

Exploit Camera Raw Data for Video Super-Resolution via Hidden Markov Model Inference

Xiaohong Liu¹, Graduate Student Member, IEEE, Kangdi Shi, Zhe Wang²,
and Jun Chen¹, Senior Member, IEEE

Abstract—To the best of our knowledge, the existing deep-learning-based Video Super-Resolution (VSR) methods exclusively make use of videos produced by the Image Signal Processor (ISP) of the camera system as inputs. Such methods are 1) inherently suboptimal due to information loss incurred by non-invertible operations in ISP, and 2) inconsistent with the real imaging pipeline where VSR in fact serves as a pre-processing unit of ISP. To address this issue, we propose a new VSR method that can directly exploit camera sensor data, accompanied by a carefully built Raw Video Dataset (RawVD) for training, validation, and testing. This method consists of a Successive Deep Inference (SDI) module and a reconstruction module, among others. The SDI module is designed according to the architectural principle suggested by a canonical decomposition result for Hidden Markov Model (HMM) inference; it estimates the target high-resolution frame by repeatedly performing pairwise feature fusion using deformable convolutions. The reconstruction module, built with elaborately designed Attention-based Residual Dense Blocks (ARDBs), serves the purpose of 1) refining the fused feature and 2) learning the color information needed to generate a spatial-specific transformation for accurate color correction. Extensive experiments demonstrate that owing to the informativeness of the camera raw data, the effectiveness of the network architecture, and the separation of super-resolution and color correction processes, the proposed method achieves superior VSR results compared to the state-of-the-art and can be adapted to any specific camera-ISP. Code and dataset are available at <https://github.com/proteus1991/RawVSR>.

Index Terms—Video super-resolution, camera raw data, hidden Markov model inference.

I. INTRODUCTION

SUPER-RESOLUTION (SR) is a promising technique that can restore high-resolution (HR) pictorial data from their low-resolution (LR) counterpart without requiring hardware upgrades. Over the past few decades, it has found applications in a wide range of areas such as medical imaging [1], [2], satellite imaging [3], [4] and surveillance [5], [6]. SR is also useful

Manuscript received July 29, 2020; revised December 3, 2020; accepted December 30, 2020. Date of publication January 13, 2021; date of current version January 25, 2021. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) through a Discovery Grant. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nikolaos Mitianoudis. (Corresponding author: Xiaohong Liu.)

Xiaohong Liu, Kangdi Shi, and Jun Chen are with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, ON L8S 4K1, Canada (e-mail: liux173@mcmaster.ca; shik9@mcmaster.ca; chenjun@mcmaster.ca).

Zhe Wang is with the David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: zhe.wang@uwaterloo.ca).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TIP.2021.3049974>, provided by the authors.

Digital Object Identifier 10.1109/TIP.2021.3049974

for improving the quality of data dedicated to high-level vision tasks [7]–[9]. There are two major categories of SR: Single Image Super-Resolution (SISR) and Video Super-Resolution (VSR). Although early studies [10]–[12] largely treat VSR as a simple extension of SISR by focusing on intra-frame spatial correlation, more attention has been paid in recent works [13]–[16] to strategically exploiting inter-frame temporal correlation (typically in the form of motion estimation and compensation) to further improve the VSR results. The availability of an extra dimension in VSR creates both opportunities and challenges. Indeed, there is considerable freedom in the ways that multiple LR frames can be leveraged to reconstruct one target HR frame, especially with the advent of deep learning techniques. Even though many effective heuristics have been proposed, a theoretical guideline for the fusion process is still lacking.

The data-driven approach has become increasingly popular in VSR research. To the best of our knowledge, the existing data-driven VSR methods [13]–[16] exclusively make use of camera processed data as inputs, despite the fact that modern cameras are capable of providing raw data which are potentially more informative. Note that camera processed data are produced by the Image Signal Processor (ISP) from raw data through several operations, including image demosaicing, denoising, sharpening, color converting, tone adjustment, compression, among others [17]–[19]. Overall, the operations in ISP are non-invertible and tend to degrade the information content of the original data. For example, the bit-depth is reduced from 12–14 bits to 8 bits [20] due to quantization, and various artifacts are introduced by compression (JPEG is the default image format for many cameras). Therefore, using camera processed data as inputs is inherently suboptimal. Moreover, such methods are inconsistent with the real imaging pipeline where VSR in fact serves as a pre-processing unit rather than a post-processing unit of ISP.

