

Two-Stream Action Recognition-Oriented Video Super-Resolution

Haochen Zhang, Dong Liu, Zhiwei Xiong

CAS Key Laboratory of Technology in Geo-Spatial Information Processing and Application System,
University of Science and Technology of China, Hefei 230027, China

zhc12345@mail.ustc.edu.cn, {dongeliu, zwxiong}@ustc.edu.cn

Abstract

We study the video super-resolution (SR) problem for facilitating video analytics tasks, e.g. action recognition, instead of for visual quality. The popular action recognition methods based on convolutional networks, exemplified by two-stream networks, are not directly applicable on video of low spatial resolution. This can be remedied by performing video SR prior to recognition, which motivates us to improve the SR procedure for recognition accuracy. Tailored for two-stream action recognition networks, we propose two video SR methods for the spatial and temporal streams respectively. On the one hand, we observe that regions with action are more important to recognition, and we propose an optical-flow guided weighted mean-squared-error loss for our spatial-oriented SR (SoSR) network to emphasize the reconstruction of moving objects. On the other hand, we observe that existing video SR methods incur temporal discontinuity between frames, which also worsens the recognition accuracy, and we propose a siamese network for our temporal-oriented SR (ToSR) training that emphasizes the temporal continuity between consecutive frames. We perform experiments using two state-of-the-art action recognition networks and two well-known datasets—UCF101 and HMDB51. Results demonstrate the effectiveness of our proposed SoSR and ToSR in improving recognition accuracy.

1. Introduction

In recent years, convolutional neural networks (CNNs) have been applied to action recognition task and obtained state-of-the-art performance over the traditional arts. For the convenience of classification, most of them adopt fully-connected layers in their architecture and thus these well-trained CNNs cannot be directly applied on low-resolution

(LR) video. As a result, their widely application could be hindered by video's resolution. Most of the datasets used for studying action recognition have a fixed resolution, e.g. UCF101 (about 320×240), HMDB51 (about 340×256), Sports-1M (about 640×360) [10]. But the resolution in real world usually varies among different sources of video capturing, inevitably being low, e.g. in surveillance scenario. There are also many situations where the video has high resolution but the region containing action is quite small.

To address the resolution problem, most of existing works choose interpolation to simply re-scale the input while some recent works in other areas propose super-resolution (SR) as an alternative solution. For example, [26] investigated the effects of SR on object detection and [37] proposed a dataset for assessing the impact of image restoration and enhancement on image classification. Most of them stop at a preliminary experimental study on existing SR methods without proposing approaches targeting on their tasks, e.g. action recognition.

Super-resolution, aiming to enhance the resolution of images or video, has long attracted the attention of researchers. In early years, the target of enhancement is mainly signal fidelity, e.g. PSNR, partly because of an intuitive assumption that PSNR is consistent with visual quality, and mean-squared-error (MSE) is extensively used during optimization. However, some recent works challenge this assumption by showing distortion and perception can be tradeoff [1], and introduce some perceptual loss [15, 20] in addition to MSE, targeting for better visual quality of super-resolved image. Nonetheless, it is still not clear whether visual quality determines the quality of visual analytics results, e.g. action recognition accuracy. Since the analytics tasks are performed by computer instead of human, we argue that the SR methods optimized for visual quality may not be optimal for action recognition task.

We investigate the video SR problem aiming to facilitate recognition quality, exemplified by action recognition, instead of visual quality. In particular, we use SR as a preprocessing step before feeding LR video into an action recognition network that is well-trained on HR video. We inves-

This work was supported by the National Key Research and Development Plan under Grant 2017YFB1002401, and by the Natural Science Foundation of China under Grant 61772483. (Corresponding author: Dong Liu.)

tigate recognition quality of different super-resolved video evaluated by a computer algorithm rather than by human. The problem we want to address is that given a data-driven classifier, exemplified by an action recognition CNN, with parameters trained on HR video, what the accuracy will be when the same classifier deals with LR video assisted with SR preprocessing. In addition, we want to find how to design a better SR network targeting on recognition to improve the accuracy.

Oriented to the popular two-stream action recognition framework [28] which learns two separate networks, one for spatial color information and the other for temporal motion information, we propose two video SR methods for these two streams respectively. For the spatial stream which can be regarded as image classification, we observe that the moving object is more related to the recognition and should be paid more attention during SR enhancement. Thus, our Spatial-oriented SR (SoSR) takes weighted mean-squared-error guided by optical flow as loss to emphasize moving objects. For the temporal stream, we observe that video SR can result in the temporal discontinuity between consecutive video frames which may harm the quality of optical flow and incur drop in recognition accuracy. Thus, in our Temporal-oriented SR (ToSR), we enhance the consecutive frames together to ensure the temporal consistency.

Our contributions can be summarized as follows. We investigate state-of-the-art image and video SR methods from the view of facilitating action recognition, assuming well-trained two-stream networks as “evaluators.” For the spatial stream, we propose an optical flow guided weighted MSE loss to guide our SoSR to pay more attention to regions with motion. For the temporal stream, we propose ToSR which enhances the consecutive frames together to achieve temporal consistency.

We perform experiments with two state-of-the-art recognition networks on two widely used datasets—UCF101 [29] and HMDB51 [19]. Comprehensive experimental results show that our SoSR and ToSR indeed improve the recognition accuracy significantly. Especially on the HMDB51 dataset, our proposed method can improve the recognition performance of LR video from 42.81% to 53.59% on spatial stream and from 56.54% to 61.5% on temporal stream. Our code is released¹.

2. Related Work

We review related works at two aspects: action recognition and image/video SR. In both fields, CNN has been the mainstream and outperforms the traditional methods significantly. Thus we only mention several CNN-based approaches that are highly related to our work.

