Super-Resolving Compressed Video in Coding Chain

Dewang Hou[1]   Yang Zhao[2,3]   Yuyao Ye[1]   Jiayu Yang[1]

Jian Zhang[1,3]   Ronggang Wang[1,3]

[1]Peking University   [2] Hefei University of Technology   [3]Peng Cheng Laboratory

Figure 1: A comparison of visual quality on low-resolution compressed video restoration (reconstructing 1080p results from 270p compressed videos). RCAN [49], EDVR [41], SRNTT [51] are the state-of-the-art methods for single image, video and reference-based super-resolution, respectively. In this paper, bicubic, RCAN and EDVR are used in the multi-resolution coding pipeline; SRNTT and the proposed R3N are used in the mixed-resolution coding pipeline.

Abstract

Scaling and lossy coding are widely used in video transmission and storage. Previous methods for enhancing the resolution of such videos often ignore the inherent interference between resolution loss and compression artifacts, which compromises perceptual video quality. To address this problem, we present a mixed-resolution coding framework, which cooperates with a reference-based DCNN. In this novel coding chain, the reference-based DCNN learns the direct mapping from low-resolution (LR) compressed video to their high-resolution (HR) clean version at the decoder side. We further improve reconstruction quality by devising an efficient deformable alignment module with receptive field block to handle various motion distances and introducing a disentangled loss that helps networks distinguish the artifact patterns from texture. Extensive experiments demonstrate the effectiveness of proposed innovations by comparing with state-of-the-art single image, video and reference-based restoration methods.

1. Introduction

Nowadays, video commands the lion's share of Internet traffic and is still climbing. On the one hand, it is continuously desirable to develop coding algorithms with a higher

compression ratio [37] and lower complexity. On the other hand, technically, a post-processing step of restoring videos at the decoder side has more room for improvement compared with the standards.

To achieve optimal video quality under bandwidth and power constraints, one effective way is to down-sample the video before compression and transmission. Therefore, resolution loss and compression artifacts are the most typical examples of degradation in practical applications. For the task of recovering a clean video from its degraded version at the decoder side, we call it restoration strategy decompression. It is a highly ill-posed problem as there exists infinite feasible solutions.

Fortunately, although super-resolution (SR) and compression artifact removal (CAR) are long-standing tasks, they are progressing rapidly driven by machine learning advances. Especially, convolutional neural networks [8, 9] have been successfully utilized for low-level vision tasks. However, the majority of previous research has focused on single well-defined problem. On the contrary, joint SR-CAR is an intricate task, where high-frequency details are restored for SR and high-frequency artifact are removed for CAR. We will elaborate it in Sec. 2.4.

Thus, when it comes to practical applications for joint SR-CAR, the difficulty is increased and less progress has been made. To make a step forward, we rethink the coding framework [3, 28, 31] and post-processing from

1

a holistic perspective. The key innovation is a novel paradigm that can leverage synergies in coding and post-processing. Specifically, we propose a modified framework, Mixed-Resolution Coding (MixedRC). The term "mixed-resolution" in this paper refers to dual bitstream, one of which contains high-frequency information from full-resolution key-frames and the other contains information from reduced-resolution frames. And we design a neural network architechture, called Reference-based Restoration Network with Refined-offset deformable alignment (R3N), to relieve performance bottlenecks of ill-posed problems. A trained R3N can efficiently transfer the textures from key-frames to LR compressed frames and achieve highly attractive performance. We illustrate comparison examples in Fig. 1. Last but not least, we improve reconstruction quality by introducing a disentangled loss that classifies pixels into texture and artifact.

In summary, the main contributions are three-fold:

- We propose MixedRC, which can be seamlessly incorporated into existing (and future) single-layer codecs to enhance their performances with minimal effort. This study also breaks new ground for reference-based networks in post-processing.

- We define a disentangled loss that helps the network distinguish the artifact patterns from texture by comparing statistics collected over the entire image. Ablation experiments further verify the effectiveness of the proposed innovations.

- We propose a reference-based network, called R3N. The core of R3N is the refined-offset deformable alignment, which is efficient to implement and is easy to optimize. Receptive field block and spatial attention in this module can improve performance by "calibrating" feature responses, which is a far more flexible alternative to handle various motion distances. And we reveal that a easy-to-hard transfer, setting on frame rate, is helpful in learning a reference-based model.

2. Related Works

2.1. Multi-Resolution Coding.

It is known that, at low bit rates, a down-sampled video visually beats the HR video when represented with the same number of bits via compression. In 2003, Bruckstein *et al.* gave a numerical analysis of the down-sampling and compression process [5]. After that, studies [10, 11, 33] combined video coding and SR techniques to improve coding efficiency.

Among them, Multi-Resolution Coding (MultiRC) [16] adaptively changes frame sizes for artifact and complexity reduction. However, the limited quality of frames, with no exception, narrows down the scope of this framework and becomes harder to super-resolve frames with artifact removal jointly, in terms of expressiveness and robustness [46, 47]. The rudiment of MixedRC is also proposed in that decade [4]. However, limited by the performance of SR at the time, these pioneering works have not been well developed. To advance this line, we propose a MixedRC based on scalable system [28, 31]. Both MultiRC and the proposed MixedRC are elaborated in Fig. 2.

2.2. Super-Resolution

Loss functions for SR. Sophisticated SR algorithms often focused on minimizing pixel-wise reconstruction errors to achieve a high peak signal-to-noise ratio (PSNR), which correlate poorly with image quality as perceived by a human observer [43]. Parallel efforts also studied adopting perceptually-motivated losses [19, 27, 52] to avoid the pixel-wise average problem, which typically leads to over-smoothed results. Orthogonal to these are adversarial losses [21, 42], pushing the super-resolved image to be of high likelihood given examples from HR domains.