The main contributions of this paper are summarized as follows:

- 1) To avoid information loss and better fit the imaging pipeline, we propose a new Raw VSR method, named RawVSR, that can directly exploit camera sensor data. Furthermore, we carefully build the first Raw Video Dataset (RawVD) for training, validation, and testing.
- 2) Assuming a Hidden Markov Model (HMM), we show that the minimal sufficient statistic of LR frames with respect to the target HR frame can be computed via

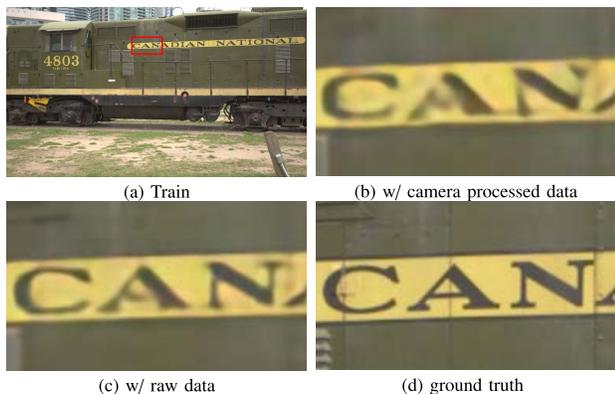


Fig. 1. Illustrations of the proposed RawVSR based on camera processed data and raw data, respectively, for 4 \times VSR using the testing video “Train” in RawVD.

an iterative pairwise fusion process; this result provides the architectural principle for the design of RawVSR, especially its Successive Deep Inference (SDI) module.

- 3) We design an Attention-based Residual Dense Block (ARDB) and leverage it to build the reconstruction module of RawVSR; this module is capable of simultaneously refining the fused feature and learning the color information needed to generate a spatial-specific transformation for accurate color correction.

Our extensive experimental results demonstrate that owing to the informativeness of the camera raw data, the effectiveness of the neural network architecture, and the separation of super-resolution and color correction processes, the proposed RawVSR achieves superior performance compared to the state-of-the-art and can be adapted to any specific camera-ISP.

Fig. 1 illustrates the proposed RawVSR based on camera processed data and raw data, respectively, for 4 \times VSR using the testing video “Train” in RawVD. It can be seen that the latter yields clearer texture and sharper edges.

II. RELATED WORK

To place our work in a proper context, we give a review of the existing VSR methods and discuss some relevant issues.

VSR can be viewed as an underdetermined problem. Many traditional methods tackle this problem via regularized optimization, aided by certain image priors [21]–[26]. The advent of deep learning has led to a paradigm shift in VSR and more generally low-level vision research [27]–[30]. Indeed, following the pioneering work by Dong *et al.* [31] on SISR, the recently proposed VSR methods are predominantly deep-learning-based [10], [11], [13]–[16], [32]–[35], achieving performances out of reach of the traditional prior-based approach. Here we just describe a few representative ones. Kappeler *et al.* [11] resort to the hand-crafted optical flow [36] for motion compensation across input frames, and adopt a pre-trained CNN for SR operation. Caballero *et al.* [33] improve the quality of estimated optical flow using a spatial transformer module to facilitate the subsequent spatial-temporal network [37] for VSR. Tao *et al.* [34] design a Sub-Pixel Motion Compensation (SPMC) layer to perform motion compensation and resolution upsampling simultaneously. A novel ConvLSTM is put

forward in [32], which can be integrated into the autoencoder structure to effectively exploit temporal correlations among input frames. The 3D convolution finds its first application to motion estimation and compensation in [14], where the spatial and temporal correlations are explored coherently via 3D convolving operations. Inspired by back-projection [38], Haris *et al.* [15] develop a fusion strategy that treats each input frame as a stand-alone source carrying distinct information and progressively embeds extracted correlation from each source-target pair into the back-projection network for motion compensation. In consideration of the fact that inter-frame displacements could spread over a wide range, Tian *et al.* [39] substitute the conventional convolution with a deformable one that is capable of exploring spatial correlation via learnable sampling points tailored to each pixel, not constrained by any prescribed kernel shape. Wang *et al.* [16] propose a refined version of deformable convolution and construct a Pyramid, Cascading and Deformable (PCD) module for motion alignment; the resulting method is the winner of the NTIRE 2019 challenge on video deblurring and super-resolution [40], and achieves the best VSR performance to date.

Below we identify several issues with the existing VSR methods that will be addressed in the present work.