CNN for action recognition. In CNN-based action

recognition, a key problem is how to properly incorporate spatial and temporal information in CNN architectures. Solutions can be divided into three categories: 3D convolution, RNN/LSTM, and two-stream. 3D CNN which learns spatio-temporal features was first presented in [13]. Later on, C3D features and 3D CNN architectures [4, 35, 36, 39] appeared. There were also several works [25, 31, 45] focusing on improvements of 3D CNNs. RNN/LSTM is believed to cope with sequential information better, and thus [5, 41, 42] attempted to incorporate LSTMs to deal with action recognition. Two-stream CNN architecture was firstly proposed in [28]. This architecture consists of two separate networks, one for exploiting spatial information from individual frames, and the other for using temporal information from optical flow; the outputs of two networks are then combined by late fusion. Several improvements were presented for two-stream [7, 8, 38]. We design SR methods specifically for two-stream networks due to two reasons. First, two-stream approach seems leading to the best performance for action recognition on several benchmarks. Second, both 3D convolution and RNN/LSTM networks are not easily decomposed, but two-stream networks have a clear decomposition, which facilitates the investigation of SR. We use two state-of-the-art methods, Temporal Segment Network (TSN) [38] and Spatio-Temporal Residual Network (ST-Resnet) [7], in our experiments.

CNN for image SR. Almost all of the existing image SR methods are designed to enhance the visual quality by adding more image details. In earlier years, PSNR is evaluated as a surrogate of visual quality and thus mean-squared-error is extensively used as loss function [6, 17, 18, 21, 22, 27, 32, 33, 44]. More recently, visual quality is considered directly and several different kinds of loss functions are proposed, such as feature loss [15] and loss defined by generative adversarial network (GAN) [9]. For example, Ledig *et al.* [20] proposed SRGAN which combined GAN loss and feature loss. It is also worth noting that PSNR and visual quality can be even contradictory [1].

CNN for video SR. Compared to single image SR, the temporal dimension provides much more information in video SR, and various methods have been proposed to exploit the temporal information [23, 24]. A majority of these methods have an explicit motion compensation module to align different frames. For example, Kappeler *et al.* [16] slightly modified SRCNN [6] and extracted features from frames that were aligned by optical flow. Caballero *et al.* [2] proposed an end-to-end SR network to learn motions between input LR frames and generate SR frames in real time. Tao *et al.* [34] introduced a new sub-pixel motion compensation (SPMC) layer to perform motion compensation and up-sampling jointly. Also several methods try to avoid the explicit motion compensation. For example, Jo *et al.* [14] proposed a network that used dynamic up-sampling

¹<https://github.com/AlanZhang1995/TwoStreamSR>

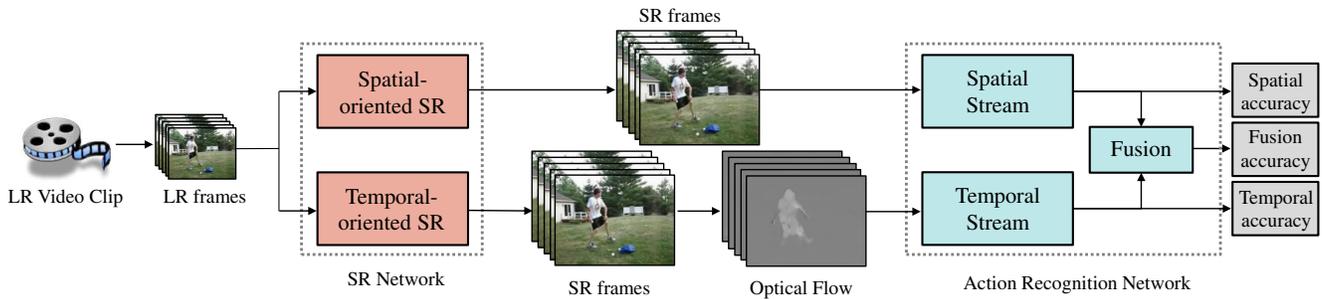


Figure 1. The pipeline of performing video SR prior to action recognition for LR video. Note that our work focuses on the SR network, we directly adopt the well-trained two-stream action recognition network without any tuning.

filters. All the aforementioned works are pursuing higher PSNR for video SR. But we consider video SR to improve action recognition accuracy. We focus on the loss functions instead of the network structures.

3. Action Recognition-Oriented SR

Pipeline. Figure 1 depicts the pipeline of using SR for action recognition by two-stream networks. Given an LR video sequence, we split it into frames on which we perform SR enhancement. We propose Spatial-oriented SR (SoSR) and Temporal-oriented SR (ToSR) for the two streams, *i.e.* we enhance the LR video twice. We then calculate optical flow from the ToSR resulting video, and feed the optical flow together with frames from the SoSR resulting video into the following recognition network.

Action Recognition Network. Our SR methods are specifically designed for two-stream action recognition networks. Specifically, we use TSN [38] and ST-Resnet [7] in our experiments. There are minor differences between the two networks: TSN uses a weighted average of the classification scores predicted from the two streams, while ST-Resnet trains a fusion sub-network together with the two streams in an end-to-end fashion. We focus on the SR part and directly use the well-trained models provided by the authors *without any tuning*.

End-to-End Optimization? According to Figure 1, it appears an appealing choice to perform an end-to-end optimization, *i.e.* training the SR network with the recognition accuracy as the objective. However, this involves a specific action recognition model. Our empirical results indicate that, training the SR with a specific action recognition model (e.g. TSN), and testing with another model (e.g. ST-Resnet), leads to much worse results. It motivates us to design specific loss functions for the SR training.