Reference-based SR. The lost information during down-sampling an HR image amounts to high-frequency components [32]. Intuitively, to restore such high-frequency components, providing similar rich textures is a more reasonable approach than generating unreal textures. Early manually-designed filters [13, 14] use an external image as guidance to adjust filter parameters, which can preserve sharp edges. These works have seminal significance on reference-based SR (Ref-SR) [35,51,53], which aims to super-resolve an LR image with the help of HR reference. The main challenge lies in the design of the alignment module: CrossNet [53] performs alignment based on optical flow estimations and warping operations. However, flow estimations [17] need additional supervision and are highly vulnerable to large motions. Later, SRNTT [51] relaxes constraint on content similarity levels by adopting brute-force patch matching at multi-scale feature space. Along this line of thought, Zhang *et al.* [50] achieved significantly improved results on a challenging task of super-resolving painting images. Shim *et al.* [35] pointed out that deformable convolution [7] is a preferred mechanism to perform implicit motion compensation. It is worth noting that the difficulty of alignment in the Ref-SR is more incredible than adjacent frame alignment in video SR [38, 41]. To this end, Shim *et al.* use the stacking of deformable convolution layers to sample more locations with a larger receptive field. However, the deformable convolution is usually tricky to train, let alone stacking multiple layers sequentially. Moreover, stacking layers "plainly" could be suboptimal for sampling locations at long distances.
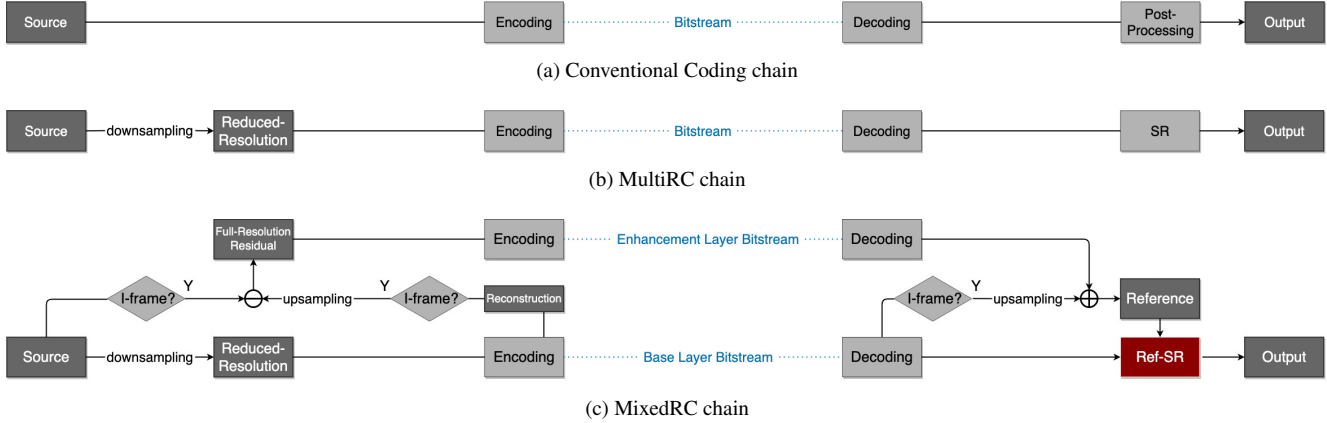
Figure 2: The scheme of mentioned coding frameworks. In MixedRC, frames scaled to different resolutions are encoded into layered bitstreams.

2.3. Compression Artifact Removal

CAR is also a long-standing low-level vision task. Compared to image CAR, learning-based video CAR appears as a mostly under-explored field. There is less complementary information between adjacent compressed frames due to the nature of coding algorithms, that quantization is enforced on the prediction error during inter-frame encoding. Moreover, the positions where blocking boundaries appear are relatively changeless in adjacent frames. To circumvent this problem, methods tend to extract temporal information that lies in the dependencies between the current frame and some valuable frames, such as peak quality frames [12] and previous restored frames [26]. While in our algorithm, these valuable frames are specifically Ref frames.

2.4. Restoration of Multiple Degradations

Scaling and compression cause the dominant degradations during video transmission. A network that can juggle SR and CAR tasks jointly is far more desirable than cascading different networks.

Recent studies [20, 24, 29, 36, 44] show the advantages of the joint-learning on multiple-degradation problem. Yu *et al.* [46] observed that the restoration of multiple degradations is not a simple composition of corresponding restorers trained on specific tasks. Unlike the rising interest in other joint-learning studies, the research devoted to restore the quality and practicality of LR compressed video is rarely mentioned. More relevant to this work, Zhang *et al.* [47] proposed SRMD, a single network to perform SR and denoising jointly. These couple of tasks would interfere with each other, which leads to visually unpleasant results. More specifically, directly super-resolving the noisy input will exacerbate the unwanted noise, rendering them visually objectionable. And simply pre-denoising also tends to lose details, result in worse SR performance. SRMD can achieve satisfactory performance if the predicted degradation maps are close to the ground truth. However, it is still not applicable in real applications as the blur kernel and noise level cannot be predicted for every image on hand.

3. Methodology

We first introduce the proposed MixedRC, a new coding chain based on scalable system. Then for joint SR-CAR problem, we carry on simplification and improvement in a reference-based network architecture, called R3N. Finally, we discuss how disentangled loss function works to suit the remedy to the case of multiple-degradation problem.