A. Raw Data Processing

The advantage of untouched camera raw data over processed data has been recognized in several areas of low-level vision. For instance, Chen *et al.* [27] make use of raw data to perform fast imaging in low-light and achieved favorable results in terms of texture detail. This success can be attributed to the primitive radiance information retained by raw data. Indeed, using camera processed data as a substitute of raw data fails to deliver comparable results due to the information loss caused by quantization in ISP, which obscures subtle differences in pixel values that are essential for exhibiting delicate textures. Ignatov *et al.* [41] demonstrate that color images suitably converted from smartphone’s raw data are visually comparable to those produced by a professional camera. A comprehensive study of raw-data-based SISR is conducted by Xu *et al.* [19], showing that the benefits of raw data are intrinsic, *i.e.*, not specific to a particular SISR method [42]. However, to the best of our knowledge, the existing VSR methods are still exclusively based on camera processed data, and there is a lack of understanding of the potential gain offered by raw data (beyond that seen in the SISR setting).

B. Motion Estimation

Due to the temporal dependencies among video frames, VSR can benefit substantially from precise motion estimation, which can be performed either explicitly or implicitly. Explicit motion estimation usually relies on optical flows [11], [13], [33], [34]. However, for video frames with large pixel displacement and object occlusion, it could be extremely challenging and arguably impossible to obtain precise optical flows. Moreover, as indicated by [43], even completely accurate optical flows might not be adequate for motion estimation since they do not fully capture all possible motion effects. Therefore, many recent VSR methods such as DUF [14], TDAN [39], and

EDVR [16] choose to estimate motion implicitly, resulting in better final reconstructions compared to the flow-based ones. In a certain sense, implicit motion estimation provides a new mechanism for multi-frame fusion. Nevertheless, many aspects of the fusion process remain unspecified, and a theoretical guideline is much needed.

C. Color Correction

Different from camera processed images, raw images only record the primitive radiance information, which has not been projected to any color space. In [27], an end-to-end mapping is learned that transforms raw images to colored images directly. However, this mapping does not have the desired flexibilities due to its camera-specific nature. [18] addresses this issue by generating a global color transformation that can cope with a variety of cameras. But it has been observed that such a transformation may cause significant local color distortions. This problem is solved by [19] through the introduction of a spatial-specific color transformation. Note that [19] makes use of a stand-alone deep convolutional network for this purpose, which increases the overall model size and computational load. In contrast, we show that spatial-specific color transformation can actually be accomplished together with HR image reconstruction by a single network. Our new design leads to more efficient implementation and facilitates joint learning for both tasks.

III. RAW VIDEO DATASET

The widely-used VSR training datasets such as Vimeo-90K [43] and REDS [40] are constructed with post-ISP videos (*e.g.*, H.264 videos) and, as a consequence, are inherently unsuitable for raw-data-based VSR. We are not aware of the existence of publicly available raw video datasets for VSR. An important contribution of this work is a new raw video dataset (RawVD) for training and benchmarking. In fact, it will be seen that RawVD consists of both HR/LR raw video pairs and their processed counterparts, thus is suitable for essentially all VSR methods, regardless whether they are raw-data-based or not. To build this dataset, we use a Canon 5D3 camera with the third-party upgrade¹ to film videos in Magic Lantern Video (MLV) format (which is a raw video format), and utilize the MLV App² to obtain corresponding raw frames in DNG format (which is a raw image format). Note that different cameras might have different Bayer patterns. For Canon 5D3, the Bayer pattern is RGGB, which means 50% sensors for green color, 25% for red color, and 25% for blue color. A degraded LR raw frame $Y_{raw} \in \mathbb{R}^{H/S \times W/S \times 1}$ is generated for each filmed HR raw frame $X_{raw} \in \mathbb{R}^{H \times W \times 1}$, where H and W stand respectively for frame height and width, and S represents the VSR scale ratio. Specifically, inspired by [19], [27], we first use AHD [44] to produce a demosaiced frame $X_{lin} \in \mathbb{R}^{H \times W \times 3}$ based on X_{raw} without any post-processing such as white balance and gamma correction (where the subscript of X_{lin} reflects the fact that each of its pixel value is a linear measurement of some radiance information contained in X_{raw}), then generate Y_{raw} from X_{lin} through a sequence

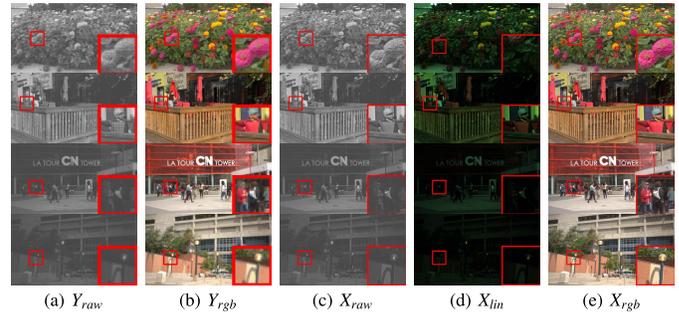


Fig. 2. Examples of videos in RawVD with the brightness and contrast of raw frames adjusted for better visualization.