3.1. Spatial-oriented SR

3.1.1 Analysis

According to the two-stream architecture, the spatial stream performs recognition from individual frames by recogniz-

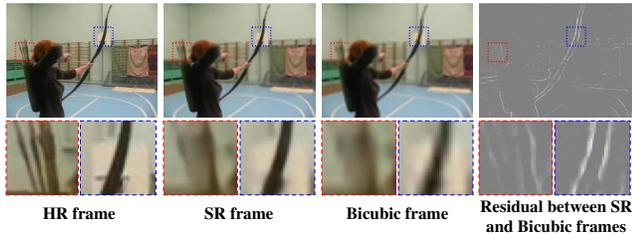
ing objects. That says, the spatial stream is equivalent to image classification. Inspired by previous work [3], we expect that SR can enhance the LR frames and add more image details with which SR helps in recognition. However, when we experiment with a representative image SR method, namely VDSR [17], we observe some counterexamples. We calculate recognition accuracy for individual classes, and find that VDSR sometimes performs worse than the simple bicubic interpolation, more interestingly, the original HR frames can be worse than super-resolved or even interpolated frames. Such examples are summarized in Table 1 and Figure 2. Considering that LR frames lose details compared to HR frames, bicubic interpolation simply up-scales frames without adding details, and SR methods usually enhance interpolated frames with image details, we conjecture that, especially in specific classes, image details can be either helpful or harmful for action recognition depending on the regions where details are added.

In Figure 2, we visually analyze some frames to confirm our conjecture. In (a), which corresponds to HR>VDSR>Bicubic, we indeed observe that many details about the bow and arrow lie in the HR frame, but are missing in bicubic frame; the SR frame adds some details on the bow (shown in the blue box), which is helpful for recognition since the bow is directly related to the class *Archery*. In (b), which corresponds to Bicubic>HR>VDSR, we observe that SR frame contains more details than bicubic frame, mostly on the background (shown in the blue box) rather than on the key object (shown in the red box); the added details seem to be harmful for action recognition. In (c), which corresponds to VDSR>HR>Bicubic, as the SR frame has more details on the human (object directly related to *Walking*) but fewer details on the background (due to LR input), the recognition accuracy is even higher than HR.

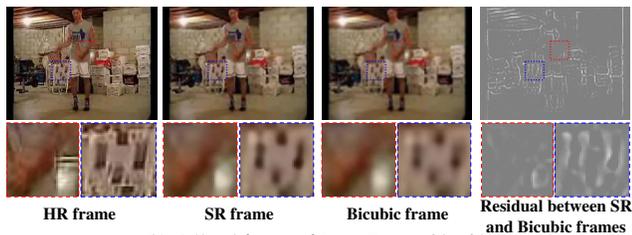
It is worth noting that, despite of these counterexamples, HR video is still the best in terms of recognition accuracy on the overall sense (as shown in Tables 4), partly because the action recognition networks are trained on HR frames. What if the networks are trained with LR video? We will study in Section 4.3.

Case	Class	Recognition Accuracy (%)		
		HR	Bicubic	VDSR
a	Archery	82.93	36.59	70.73
	PlayingFlute	97.92	72.92	79.17
b	JumpRope	39.47	42.11	7.89
	SalsaSpin	79.07	83.72	53.49
c	FrontCrawl	64.86	32.43	78.38
	HandstandWalking	35.29	29.41	41.18

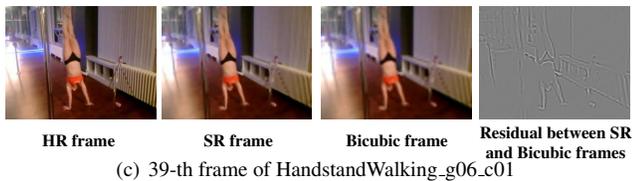
Table 1. We observe different cases in recognition accuracy for the classes in UCF101 using the TSN network. In case (a), HR>VDSR>Bicubic. In case (b), Bicubic>HR>VDSR. In case (c), VDSR>HR>Bicubic. Some representative classes are presented in this table. See Figure 2 for visual inspection. (The scaling factor is 4 throughout the entire paper unless otherwise indicated.)



(a) 10-th frame of Archery_g01_c07



(b) 152-nd frame of JumpRope_g02_c02



(c) 39-th frame of HandstandWalking_g06_c01

Figure 2. Examples show how image details added by VDSR [17] influence the recognition accuracy. In (a), VDSR adds details on the bow, and since bow directly relates to Archery, VDSR improves recognition than bicubic. In (b), VDSR adds details on the background but not on the object, resulting in even lower accuracy than bicubic. In (c), VDSR adds details on the walking woman but not on the background, resulting in even higher accuracy than HR. See Table 1 for the accuracy values.

3.1.2 Method

Based on the observation, we propose an SR method to selectively enhance the image regions that are highly related to action recognition. These regions usually have high motion, such as the bow in Figure 2 (a), the rotating rope in Figure 2 (b), and the walking woman in Figure 2 (c). We

Structure	MSE/W MSE	Feature	Adversarial	Accuracy
VDSR	MSE	-	-	46.6%
VDSR	WMSE	-	-	47.91%
VDSR	WMSE	✓	-	50.39%
ESRGAN	WMSE	✓	-	52.55%
ESRGAN	MSE	✓	✓	52.48%
ESRGAN	WMSE	✓	✓	53.59%

Table 2. Ablation study for SoSR using different network structures and different loss functions, with TSN [38] on HMDB51 dataset.

select these regions according to the optical flow since optical flow is a commonly chosen representation for motion information. Note that high motion does not necessarily relate to action. It is a much simplified implementation, but seems working well in our experiments.