3.1. Coding Framework

The overall pipeline of the proposed coding framework is shown in Fig. 2. Used by the proposed framework, spatial scalability [31] describes cases in which the base layer (BL) is used to encode a lower resolution signals of the video stream, and the enhancement layers (ELs) are used to encode high-frequency information in key-frames.

Another point is that video encoders divide up a video into sets of frames called group-of-pictures (GOP), specifies the order in which intra- and inter-frames are arranged. And in the proposed framework, only Intra-coded (I-) frames exist in the ELs. In addition to lower complexity, it further reduces the bit rate and the syntax overhead for enhancement layers. Finally, the I-frame decoded by layered coding in a GOP will be fed to a Ref-SR model as reference, to super-resolve other frames in the same GOP. And for simplicity, backward or bi-directional prediction is not used in this paper. MixedRC is not meant to be an alternative to existing codecs, but rather a useful complement to any codec.

3.2. Network Architecture

In the proposed coding framework, R3N works as a post-processing tool at the decoder side, which helps extract details lost in LR features but existed in Ref features. The
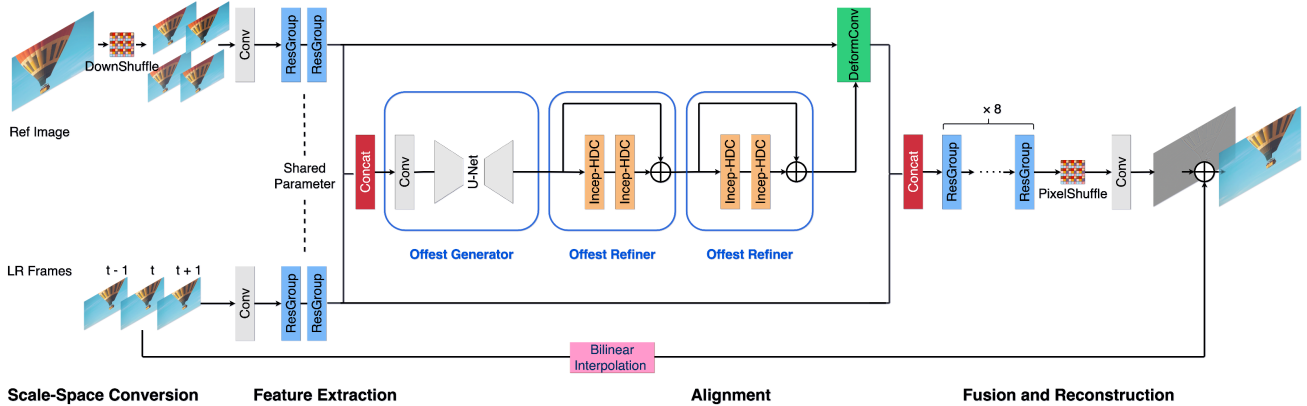
Figure 3: Network architecture of R3N. It is a pipeline that consists of scale-space conversion, feature extraction, alignment, fusion and reconstruction. we adopt the *ResGroups* proposed in RCAN [49] to the feature extraction and reconstruction parts. And for simplicity, we only show the architecture for three LR frames and one Ref image as input.

overall structure of R3N is shown in Fig. 3. We formulate reference-based video restoration as an integrative process of transferring plausible textures conditioned on Ref images I^{Ref}, and at last restoring the middle frame I_t^{LR} among $2N + 1$ consecutive LR compressed frames $I_{[t-N:t+N]}^{LR}$.

Next, we detail the individual components of the proposed R3N.

Down-Shuffle. Scale-conversion meant to make the contents of Ref image and LR frames be on the same scale. In contrast to the previous works, which usually use strided convolution or interpolation, we apply down-shuffle to accomplish scale-conversion. PixelShuffle [34] is proposed for image SR, and it is also the latest up-sampling scheme used in the state-of-the-art. We use its inverse process, *i.e.*, down-shuffle, to perform lightweight scale-space conversion. The down-shuffle operation systematically rearranges spatial pixels into channels, keeping the high-frequency amount intact, hence providing sufficient information for the following convolutional layers. Note that down-shuffle is parameter-free, and it can naturally handle geometric transformations through the combinations of multiple channels with shifted features. And it turns out that down-shuffle is simple but enough to achieve competitive performance without extra cost.

Refined-offset Deformable Alignment. Deformable convolution has recently shown compelling performance in aligning frames. Furthermore, Chan *et al.* [6] shed light on the underlying mechanism of deformable alignment, suggesting that the increased diversity in deformable alignment promises better restoration performance than flow-based alignment. The key idea of deformable alignment is to displace the sampling locations of standard convolution by some learned offsets. Through experiments, we observed that sampling more locations is not as easy as stacking more layers, which can cause instability problems of learned offsets.

To this end, we propose a refined-offset deformable alignment to ease the training of deformable convolutional layer and steadily expand the receptive field to sample long-range locations, whose formulation is

$$O^1 = G(F^{Ref}, F^{LQ}), \qquad (1)$$

$$O^r = R(O^{r-1}) + O^{r-1}, 2 \le r \le n, \qquad (2)$$

where G refers to the offset generator and R refers to the offset refiner, as shown in Fig. 3. F^{Ref} and F^{LQ} refer to features of reference and LR compressed frames respectively. Precisely, the core alignment module consists of n offset refiners such that the $(r-1)^{th}$ updates the offset residues between the estimated and the previous to obtain offsets O^r. The rationale behind this algorithm design is residual learning [15]. We hypothesize that it is easier to update the offset residues than to estimate unreferenced offsets. To the extreme, if the previous offsets were optimal, it would be easier to push the residues to zero, compared with stacking deformable convolution layers, which regresses to learn "ordinary" convolutional layers and further performs identity mapping. Through these benefits, refined-offset deformable alignment can handle larger motion with a wider receptive field and its coarse-to-fine strategy.