of degradation operations involving, among others, blurring, downsampling, and noising. More precisely, we have

$$Y_{raw} = f_{Bayer}(f_{Down}(X_{lin} * K_{Blur})) + n. \quad (1)$$

Here K_{Blur} is a uniformly distributed defocus blur kernel, which simulates the out-of-focus effect in camera; $*$ denotes the convolutional operation; n is a heteroscedastic Gaussian noise [19], [45] with its variance specified by two parameters σ_1^2 and σ_2^2 ; f_{Down} and f_{Bayer} stand respectively for downsampling and mosaicing operations (in particular, f_{Bayer} converts a three-channel frame back to a one-channel frame that obeys the RGGB pattern). Finally, given X_{raw} and Y_{raw} , we use Rawpy, a Python version of LibRaw (with parameters adjusted to closely approximate Canon-ISP), to generate their corresponding color frames $X_{rgb} \in \mathbb{R}^{H \times W \times 3}$ and $Y_{rgb} \in \mathbb{R}^{H/S \times W/S \times 3}$. An illustration of Y_{raw} , Y_{rgb} , X_{raw} , X_{lin} , and X_{rgb} can be found in Fig. 2; note that the color of X_{lin} is biased towards green because the dominant measurement recorded in raw data is from green sensors. In total, 110 videos (with a resolution of 1920×1080 and at least 100 frames each) are filmed. These videos capture a wide variety of scenes. We divide these 110 videos into three groups: 100 videos for training, 5 for validation, and the rest 5 for testing. The 5 test videos are named “Store”, “Painting”, “Train”, “City”, and “Walk” according to their respective content.

As mentioned above, each video has four different versions (in addition to the linear measurement version (X_{lin})): the original HR raw version (X_{raw}), the LR raw version (Y_{raw}), the HR color version (X_{rgb}), and the LR color version (Y_{rgb}). The LR raw version and the LR color version serve as the inputs of raw-data-based VSR methods and processed-data-based VSR methods, respectively. We adopt the HR color version as the ground truth for both types of methods to facilitate direct comparisons of their outputs. The linear measure version is not directly used for supervised training in the present work; nevertheless, we choose to include it in our RawVD as it visually reveals the intimate connections among the other four versions and potentially provides valuable insights that can be fruitfully exploited in future research.

Remark: We have also considered collecting real-world raw LR-HR pairs through optical zooming [46]. However, this approach has several limitations. First, each LR-HR pair should have exactly the same motion (otherwise the pixel-level alignment is difficult to accomplish), which severely limits the varieties of suitable videos that can be collected in practice.

¹<https://magiclantern.fm/>

²<https://mlv.app/>

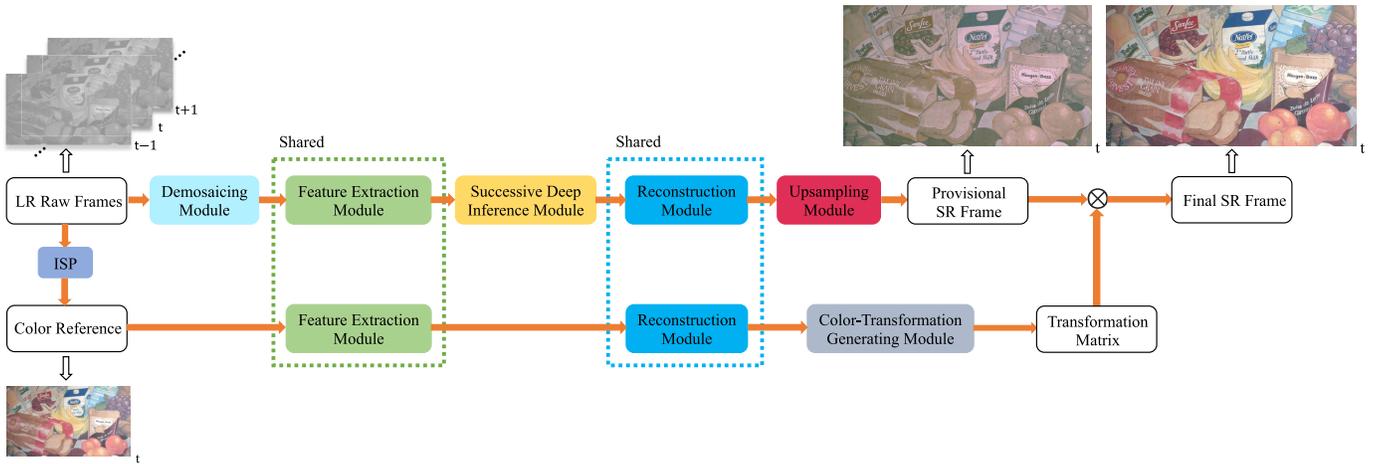


Fig. 3. The architecture of the proposed RawVSR.