Most of SR networks use mean-squared-error (MSE) as loss function, which is to assume equal importance of every pixel. In contrast, we propose to use a weighted MSE (WMSE) based on optical flow to emphasize some pixels that are more important than others. In short, the loss function we used here is

$$WMSE = \frac{1}{N} \sum_{p=1}^N \left\| I(p) - \hat{I}(p) \right\|^2 \cdot \sqrt{u^2(p) + v^2(p)}, \quad (1)$$

where I and \hat{I} are HR and SR frames respectively, p is the pixel index, and N is the total number of pixels. u and v represent the magnitude of optical flow in the horizontal and vertical directions respectively. Here, the optical flow is calculated offline from the HR video using Flownet 2.0 [12], which we observe is slightly better than using TVL¹ [43]. In this way, the loss can guide the network in a pixel-wise manner: pixels with larger motion correspond to larger loss weights and thus are paid more attention during SR enhancement.

In addition to WMSE, we further investigate two kinds of perceptual loss: feature loss and adversarial loss, which have been widely used in recent SR methods for improving visual quality [15, 20]. Using feature loss is to minimize the difference of high-level image features between SR image and HR image and using adversarial loss is to generate SR image which is closer to HR image in terms of distribution.

As mentioned before, the spatial stream is equivalent to image classification. We anticipate that single frame SR can perform well for the spatial stream and also has lower complexity than multi-frame SR. So here, we investigate two image SR network structures: One is VDSR [17] structure and the other is based on ESRGAN [40].

We conduct an ablation study about the proposed loss function and different network structures. As shown in Table 2, feature loss, adversarial loss as well as advanced network structure are all beneficial to the final recognition accuracy, and our proposed WMSE further improves the

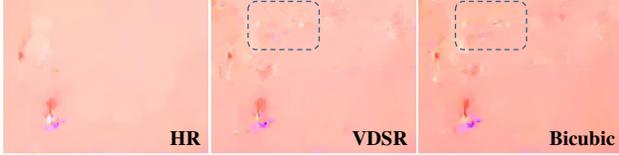


Figure 3. Optical flow maps calculated from HR, VDSR [17], and bicubic video, respectively. Artifacts can be found in the circled regions. Zooming-in inspection can observe that SR video has more artifacts than bicubic one. In this example, VDSR incurs lower recognition accuracy than bicubic on the temporal stream.

recognition performance. For more analyses please refer to Section 4.2.

As a result, we train our SoSR network based on ESRGAN with our training data (details in Section 4.1) and the following loss function:

$$\mathcal{L}_{\text{SoSR}} = \alpha \mathcal{L}_{\text{WMSE}} + \beta \mathcal{L}_{\text{Feature}} + \gamma \mathcal{L}_{\text{Adversarial}}, \quad (2)$$

where α , β and γ are weights.

3.2. Temporal-oriented SR

3.2.1 Analysis

We now switch to the temporal stream. As described in the two-stream architecture, the temporal stream takes optical flow as input to utilize temporal information. Thus, the question should be how SR affects the quality of optical flow. We again experiment with the representative image SR method—VDSR [17]. Figure 3 shows the optical flow maps calculated from HR video, SR video, and bicubic interpolated video, respectively. Here the optical flow is calculated by the TVL¹ method [43] as most action recognition works do. From Figure 3, we can find that the optical flow from bicubic video has a lot of artifacts, and VDSR even worsens the optical flow. Thus, VDSR incurs less appealing results of recognition accuracy.

The reason for above results should be attributed to VDSR being an image SR network that enhances video frames individually and causes temporal inconsistency. For high-quality optical flow, we need to ensure the temporal consistency between frames, which has also been studied in previous video SR works. For example, [2, 14] discussed the temporal consistency and its relation to visible flickering artifacts when displaying SR video. In Figure 4, we adopt the visualization method known as *temporal profiles* with which [2, 14] display the flickering artifacts. As seen, VDSR indeed incurs more temporal discontinuity.

3.2.2 Method

Through the above observation, we find a relation between optical flow-based recognition accuracy and the temporal

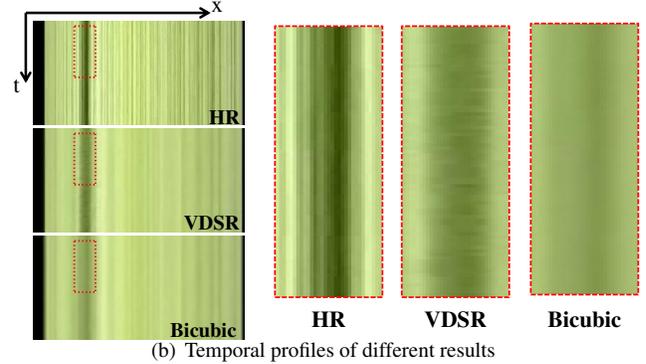
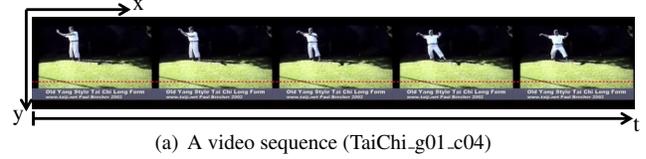


Figure 4. (a) An example video sequence. We sample one row at the same location (indicated by the red dot line) from each frame and concatenate the rows to produce (b) the temporal profiles. Obviously, bicubic video has the least image details, VDSR video has some details but displays temporal discontinuity that will cause flickering artifacts. In this example, VDSR incurs lower recognition accuracy than bicubic on the temporal stream.

consistency in the SR video. Since the existing video SR schemes usually perform SR frame by frame, they have difficulty in guaranteeing the consistency between SR frames. We consider a siamese network reconstructing the consecutive frames together for training video SR.