Here, layers built for residual learning also consider the eccentricity of receptive fields. To enable the offset to cover a wide range of areas, sampling locations from near to far distances, we introduce a basic module, termed Inception Hybrid Dilated Convolution (Incep-HDC).

Incep-HDC. As shown in Fig. 4, The inner structure of Incep-HDC can be divided into two components: the multi-branch convolution layers with different dilation rates and the spatial attention module.

Like seminal works in semantic segmentation and object detection [25, 40], recent advances in video SR [18] also carefully design the combinations of different dilation
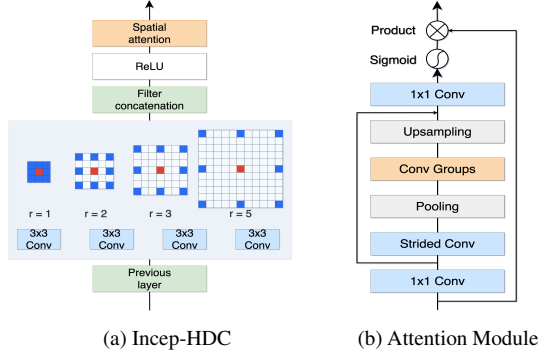
(a) Incep-HDC   (b) Attention Module

Figure 4: Illustration of Incep-HDC and spatial attention module.

rate, according to the motion distance in divided temporal groups. However, since both large and small motion may occur in different region of the same frame, we are motivated to utilize spatial attention mechanism. We conduct investigation between channel attention and spatial attention, which will be detailed in Table 4. And we come to the conclusion that, to reach various and distant positions, multi-branch convolutions with different receptive fields are best to be used in conjunction with the spatial attention layer.

3.3. Disentangled Loss

A trained R3N has also carried on the synthetical improvement at the aspect of loss function design.

When encountering severe noise/artifact interference, SR algorithms will suffer a huge performance degradation. There is a strong motivation on avoiding or mitigating such problems. For existing methods, ranging from loop filters to data-driven algorithms, there is no mechanism or formulization to suppress various noises in signals. Specifically, algorithms could be either too aggressive and amplify erroneous high-frequency components, or too conservative and tend to smooth over ambiguous components, both resulting in bad cases that seriously affect subjective visual impression. To alleviate such problems, we turn to additionally introduce a novel loss function into network optimization.

Definition: Returning to the design of disentangled loss, the main idea is to introduce the texture analysis into network optimization by measuring statistics collected over the entire image. The disentangled loss function between the ground truth image Y, and the prediction $\hat{Y} = F_{SR}(X)$ is given by

$$\mathcal{L}_{distan}(X, \hat{Y}, Y) = \ell_p(D(Y, X), D(\hat{Y}, X))$$
$$= \left\| D(Y, X) - D(\hat{Y}, X) \right\|_p, \quad (3)$$

where $D(\cdot)$ is the texture analysis procedure that distinguishes the pattern of signal and noise. The proposed
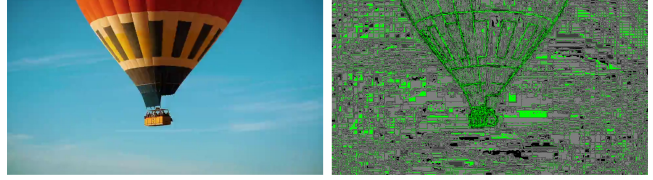

Figure 5: **Left**: A frame captured from a compressed video. **Right**: Visualization of an analysis map on the left image. The mixture of artifact, texture and flat regions are partitioned into black, green and gray pixels respectively. For instance, the ringing and motion prediction error near the boundaries that look like "mosquitos" flying around the ballon.

method tend to disentangle high frequency components correctly by minimizing a distance, for instance ℓ_p with p = 1 or p = 2. We next formulate $D(\cdot)$ step by step mathematically. Let $\uparrow$ denotes applying the bicubic upsampler, and F_{lp} denotes low-pass filtering, which is implemented in a differentiable manner with *Kornia* [30]. we derive the following procedure:

$$\mathcal{R}_{lp} = F_{lp}(x_\uparrow) - x_\uparrow, \quad (4)$$

$$\mathcal{R}_y = y - x_\uparrow, \quad (5)$$

$$D(y, x) = sgn(\mathcal{R}_y \odot \mathcal{R}_{lp}), \quad (6)$$

where $sgn(\cdot)$ denotes signum function that extracts the sign of a given pixel value; $\odot$ is the element-wise product. Finally, Eq. (6) gives pixel-wise classification results discretized into values of -1, 0 and 1. To facilitate understanding, we visualize an analysis result in Fig. 5.

4. Dataset

VSRE-set. Youku-VSRE [1] collects 1,000 HD visual lossless videos from *Youku* media database, and this dataset contains diverse contents that cover a variety of categories. It meets requirement of research on video transmission and post-processing algorithms. We reprocessed the source video from VSRE-set, and contribute a large-scale dataset for processing LR compressed video with combinations of encoding parameters, along with reference and ground truth. So that it can be used to train and benchmark SR on compressed videos.

Our distortions. The degradation methods of VSRE-set are described as following. We down-sampled HD video clips from VSRE-set with bicubic degradation on scaling factors of 2× and 4×. To obtain the target degraded clips, we generated the compressed frames through two coding settings, i.e., the latest HEVC [37] standard using HM 16.0 with the Low Delay P (LDP) mode. And the quantization parameter qp = 37 on frames with scaling factors of 2×, qp = 28 on frames with scaling factors of 4×. The above processing has constituted sets for 2× and 4× SR respectively.