Second, since raw data only record primitive radiance information, some back-and-forth conversions between raw data and processed data (e.g., sRGB) are typically needed for raw LR-HR pair alignment, causing information loss. As such, this approach is tailored to SISR (without inter-frame motion) with sRGB data. In contrast, our approach does not suffer from these problems and is better suited to raw VSR. It is also worth noting that even though AHD is not the state-of-the-art for demosaicing, it performs very stable and is widely adopted in practice [19]. Indeed, upon close scrutiny, no ground-truth video in the constructed RawVD suffers evident demosaicing artifacts.

IV. METHOD

Now we are in a position to present the proposed raw-data-based VSR method, *i.e.*, RawVSR. A general overview is provided in Section IV-A. Sections IV-B and IV-C are devoted to describing two key components of RawVSR: the SDI module and the reconstruction module. We address some potential practical concerns regarding raw-data-based VSR in Section IV-D. The implementation details can be found in the supplementary material.

A. Overview

To effectively exploit LR raw videos, one might be inclined to build a deep neural network to learn an end-to-end mapping that can generate HR color videos directly. However, in reality, raw videos only record the radiance information read from camera sensors and the associated color videos are in fact ISP-dependent. That is to say, different ISPs may produce different color videos based on the same raw video. As such, an end-to-end mapping can only be tailored to a specific type of camera ISP and lacks the flexibilities needed for accommodating diverse VSR requirements. To tackle this problem, we use one certain color reference generated by an ISP-adaptive operation to guide the transformation from the raw space to the color space.

The proposed RawVSR consists of several components: the demosaicing module, the feature extraction module, the SDI module, the reconstruction module, the upsampling module, and the color-transformation generating module. It takes 7

consecutive LR raw frames $Y_{raw}^{t-3}, \dots, Y_{raw}^{t+3}$ and outputs a SR color frame $X_{sr}^t \in \mathbb{R}^{H \times W \times 3}$. The goal is to make X_{sr}^t as close to the HR color frame X_{rgb}^t as possible. We use Y_{rgb}^t as the color reference frame to ensure that the provided color information is consistent with that of X_{rgb}^t . Note that the color reference frame is derived from the middle frame in the raw video sequence through an ISP-adaptive operation, and thus should not be viewed as an additional input. Indeed, except some ISP-specific parameters needed for color transformation and correction, this middle frame already contains all necessary information to produce the color reference frame.

As shown in Fig. 3, RawVSR has two branches: the texture restoration branch (upper branch) produces a provisional SR frame $\hat{X}_{sr}^t \in \mathbb{R}^{H \times W \times 3}$ while the color correction branch (lower branch) generates a spatial-specific color transformation T . The final SR frame X_{sr}^t is obtained from \hat{X}_{sr}^t through pixel-wise color correction specified by T . The orange arrow shows the image/feature flow direction, and the black arrow points to the illustration of input/color reference/provisional SR/final SR frame respectively using the ‘‘Painting’’ video in RawVD.

In the texture restoration branch, the input LR raw frames are first processed by the demosaicing module, where each LR raw frame is converted to a 64-channel frame (with the spatial size unchanged) via packing, convolution, and pixel shuffling. We utilize the U-net [47] for frame-by-frame multi-scale feature extraction. The outputs of the feature extraction module induced by the given seven LR raw frames are jointly fed into the SDI module to perform feature alignment and fusion. The fused feature is then refined by the reconstruction module. The upsampling module consists of three convolutional layers and one pixel-shuffle operator. It enlarges the spatial size of the refined feature to generate a provisional SR frame. It can be seen that although the provisional SR frame suffers severe color distortion, the texture details are faithfully restored.

In the color correction branch, both the feature extraction module and the reconstruction module are shared from the texture restoration branch. To ensure dimension compatibility, the color reference frame is pre-processed by a single convolutional layer before fed into the feature extraction module. Different from its main role in the texture restoration branch

(which is feature refinement), the reconstruction module is used in the color correction branch to learn the ISP information needed for generating a spatial-specific color transformation. The structure of the color-transformation generating module is the same as that of the upsampling module except for the last convolutional layer, which outputs a 9-channel feature map rather than a 3-channel one. The output feature map is reshaped from $H \times W \times 9$ to $H \times W \times 3 \times 3$, and the reshaped map is denoted by T . Note that T can be viewed as a collection of 3×3 matrices $\{T(i, j)\}$ with each (i, j) pair specifying a pixel position. These matrices are leveraged to perform a spatial-specific color transformation on the provisional SR frame \hat{X}_{sr}^t to produce the final SR frame X_{sr}^t as follows:

$$X_{sr}^t(i, j) = T(i, j) \otimes \hat{X}_{sr}^t(i, j), \quad (2)$$

where $X_{sr}^t(i, j) \in \mathbb{R}^{3 \times 1}$ ($\hat{X}_{sr}^t(i, j) \in \mathbb{R}^{3 \times 1}$) is the RGB vector of the (i, j) pixel of X_{sr}^t (\hat{X}_{sr}^t), and \otimes denotes matrix multiplication.

