934–8943 (2018) [7](#), [11](#)
29. Tao, X., Gao, H., Liao, R., Wang, J., Jia, J.: Detail-revealing deep video super-resolution. In: *Proceedings of the IEEE International Conference on Computer Vision* (2017) [2](#), [3](#)
30. Tommasi, T., Patricia, N., Caputo, B., Tuytelaars, T.: A deeper look at dataset bias. In: *Domain adaptation in computer vision applications*, pp. 37–55. Springer (2017) [10](#)
31. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7794–7803 (2018) [13](#)

- 32. Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: Edvr: Video restoration with enhanced deformable convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019) [2](#), [3](#), [7](#), [10](#)
- 33. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: Esrgan: Enhanced super-resolution generative adversarial networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 0–0 (2018) [8](#)
- 34. Xue, T., Chen, B., Wu, J., Wei, D., Freeman, W.T.: Video enhancement with task-oriented flow. arXiv (2017) [2](#), [10](#)
- 35. Xue, T., Chen, B., Wu, J., Wei, D., Freeman, W.T.: Video enhancement with task-oriented flow. International Journal of Computer Vision **127**(8), 1106–1125 (2019) [4](#), [9](#)