As our objective is to achieve high quality optical flow, it is straightforward to calculate the optical flow between SR frames and compare it with that between HR frames. However, this would require an optical flow estimation network to support end-to-end training. But recent optical flow networks [11, 12, 30] are too deep to be efficiently trained with the standard error back-propagation technology. We take a warping approach to estimate the temporal continuity.

The siamese network for training ToSR is shown in Figure 5. We use two copies of an SR network to enhance two consecutive frames respectively. First of all, we want to achieve SR frames with high quality, and use two MSE losses for the two frames respectively, *i.e.* $\mathcal{L}_{\text{SR}} = \|I_t - \hat{I}_t\|_F^2 + \|I_{t+1} - \hat{I}_{t+1}\|_F^2$. Moreover, we want to ensure the temporal continuity between SR frames. So we adopt the optical flow from HR video, which can be calculated beforehand, to perform warping between two SR frames. Let the optical flow be $F_{t \rightarrow t+1}$, we use the relation $\tilde{I}_t(p) = \hat{I}_{t+1}(p + F_{t \rightarrow t+1}(p))$ to warp the SR frame \hat{I}_{t+1} . Warping is implemented by bilinear interpolation that is free of parameters. The warped result \tilde{I}_t is compared against both SR and HR frames of the previous timestamp. Accordingly, we define two losses: $\mathcal{L}_{\text{warp-SR}} = \|\tilde{I}_t - \hat{I}_t\|_F^2$

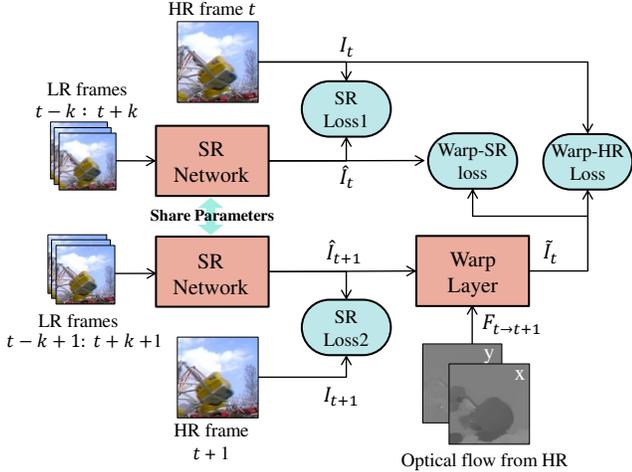


Figure 5. Our proposed ToSR uses a siamese network for training. We jointly consider two consecutive frames and design four loss terms to ensure the quality of individual frames as well as the temporal consistency between them.

$$\text{and } \mathcal{L}_{\text{warp-HR}} = \|I_t - \tilde{I}_t\|_F^2.$$

In summary, the loss function for ToSR is

$$\mathcal{L}_{\text{ToSR}} = \alpha \mathcal{L}_{\text{SR}} + \beta \mathcal{L}_{\text{warp-SR}} + \gamma \mathcal{L}_{\text{warp-HR}}, \quad (3)$$

where α, β, γ are weights.

Any existing image or video SR network can implement ToSR. We investigate two choices. The first is based on the VDSR network [17], which performs SR for frames individually. The second is based on the VSR-DUF with 16 layers [14], which utilizes multiple LR frames for SR.

The ablation study about the two network structures as well as the proposed loss function is reported in Table 3. The performance of VDSR network is limited by lack of information about adjacent frames while multi-frame SR network performs better on temporal stream. On different networks, our proposed warp loss could benefit to a large extent.

4. Experiments

4.1. Experimental settings

Datasets. We perform experiments using three datasets: one natural video dataset CDVL-134 for training SR network, and two datasets, UCF101 and HMDB51, for testing SR with action recognition networks. For the video SR task, there is no commonly used dataset. CDVL-134 is a dataset collected by ourselves from CDVL², and contains 134 natural video sequences with various content, including landscapes, animals, activities and so on. Because the resolution of these sequences varies from 480×360 to

²<https://www.cdvl.org/>

Structure	Warp loss	Accuracy
VDSR	-	55.1%
VDSR	✓	58.76%
VSR-DUF-16	-	59.48%
VSR-DUF-16	✓	61.5%

Table 3. Ablation study for ToSR using different network structures and different loss functions, with TSN [38] on HMDB51 dataset.

1920×1080 , we resize them to around 320×240 (similar to UCF101 and HMDB51) with bicubic interpolation while maintaining their aspect ratios. We further down-sample these resized sequences by a factor of 4 to generate LR video for training. As for UCF101 and HMDB51, they are popular action recognition datasets. The former dataset contains 13,320 video clips belonging to 101 action categories, and the latter is composed of 51 action categories and 6,766 video clips. Both datasets provide three training/testing splits and we here only use the first split as a representative. For more details, please refer to [29] and [19] respectively. In our experiment, we use a $4 \times$ down-sampled version of these two datasets as the input LR video for testing, and we use the original resolution of the two datasets (denoted by HR) as a reference.

Spatial-oriented SR. We randomly select HR frames from training video, and use FlowNet2.0 [12] on them to calculate optical flow which is then converted into weight maps. We crop HR frames and weight maps into 128×128 aligned patches and generate LR patches by bicubic interpolation. In particular, we select 120 frames from each video of CDVL-134 dataset and choose the top 10 crops with the largest area of motion. Manually excluding some obviously low-quality patches, there are totally 144,306 patches for training. We use the deep learning framework PyTorch to perform experiments and the optimization settings, such as learning rate and batch size, are recommended by [40]. The loss weight used in our SoSR training is $\alpha = 1, \beta = 1, \gamma = 0.005$.