It is worth noting that due to the sharing of the feature extraction module and the reconstruction module as well as the lightweight design of the color-transformation generating module, the number of independent parameters that come from the color correction branch is negligible (accounting for about 5% of the model size of RawVSR). On the other hand, from the perspective of building an interpretable device-independent network, it is preferable to have a dual-branch architecture that disentangles the super-resolution process from the color correction process, and one may argue that weight sharing does not suit this purpose. However, since the generated transformation matrix needs to be properly aligned with the provisional SR frame to perform pixel-wise color correction, it is essential to have a mechanism to facilitate the coordination of two processes. In this regard, weight sharing contributes to this coordination by helping maintain the latent consistency of intermediate features in these two branches. Moreover, it will be seen from the experimental results in Section V that weight sharing actually does not jeopardize the interpretability and device independence of the resulting design.

B. Successive Deep Inference Module

Although many existing VSR methods exploit temporal redundancy, the underlying alignment and fusion rules are often heuristic in nature. In contrast, we shall propose a systematic approach based on a canonical decomposition result for HMM inference.

First note that the VSR problem can be mathematically formulated as

$$\tilde{X}^t \triangleq \min_f \mathbb{E}[d(X^t - f(Y^{t-3}, \dots, Y^{t+3}))], \quad (3)$$

where $(Y^{t-3}, \dots, Y^{t+3}) \triangleq (Y_{raw}^{t-3}, \dots, Y_{raw}^{t+3})$, $X^t \triangleq X_{rgb}^t$, $\tilde{X}^t \triangleq X_{sr}^t$, and $d(\cdot, \cdot)$ is the loss function. However, \tilde{X}^t in general depends on the choice of $d(\cdot, \cdot)$. To gain insight into the fundamental architectural principle that is applicable under any loss function, it is instructive to consider the following alternative formulation with hard reconstruction \tilde{X}^t replaced by soft reconstruction

$$\tilde{\Theta}^t \triangleq p(X^t | Y^{t-3}, \dots, Y^{t+3}), \quad (4)$$

where $p(X^t | Y^{t-3}, \dots, Y^{t+3})$ denotes the conditional distribution of X^t given $(Y^{t-3}, \dots, Y^{t+3})$. The SDI module is designed to solve (4), *i.e.*, to compute $\tilde{\Theta}^t$ based on $(Y^{t-3}, \dots, Y^{t+3})$. Note that $\tilde{\Theta}^t$ is a minimal sufficient statistic of $(Y^{t-3}, \dots, Y^{t+3})$ with respect to X^t in the sense that $(Y^{t-3}, \dots, Y^{t+3})$ and X^t are conditionally independent given $\tilde{\Theta}^t$ and $\tilde{\Theta}^t$ is a function of any statistic of $(Y^{t-3}, \dots, Y^{t+3})$ with this conditional independence property. As a consequence, it is possible to deduce \tilde{X}^t from $\tilde{\Theta}^t$ once the loss function is specified; for example, \tilde{X}^t is simply the mean of $\tilde{\Theta}^t$ (*i.e.*, $\tilde{X}^t = \mathbb{E}[X^t | Y^{t-3}, \dots, Y^{t+3}]$) if the loss function is chosen to be Mean Squared Error (MSE). This deduction is basically accomplished by the reconstruction module to be described in Section IV-C. It is also worth noting that (4) is invariant under isomorphic representations of $(Y^{t-3}, \dots, Y^{t+3})$. This is the reason why the features extracted from the LR raw frames can be used in lieu of the LR frames themselves as the input of the SDI module. Similarly, the output of the SDI module can be any equivalent representation of $\tilde{\Theta}^t$.