In addtion, we generated the Ref images through SHM 6.1 with the All Intra (AI) mode at qp = 28. In short, each data sample consists of: source video clip, LR compressed video clip and Intra-coded video clip. LR video clip is the input and source video clip is the ground-truth.

For training Ref-SR models, the Ref image can be choosen from the Intra-coded video clip. And here for testing Ref-SR, the Ref image always refers to the first frame of the Intra-coded video clip.

Test set. In VSRE-set, 900 video clips with a total of 90, 000 frames are used as training set. The remaining video clips are re-grouped, and can be used as the validation and test set. Among them, we select 40 representative clips into our test set. We additionally cover more tests on 60 videos from JCT-VC [2] and VideoSet [39]. Finally, test set consists of 100 videos, denoted by *Video-Transmission-100 (VT-100)*. And HEVC sequences (Class A, B, E) have also been compressed with a wide range of QP spans for a wider test.

Note that, it is difficult to collect a dataset with strict bitrate control and alignment. The reprocessed VSRE-set mainly aims to simulate the real and complex degradation in video transmission, and to train and verify the restoration algorithms. However, we controlled the bitrate to the same range for experiments in Sec. 5.1. Experiments have shown that a R3N trained on VSRE-set can effectively be generalized to other datasets.

5. Experiments

5.1. Coding Setup and Performance

For our simulation framework, the HEVC base layer was encoded by software HM 16.0 while the HEVC enhancement layer was encoded by SHM 6.1, and Bidirectionally predicted frames (B-frames) are disabled in codec. In practical, GOP typically contains between 6 to 16 frames that are visually similar with no scene transitions. We also set the GOP with 16 frames, so the proportion of I-frames to the total number of frames is 1/16. And we conduct experiments compared with the conventional and the MultiRC coding chain. Their pipelines are shown schematically in Fig. 2.

In order to evaluate the proposed framework, we test HEVC sequences with a wide bitrate range controlled by Constant QP (CQP). For the conventional coding chain, the QP values are set between 28 and 48. And in order to create Rate-Distortion (RD) curves for the MultiRC and MixedRC that are located in the same bitrate range, QPs from 20 to 40 are used to compress the 2× down-sampled test sequences. For MultiRC, we use EDVR (2×) [41] at the decoder side. For clarity, the RD curves of *Traffic* and *ParkScene* sequences are drawn exemplary in Fig. 6 with PSNR as quality metric. And overall BD-rates are reported
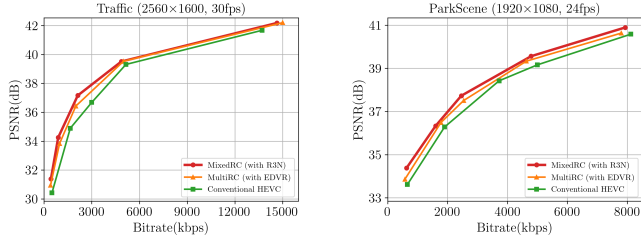

Figure 6: PSNR over bitrate for *Traffic* and *ParkScene*.

base & anchor codec	HEVC sequences	PSNR	VMAF
Conventional HEVC	Class A	-14.43%	-25.20%
	Class B	-16.35%	-26.77%
	Class E	-8.35%	-22.91%
MultiRC (with EDVR)	Class A	-7.08%	-18.20%
	Class B	-6.38%	-16.09%
	Class E	-3.91%	-7.53%

Table 1: Coding performance comparison of MixedRC over HEVC anchor (HM) and MultiRC.

in Table 1, with PSNR and Video Multimethod Assessment Fusion (VMAF) [22] as quality metrics.

Experimental results show that both MixedRC and MultiRC chains surpass results of the conventional HEVC up to a certain bitrate around 7000 kbit/s, as well as MixedRC raise a curve at the upper left of the conventional and the MultiRC. Moreover, the proposed MixedRC shows superior results especially with VMAF as quality metric.

5.2. Network Settings

The deformable alignment module adopts one offset generator and three offset refiner, and each offset refiner contains two Incep-HDC blocks. The network takes three consecutive LR frames as inputs in 2× model, and five consecutive frames in 4× model. As we observed through experimental results that the worse compression quality, the less complementary information in adjacent decoded frames.

We train our model with Adam optimizer by setting $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The learning rate is initialized as 10^{-4} and then decreases to half whenever the validation loss stops decreasing for more than 5 epochs. Mini-batch size is set to 16. In each training mini-batch, patches randomly cropped with size 192×192 are used. By analyzing the dataset, we conclude that a large patch size is necessary for the Ref-SR task. And we train R3N following an intuitive idea of "easy-hard transfer", which is discussed in Sec. 6. We implement our models with the PyTorch framework and train them using 8 NVIDIA Tesla P40 GPUs.