Temporal-oriented SR. All the training samples are prepared similarly as for SoSR, except that TVL¹ [43] is applied on HR frames to calculate optical flow for warping (for a fair comparison, because the recognition networks use TVL¹), and we have 143,250 patch pairs for training, and 10,386 pairs for validation. Our ToSR is implemented on TensorFlow. The initial learning rate is 0.01 and multiplied by 0.1 every 10 epochs as recommended in [14]. We use batch size 16 and fine-tune from the model provided by the authors of [14]. As for loss weights, we have $\alpha = 1, \beta = 0.8, \gamma = 0.1$.

4.2. Recognition Results

All experimental results are obtained by different SR methods with the same scaling factor 4. Baseline methods

Method	HMDB51						UCF101					
	TSN			ST-Resnet			TSN			ST-Resnet		
	Spatial	Temporal	Fusion	Spatial	Temporal	Fusion	Spatial	Temporal	Fusion	Spatial	Temporal	Fusion
Bicubic	42.81	56.54	63.53	43.59	53.76	59.48	71.25	81.08	87.87	72.01	78.28	84.62
VDSR [17]	46.6	55.1	63.59	49.18	54.44	60.2	67.09	79.81	86.84	72.27	79.43	84.48
RCAN [44]	48.76	56.8	66.21	51.76	55.72	62.61	67.18	82.12	88	72.23	80.52	85.01
SRGAN [20]	48.82	49.87	63.01	51.41	47.22	60.85	81.33	75.45	87.55	83.31	70.16	86.97
ESRGAN [40]	52.48	51.5	63.4	53.79	49.72	61.83	82.97	75.32	87.75	83.81	70.64	86.62
SoSR	53.59	50.26	64.51	54.77	48.27	63.01	83.11	74.1	86.63	83.92	69.68	85.77
SPMC [34]	48.95	56.41	64.31	53.14	53.53	63.66	70.42	80.19	87.15	74.45	77.44	84.09
VSR-DUF-16 [14]	48.37	59.48	66.08	50.62	55.07	61.11	68.56	84.89	89.36	72.11	80.06	83.9
VSR-DUF-52 [14]	48.5	60.52	66.86	52.84	57.61	65.23	70.54	85.09	89.85	74.49	80.16	84.88
ToSR	47.45	61.5	66.08	51.54	58.92	64.77	64.79	85.29	88.46	70.88	81.07	83.82
SoSR+ToSR	/	/	68.3	/	/	67.32	/	/	92.13	/	/	90.19
HR	54.58	62.16	69.28	56.01	59.41	68.1	86.02	87.63	93.49	88.01	85.71	92.94

Table 4. Recognition accuracy (%) of $4\times$ super-resolved video from UCF101 and HMDB51 dataset using two action recognition network, TSN and ST-Resnet. Number of VSR-DUF [14] indicates number of layers. Accuracy of HR video is provided for reference. (Please refer to the supplementary material for PSNR and SSIM results of different methods.)

are four single image SR methods: VDSR [17], RCAN [44], SRGAN [20], ESRGAN [40], and two video SR methods: SPMC [34] and VSR-DUF [14]. TSN [38] and ST-Resnet [7] are used to obtain the recognition accuracy, shown in Table 4.

On spatial stream, firstly, comparing VDSR and RCAN, SRGAN and ESRGAN respectively shows that advanced design of network structure could benefit recognition quality of super-resolved video. This result is intuitive because more advanced SR method would generate SR frames with more details and be more helpful to recognition *on average*. Secondly, by comparing VDSR/RCAN and SRGAN/ESRGAN, the former methods optimize MSE only while the latter methods take use of perceptual loss, we can see the perceptual loss could also improve the recognition performance to some extent. Thirdly, our SoSR achieves the highest recognition accuracy even outperforming ESRGAN that is believed to achieve the best perceptual index [40]; this demonstrates that the SR methods optimized for visual quality are not optimal for action recognition task. For visual analyses please refer to Section 4.3.

Switching to temporal stream, where it is obvious that SRGAN/ESRGAN perform worse than VDSR/RCAN among single image SR methods. This difference should result from the perceptual loss, as Ledig *et al.* explained in [20]: MSE-based result is the pixel-wise average of possible results in pixel space, while GAN drives the reconstruction towards the natural image manifold. Accordingly, MSE-based result has better temporal consistency between adjacent SR frames. Among multi-frame SR methods, VSR-DUF outperforms SPMC significantly, even VSR-DUF-16 beats SPMC with a large gap. This difference may be attributed to the design of network structure. SPMC performs explicit warping with optical flow estimated from LR frames, which may introduce errors and undermine tempo-

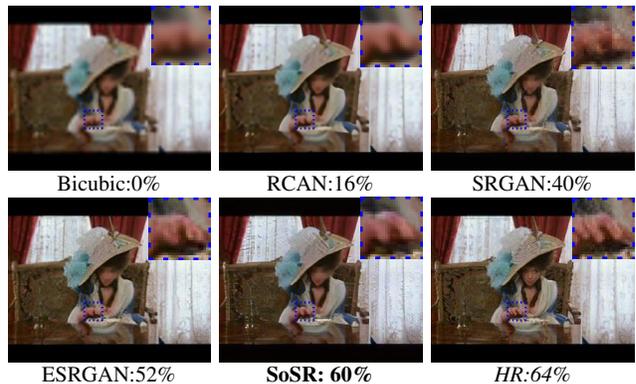


Figure 6. Visual quality comparison, and the numbers indicate video-level recognition accuracy using TSN. We observe from the SoSR result that the woman’s hand looks over-smooth but the background texture appears sharp. This is due to the interaction between WMSE and perceptual losses.

ral consistency, while VSR-DUF uses 3D convolution directly operating on consecutive LR frames to predict dynamic filters that are then used to up-sample the central LR frame. The structure without explicit motion compensation may be the key for VSR-DUF to achieve good performance. Last, owing to the proposed siamese network, the performance of our ToSR is even better.