However, solving (4) through a data-driven approach is very difficult. For one thing, it requires learning joint patterns existent in the (features of) LR raw frames. Presumably the number of such patterns is already very large even for 2 frames and the number is likely to increase by several orders of magnitude with every additional frame. Therefore, it is unrealistic to expect that one can learn enough joint patterns needed to solve (4) reliably based on limited training data. Fortunately, it turns out that the learning complexity can be greatly reduced under a proper fusion strategy. For ease of exposition, we assume that the HR raw frames $X^{t-3} \leftrightarrow \dots \leftrightarrow X^{t+3}$ form a Markov chain in this order. Since each LR raw frame (or its feature) is generated from the corresponding HR raw frame through an independent degradation process, $(X^{t-3}, \dots, X^{t+3})$ and $(Y^{t-3}, \dots, Y^{t+3})$ jointly constitute an HMM. This HMM enables us to decompose (4) into a sequence of simpler problems. Note that $\Theta^{t-2} \triangleq p(X^{t-2} | Y^{t-3}, Y^{t-2})$ is a minimal sufficient statistic of (Y^{t-3}, Y^{t-2}) with respect to X^{t-2} . Therefore,

$$\tilde{\Theta}^t = p(X^t | \Theta^{t-2}, Y^{t-1}, \dots, Y^{t+3}). \quad (5)$$

Moreover, removing X^{t-3} and replacing (Y^{t-3}, Y^{t-2}) with Θ^{t-2} resulting in a new HMM. One can iterate the above argument to show that

$$\tilde{\Theta}^t = p(X^t | \Theta^t, \Theta^{t+1}), \quad (6)$$

where $\Theta^{t+\ell} \triangleq p(X^{t+\ell} | \Theta^{t+\ell-1}, Y^{t+\ell})$ for $\ell = -2, -1, 0$, and $\Theta^{t+\ell} \triangleq p(X^{t+\ell} | Y^{t+\ell}, \Theta^{t+\ell+1})$ for $\ell = 1, 2$, with $\Theta^{t-3} \triangleq Y^{t-3}$ and $\Theta^{t+3} \triangleq Y^{t+3}$. Fig. 4 provides an intuitive illustration of this iterative argument using probabilistic graphical models. It is worth mentioning that this is similar to the reasoning underlying the well-known forward-backward algorithm for HMM inference. The main difference is that here we are mostly concerned with the fundamental architectural principle rather than the computational complexity aspect.

Our fusion strategy has several advantages over the existing ones [16], [48]. First, an important implication of (6) is that (4) can be solved by repeatedly performing pairwise fusion, which

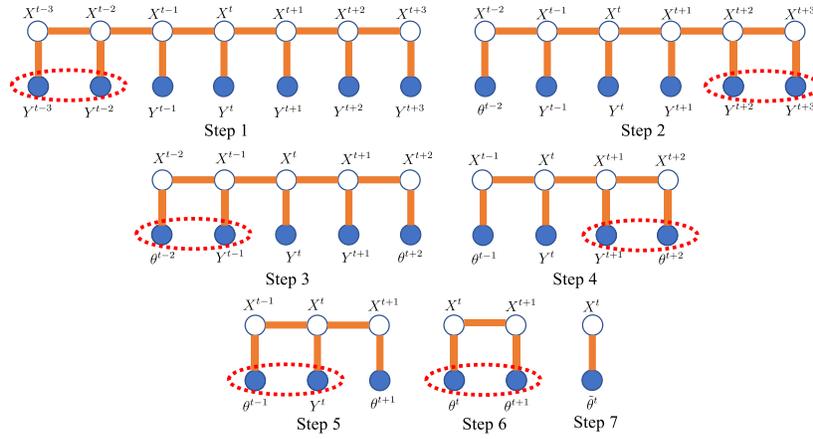


Fig. 4. A graphical illustration of HMM inference via iterative pairwise fusion.