5.3. Network Performance

We choose a few representative methods for comparisons: RCAN [49], EDVR [41], SRNTT [51] which demonstrate different network architectures for single image SR,

Codec Chain	Restoration Methods	540p-to-1080p (2×)			270p-to-1080p (4×)		
		SSIM↑	PSNR↑	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓
MultiRC	Bicubic	0.933	30.11	0.195	0.912	27.83	0.312
MultiRC	RCAN [49]	0.939	30.96	0.176	0.916	28.35	0.268
MultiRC	EDVR [41]	0.944	31.80	0.159	0.925	29.51	0.220
MixedRC	SRNTT [51]	0.940	31.34	0.142	0.921	29.29	0.181
MultiRC	MF-CNN [45] + Bicubic	0.941	31.65	0.163	0.916	28.21	0.283
MultiRC	MF-CNN [45] + EDVR [41]	0.942	31.97	0.156	0.922	29.36	0.242
MixedRC	R3N(Ours)	**0.951**	**32.49**	**0.138**	**0.935**	**30.58**	**0.177**

Table 2: Quantitative results of coding chains with corresponding SR and CAR methods, for video restoration on our collected VT-100 (2× and 4×). Best and second best results are highlighted and underlined, and note that a lower LPIPS score indicates better image quality.
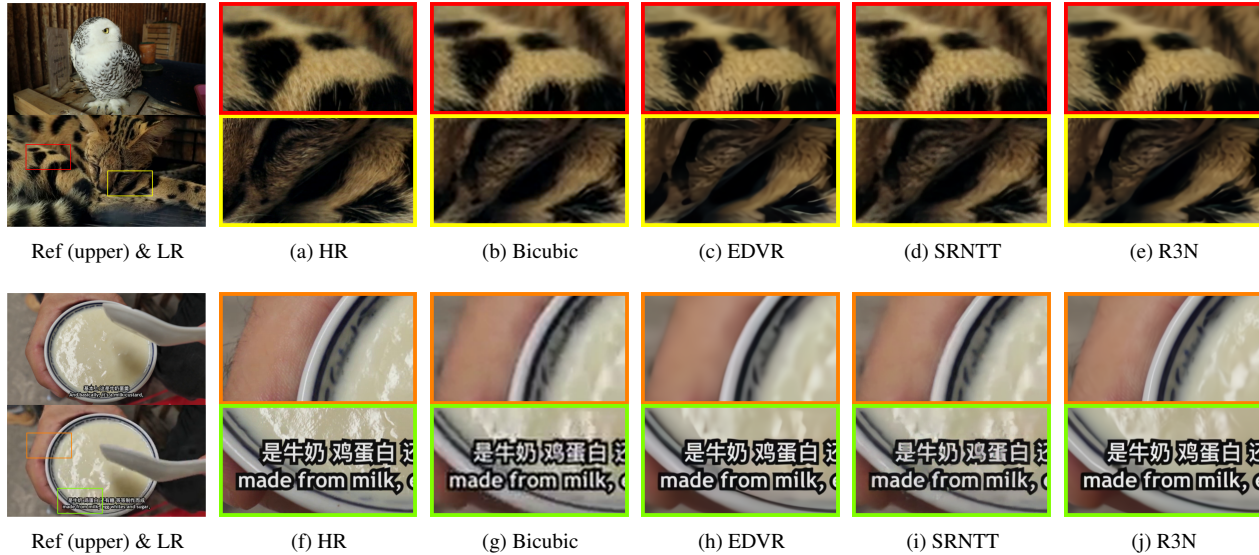

Figure 7: R3N (ours) is compared to EDVR [41] (a state-of-the-art video SR method) and SRNTT [51] (a state-of-the-art Ref-SR method). In the first row, the Ref image has irrelevant content to the LR input. In the second row, the Ref image and the LR input are in the same scene.

video SR and Ref-SR, and MF-CNN [12] which is designed to enhance the quality of compressed video. The structure of SR methods have also been proven to achieve good performance in other general low-level tasks. Therefore, we retrained these models to perform SR and CAR jointly. It also because, to our knowledge, that this study is the first attempt to explore effectiveness of DCNN in joint SR-CAR problem in literature. And for a fair comparison, we use the same data processing method and training data from VSREset. Since MF-CNN is not a SR model, we down-sampled ground-truth as the label for training it. And we enable memory-efficient forward so that SRNTT can be tested on HD video sequences.

To quantitatively evaluate the proposed method, we use the standard Structural Similarity (SSIM), Peak Signal To Noise Ratio (PSNR) and Learned Perceptual Image Patch Similarity (LPIPS) [48]. The results are reported in Table 2. On VT-100, R3N and SRNTT achieve the best and second performance in terms of LPIPS. One can clearly see that

Restoration Methods	RCAN	EDVR	SRNTT	R3N (Ours)
Times (second)	0.321	0.694	30.163	0.473

Table 3: Test the speeds of restoration methods (reconstructing a frame from 270p to 1080p using a GPU of NVIDIA Tesla V100).

R3N outperforms other methods by a large margin in terms of PSNR, SSIM and LPIPS.

Moreover, visual quality comparisons are presented in Fig. 7. Our proposed method recovers more high-fidelity visual information compared to others, especially in Fig. 7 (j), where the skin texture is only restored by R3N. By contrast, SRNTT generates more twisty artifacts than ours. Our method could still achieve state-of-the-art SR performance when scene switching occurs, shown in Fig. 7 (e), demonstrating the robustness of R3N. Furthermore, R3N is superior to EDVR and SRNTT on the implementation speed, shown in Table 3.

(a) Ref & LR   (b) w/o refined-offset alignment
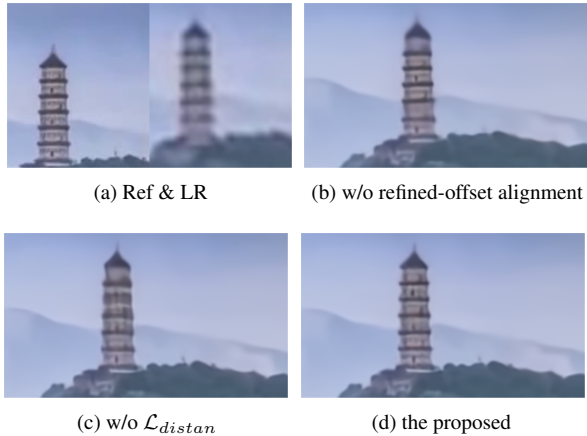
(c) w/o $\mathcal{L}_{distan}$ (d) the proposed

Figure 8: Visual comparison for ablation testing. Note that there are significant motion between Ref and LR frames, which could be seen clearly by embedding their patches next to each other in (a).