Combining the two streams, SoSR plus ToSR gives out the highest accuracy. Results of other combinations are provided in the supplementary material.

4.3. Analyses

Figure 6 shows an example for comparing the visual quality as well as recognition accuracy of different SR methods. By comparing RCAN, SRGAN, ESRGAN results with the HR frame, we can find the recognition accuracy in-

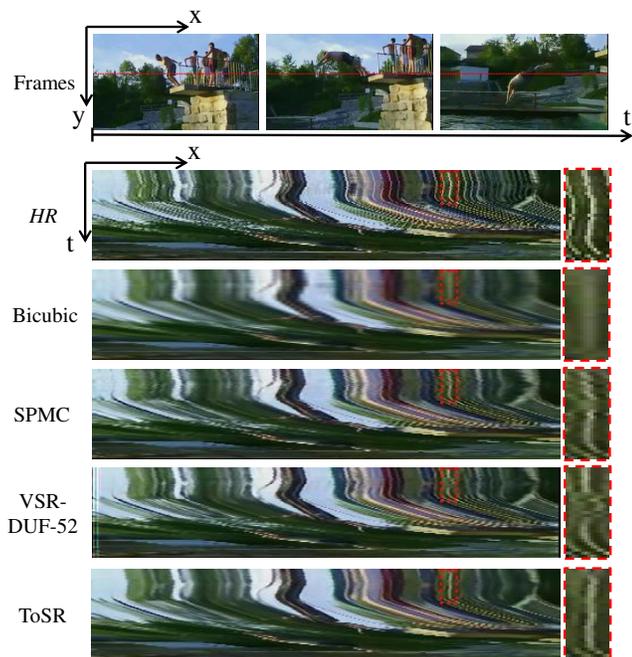


Figure 7. Temporal profiles of different SR results. Bicubic failed to produce vivid image details while SPMC and VSR-DUF incur obvious temporal discontinuity.

creases as the visual quality improves. When adding SoSR result into consideration, we observe that its visual quality is not consistent at different regions. As shown by the inset, the woman’s hand is over-smooth, but the background texture appears quite sharp. This is due to the interaction between WMSE and perceptual loss: WMSE emphasizes the MSE loss on the regions with large motion, and the MSE loss leads to over-smooth result as claimed in [20]; on the regions with little motion (*e.g.* the background), perceptual loss dominates the optimization target, which produces vivid but not true texture. In addition, we also observed cases where visual quality and recognition accuracy are not consistent, please refer to the supplementary material.

Figure 7 shows temporal profiles of video obtained by different SR methods, from which we can find bicubic interpolated frames do not have enough image details, while SPMC and VSR-DUF results look sharp. However, both of them incur severe temporal discontinuity. Our ToSR produces the best temporal profile. For more visual results about temporal profiles as well as artifacts in optical flow, please refer to the supplementary material.

In all of the previous experiments, we adopt action recognition network pretrained with HR video and evaluate different SR methods. One may question about this setting and consider whether it would be different if the action recognition network is trained *with video of different resolutions*. We investigate the joint influence of data augmentation when training recognition network and SR preprocess-

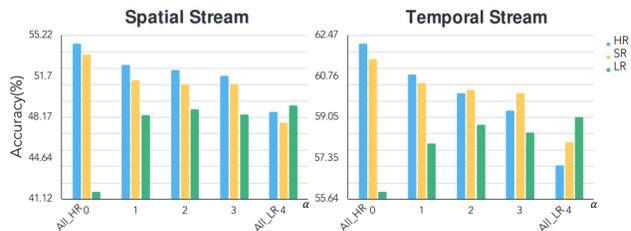


Figure 8. Recognition accuracy of models trained with different data augmentation configurations (denoted by α) and tested on HR, SR, LR video respectively.

ing when using the trained recognition network.

In our experiment, we train several TSN models using mixed HR and LR video from the HMDB51 dataset. Here the LR video is generated by $4\times$ down-sampling. The LR and HR video sequences are mixed with a ratio $\alpha : (4 - \alpha)$, *e.g.* $\alpha = 0$ means HR only and $\alpha = 4$ means LR only. Then we test the recognition performance of each model on HR video, LR video and super-resolved video (using SoSR and ToSR) respectively and report the results in Figure 8.

Firstly, using data augmentation can make the network pay more attention to the features shared by HR and LR, and improve the performance on LR video. However, the large difference between HR and LR may cause the network to neglect useful features unique to HR, and incur a decline in HR performance. Secondly, the network trained on LR video only ($\alpha = 4$) performs the best on LR video input, where the accuracy on LR is even higher than that on HR and SR. But this network performs the worst on HR video input, and should not be a good choice in practice. Thirdly, excluding the LR only case, there are still cases where SR outperforms HR (on the temporal stream). Thus, we anticipate that a joint consideration of SR network and action recognition network may lead to even better performance, which will be our future work.

Training different models is a straightforward solution for different resolutions, but has several limitations. First, it needs to train and maintain multiple models that can be costly. Second, how to select the appropriate model to match the resolution for a given input video is a problem.

5. Conclusion

We consider the video SR problem not for visual quality, but for facilitating action recognition accuracy. Tailored for two-stream action recognition networks, we propose SoSR with optical flow guided weighted MSE loss, and ToSR with a siamese network to emphasize temporal consistency. Experimental results demonstrate the advantages of our proposed SoSR and ToSR methods. In the future, we plan to combine SoSR and ToSR into a single step, and study the tradeoff between visual quality and recognition accuracy.

References

- [1] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *CVPR*, pages 6228–6237, 2018.
- [2] Jose Caballero, Christian Ledig, Andrew P Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *CVPR*, volume 1, pages 4778–4787, 2017.
- [3] Dengxin Dai, Yujian Wang, Yuhua Chen, and Luc Van Gool. Is image super-resolution helpful for other vision tasks? In *WACV*, pages 1–9, 2016.
- [4] Ali Diba, Vivek Sharma, and Luc Van Gool. Deep temporal linear encoding networks. In *CVPR*, pages 2329–2338, 2017.
- [5] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pages 2625–2634, 2015.
- [6] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016.
- [7] Christoph Feichtenhofer, Axel Pinz, and Richard Wildes. Spatiotemporal residual networks for video action recognition. In *NIPS*, pages 3468–3476, 2016.
- [8] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, pages 1933–1941, 2016.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.
- [10] Samitha Herath, Mehrtash Harandi, and Fatih Porikli. Going deeper into action recognition: A survey. *Image and Vision Computing*, 60:4–21, 2017.
- [11] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Lite-FlowNet: A lightweight convolutional neural network for optical flow estimation. In *CVPR*, pages 8981–8989, 2018.
- [12] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, volume 2, pages 2462–2470, 2017.
- [13] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3D convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, 2013.
- [14] Younghyun Jo, Seoung Wug Oh, Jaeyeon Kang, and Seon Joo Kim. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *CVPR*, pages 3224–3232, 2018.
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711, 2016.
- [16] Armin Kappeler, Seunghwan Yoo, Qiqin Dai, and Aggelos K Katsaggelos. Video super-resolution with convolutional neural networks. *IEEE Transactions on Computational Imaging*, 2(2):109–122, 2016.
- [17] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, pages 1646–1654, 2016.
- [18] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *CVPR*, pages 1637–1645, 2016.
- [19] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: A large video database for human motion recognition. In *ICCV*, pages 2556–2563, 2011.
- [20] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P Aitken, Alykhan Tejani, Johannes Totz, and Zehan Wang. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, pages 4681–4690, 2017.
- [21] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, number 2, pages 136–144, 2017.
- [22] Ding Liu and Thomas S Huang. Single image super-resolution: From sparse coding to deep learning. In *Deep Learning through Sparse and Low-Rank Modeling*, pages 47–86. Elsevier, 2019.
- [23] Ding Liu, Zhaowen Wang, Yuchen Fan, Xianming Liu, Zhangyang Wang, Shiyu Chang, and Thomas Huang. Robust video super-resolution with learned temporal dynamics. In *ICCV*, pages 2507–2515, 2017.
- [24] Ding Liu, Zhaowen Wang, Yuchen Fan, Xianming Liu, Zhangyang Wang, Shiyu Chang, Xinchao Wang, and Thomas S Huang. Learning temporal dynamics for video super-resolution: A deep learning approach. *IEEE Transactions on Image Processing*, 27(7):3432–3445, 2018.
- [25] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3D residual networks. In *ICCV*, pages 5534–5542, 2017.
- [26] Jacob Shermeyer and Adam Van Etten. The effects of super-resolution on object detection performance in satellite imagery. *arXiv preprint arXiv:1812.04098*, 2018.
- [27] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *CVPR*, pages 1874–1883, 2016.
- [28] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, pages 568–576, 2014.
- [29] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [30] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, pages 8934–8943, 2018.
- [31] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *ICCV*, pages 4597–4605, 2015.
- [32] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *CVPR*, number 2, pages 3147–3155, 2017.

- [33] Ying Tai, Jian Yang, Xiaoming Liu, and Chunyan Xu. MemNet: A persistent memory network for image restoration. In *CVPR*, pages 4539–4547, 2017.
- [34] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In *ICCV*, pages 22–29, 2017.
- [35] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [36] Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1510–1517, 2018.
- [37] Rosaura G Vidal, Sreya Banerjee, Klemen Grm, Vitomir Struc, and Walter J Scheirer. UG2: A video benchmark for assessing the impact of image restoration and enhancement on automatic visual recognition. In *WACV*, pages 1597–1606, 2018.
- [38] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36, 2016.
- [39] Xuanhan Wang, Lianli Gao, Peng Wang, Xiaoshuai Sun, and Xianglong Liu. Two-stream 3D convnet fusion for action recognition in videos with arbitrary size and length. *IEEE Transactions on Multimedia*, 20(3):634–644, 2018.
- [40] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. ESRGAN: Enhanced super-resolution generative adversarial networks. In *ECCVW*, pages 63–79, 2018.
- [41] Zuxuan Wu, Xi Wang, Yu-Gang Jiang, Hao Ye, and Xiangyang Xue. Modeling spatial-temporal clues in a hybrid deep learning framework for video classification. In *ACM MM*, pages 461–470, 2015.
- [42] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, pages 4694–4702, 2015.
- [43] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime TV-L¹ optical flow. In *Joint Pattern Recognition Symposium*, pages 214–223, 2007.
- [44] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, pages 1–16, 2018.
- [45] Yizhou Zhou, Xiaoyan Sun, Zheng-Jun Zha, and Wenjun Zeng. MiCT: Mixed 3D/2D convolutional tube for human action recognition. In *CVPR*, pages 449–458, 2018.