Method	(a)	(b)	(c)	(d)	(e)
Stacking DeformConv	✓				
Refined-offset		✓	✓	✓	✓
Channel attention module		✓			
Spatial attention module	✓		✓	✓	✓
Disentangled loss				✓	✓
Easy-hard transfer					✓
VT-100 (4×)	29.46	29.87	30.12	30.31	30.58

Table 4: Ablation study on the proposed methods. Our results are validated by the PSNR score.

6. Ablation Study

We study the effects of each component in the proposed method, as reported in Table 4.

Effectiveness of Refined-offset Deformable Alignment. As illustrated in Table 4, we introduce two baselines: (a) and (c), where the model (a) stacks three layers of deformable convolution. For fairness, we set the same total number of Incep-HDC blocks as 16 in (a) and (c). As expected, R3N has at least an average PSNR gain of 0.6dB over stacking deformable convolution layers due to the residual learning manner. Visual comparisons are provided in Fig. 6 (b) and (d).

Effectiveness of Disentangled Loss. To demonstrate the effectiveness of disentangled loss $\mathcal{L}_{distan}$, we train one of our models with ℓ_1 only and another same model with $\mathcal{L}_{distan}$ and ℓ_1, and here p = 1 in Eq. (3). In Fig. 6 (c), we can see result without $\mathcal{L}_{distan}$ would suffer from some erroneous texture artifacts. Such observations futher prove the effectiveness of our proposed disentangled loss.

Effect of Attention Module. When we compare the results of (b) and (c) in Table 4, we find that Incep-HDC with spatial attention would perform better than those with channel

attention. The implementation of spatial and channel attention is in line with RFANet [23] and RCAN [49].

Effect of Transfer Learning. Transfer learning in deep models provides a good starting point that could help train a better network with higher convergence speed. In frame alignment of Ref-SR and video SR, it is observed that using training data with large motion would encounter the problem of convergence. We are also met with this difficulty during the training process of R3N. To this end, we conduct the transfer learning settings as follows: we trained a model whose Ref is selected randomly between the I_{t-d}^{Ref} and I_{t+d}^{Ref}, where d=8 at first. Then we reuse the features learned in the former trained model to initialize a harder network, where d is increased to 16, 24 and 32 gradually. Results in Table 4 show that easy-hard transfer can bring a gain of 0.27dB, compared with the model trained directly with d=32.

7. Conclusion

This paper explains that advancement in deep convolutional network opens up expanded design space and has far-reaching implications in the video transmission ecosystem. First, we introduced the mixed-resolution video coding which not only improves the coding efficiency, but also leaves more improvement room for restoration algorithms. However, the picture quality loss in video transmission mainly comes from resolution loss and compression artifacts. Existing SR algorithms suffer a huge peformance degradation when encountering this intricate task. So it is of particular interest to bring a joint SR-CAR solution.

Aiming at this problem, we embrace and carry on improvement on Ref-SR methods. We propose a refined-offset deformable alignment module, at the mean time, we propose a disentangled loss to distinguish different components in high-frequency regions. It is worth mentioning that the proposed loss also has potential for the denoising task. We will leave them to further work.

References

[1] Youku-VSRE: youku video super-resolution and enhancement challenge dataset, 2019. https://tianchi.aliyun.com/dataset/dataDetail?dataId=39568. 5

[2] Frank Bossen et al. Common test conditions and software reference configurations. In *JCTVC-L1100*, volume 12, 2013. 6

[3] Jill M Boyce, Yan Ye, Jianle Chen, and Adarsh K Ramasubramanian. Overview of SHVC: Scalable extensions of the high efficiency video coding standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(1):20–34, 2015. 1

[4] Fernanda Brandi, Ricardo de Queiroz, and Debargha Mukherjee. Super-resolution of video using key frames and

motion estimation. In *2008 15th IEEE International Conference on Image Processing*, pages 321–324. IEEE, 2008. 2

[5] Alfred M Bruckstein, Michael Elad, and Ron Kimmel. Down-scaling for better transform compression. *IEEE Transactions on Image Processing*, 12(9):1132–1144, 2003. 2

[6] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Understanding deformable alignment in video super-resolution. *arXiv preprint arXiv:2009.07265*, 2020. 4

[7] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. 2

[8] Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 576–584, 2015. 1

[9] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199. Springer, 2014. 1

[10] Kristian Fischer, Christian Herglotz, and André Kaup. On versatile video coding at uhd with machine-learning-based super-resolution. In *2020 Twelfth International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–6. IEEE, 2020. 2

[11] Georgios Georgis, George Lentaris, and Dionysios Reisis. Reduced complexity superresolution for low-bitrate video compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(2):332–345, 2015. 2

[12] Zhenyu Guan, Qunliang Xing, Mai Xu, Ren Yang, and Zulin Wang. MFQE 2.0: A new approach for multi-frame quality enhancement on compressed video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP(99):1–1, 2019. 3, 7

[13] Bumsub Ham, Minsu Cho, and Jean Ponce. Robust image filtering using joint static and dynamic guidance. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4823–4831, 2015. 2

[14] Kaiming He, Jian Sun, and Xiaoou Tang. Guided image filtering. *IEEE transactions on pattern analysis and machine intelligence*, 35(6):1397–1409, 2012. 2

[15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4

[16] Thomas W Holcomb, Shankar Regunathan, Chih-lung Bruce Lin, and Sridhar Srinivasan. Multi-resolution video coding and decoding, May 27 2008. US Patent 7,379,496. 2

[17] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. 2

[18] Takashi Isobe, Songjiang Li, Xu Jia, Shanxin Yuan, Gregory Slabaugh, Chunjing Xu, Ya-Li Li, Shengjin Wang, and Qi Tian. Video super-resolution with temporal group attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8008–8017, 2020. 4

[19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 2

[20] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Deep SR-ITM: Joint learning of super-resolution and inverse tone-mapping for 4k uhd hdr applications. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3116–3125, 2019. 3

[21] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 2

[22] Zhi Li, Christos Bampis, Julie Novak, Anne Aaron, Kyle Swanson, Anush Moorthy, and JD Cock. Vmaf: The journey continues. *Netflix Technology Blog*, 25, 2018. 6

[23] Jie Liu, Wenjie Zhang, Yuting Tang, Jie Tang, and Gangshan Wu. Residual feature aggregation network for image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2359–2368, 2020. 8

[24] Lin Liu, Xu Jia, Jianzhuang Liu, and Qi Tian. Joint demosaicing and denoising with self guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2240–2249, 2020. 3

[25] Songtao Liu, Di Huang, et al. Receptive field block net for accurate and fast object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 385–400, 2018. 4

[26] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Zhiyong Gao, and Ming-Ting Sun. Deep kalman filtering network for video compression artifact reduction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 568–584, 2018. 3

[27] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. 2

[28] Florian Maurer, Stefano Battista, Lorenzo Ciccarelli, Guido Meardi, and Simone Ferrara. Overview of mpeg-5 part 2–low complexity enhancement video coding (lcevc). *ITU Journal: ICT Discoveries*, 3(1), 2020. 1, 2

[29] Guocheng Qian, Jinjin Gu, Jimmy S Ren, Chao Dong, Furong Zhao, and Juan Lin. Trinity of pixel enhancement: a joint solution for demosaicking, denoising and super-resolution. *arXiv preprint arXiv:1905.02538*, 2019. 3

[30] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3674–3683, 2020. 5

[31] Heiko Schwarz, Detlev Marpe, and Thomas Wiegand. Overview of the scalable video coding extension of the H.

264/AVC standard. *IEEE Transactions on circuits and systems for video technology*, 17(9):1103–1120, 2007. 1, 2, 3

[32] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949. 2

[33] Minmin Shen, Ping Xue, and Ci Wang. Down-sampling based video coding using super-resolution technique. *IEEE Transactions on Circuits and Systems for Video Technology*, 21(6):755–765, 2011. 2

[34] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016. 4

[35] Gyumin Shim, Jinsun Park, and In So Kweon. Robust reference-based super-resolution with similarity-aware deformable convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8425–8434, 2020. 2

[36] Masanori Suganuma, Xing Liu, and Takayuki Okatani. Attention-based adaptive selection of operations for image restoration in the presence of unknown combined distortions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9039–9048, 2019. 3

[37] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12):1649–1668, 2012. 1, 5

[38] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. TDAN: Temporally-deformable alignment network for video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3360–3369, 2020. 2

[39] Haiqiang Wang, Ioannis Katsavounidis, Jiantong Zhou, Jeonghoon Park, Shawmin Lei, Xin Zhou, Man-On Pun, Xin Jin, Ronggang Wang, Xu Wang, et al. Videoset: A large-scale compressed video quality dataset based on jnd measurement. *Journal of Visual Communication and Image Representation*, 46:292–302, 2017. 6

[40] Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, and Garrison Cottrell. Understanding convolution for semantic segmentation. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 1451–1460. IEEE, 2018. 4

[41] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. EDVR: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1, 2, 6, 7

[42] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 2

[43] Zhou Wang and Alan C Bovik. Mean squared error: Love it or leave it? a new look at signal fidelity measures. *IEEE signal processing magazine*, 26(1):98–117, 2009. 2

[44] 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*, pages 3370–3379, 2020. 3

[45] Ren Yang, Mai Xu, Zulin Wang, and Tianyi Li. Multi-frame quality enhancement for compressed video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6664–6673, 2018. 7

[46] Ke Yu, Chao Dong, Liang Lin, and Chen Change Loy. Crafting a toolchain for image restoration by deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2443–2452, 2018. 2, 3

[47] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Learning a single convolutional super-resolution network for multiple degradations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3262–3271, 2018. 2, 3

[48] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 7

[49] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 286–301, 2018. 1, 4, 6, 7, 8

[50] Yulun Zhang, Zhifei Zhang, Stephen DiVerdi, Zhaowen Wang, Jose Echevarria, and Yun Fu. Texture hallucination for large-factor painting super-resolution. 2

[51] Zhifei Zhang, Zhaowen Wang, Zhe Lin, and Hairong Qi. Image super-resolution by neural texture transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7982–7991, 2019. 1, 2, 6, 7

[52] Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss functions for image restoration with neural networks. *IEEE Transactions on computational imaging*, 3(1):47–57, 2016. 2

[53] Haitian Zheng, Mengqi Ji, Haoqian Wang, Yebin Liu, and Lu Fang. Crossnet: An end-to-end reference-based super resolution network using cross-scale warping. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 88–104, 2018. 2