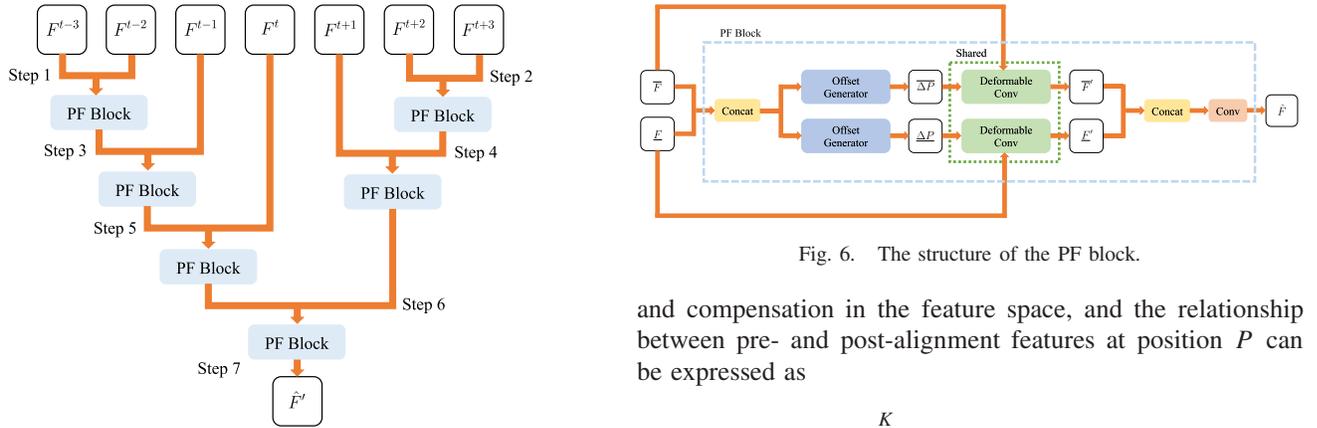


Fig. 5. The architecture of the SDI module.

is much less demanding in terms of the amount of training data needed for reliable learning. Second, the two inputs are of comparable importance with respect to the target object and consequently are more likely to be thoroughly exploited. In contrast, for other fusion strategies where the inputs under consideration have a clear difference in importance (for example, Y^t is clearly more relevant than Y^{t-3} for the purpose of estimating X^t), the subtle details contained in the less important ones may easily get overlooked.

As shown in Fig. 5, the SDI module is designed according to the architectural principle suggested by (6) (note that the slight asymmetry in (6) also manifests in the structure of this module). It takes the features F^{t-3}, \dots, F^{t+3} extracted respectively from the LR raw frames $Y_{raw}^{t-3}, \dots, Y_{raw}^{t+3}$ and outputs a fused feature \hat{F}^t . Each pairwise fusion (PF) step is realized by a PF block (see Fig. 6 for its structure). Specifically, the PF block first concatenates its two input features \bar{F} and \underline{E} ; it then uses two structurally-identical offset generators (without weight-sharing) to learn positional offsets $\bar{\Delta P} \triangleq (\bar{\Delta P}_1, \dots, \bar{\Delta P}_K)$ and $\underline{\Delta P} \triangleq (\underline{\Delta P}_1, \dots, \underline{\Delta P}_K)$, which are leveraged to align \bar{F} and \underline{E} using the modulated deformable convolution [49]; finally, the aligned features \bar{F}' and \underline{E}' are concatenated and passed through two convolutional layers to produce the fusion result \hat{F} . Note that the alignment procedure can be interpreted as implicit motion estimation

and compensation in the feature space, and the relationship between pre- and post-alignment features at position P can be expressed as

$$\bar{F}'(P) = \sum_{i=1}^K \bar{w}_i \cdot \bar{F}(P + \bar{\Delta P}_i) \cdot \bar{\Delta S}_i, \quad (7)$$

$$\underline{E}'(P) = \sum_{i=1}^K \underline{w}_i \cdot \underline{E}(P + \underline{\Delta P}_i) \cdot \underline{\Delta S}_i, \quad (8)$$

where \bar{w}_i (\underline{w}_i) and $\bar{\Delta S}_i$ ($\underline{\Delta S}_i$) denote the weight and the modulation scalar of the deformable convolution for $\bar{F}'(P)$ ($\underline{E}'(P)$) with respect to offset $\bar{\Delta P}_i$ ($\underline{\Delta P}_i$), $i = 1, \dots, K$ (we set $K = 9$ in this work). It is worth noting that different from the flow-based method, the proposed SDI module does not require any pre-training and can be embedded into another network for end-to-end training. We would also like to point out that the first-order Markov assumption for HR raw frames is not required for the derivation of the architecture in Fig. 5. Indeed, a higher-order Markov chain (which is more suitable to model scenarios with occlusion and disocclusion) can be reduced to a first-order one by grouping several consecutive states in a sliding-window fashion, after which our previous argument can be invoked with no essential change. In fact, the only constraining factor for the SDI module is the dimension of the output of each PF block in the sense that higher dimensional outputs are likely needed to accommodate more sophisticated Markov structures.

C. Reconstruction Module

The reconstruction module is elaborately designed with two tasks in mind: 1) refining the fused feature associated with the

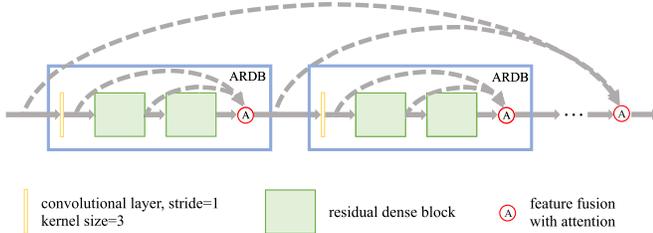


Fig. 7. The architecture of the reconstruction module.

input LR raw frames and 2) learning the ISP-specific information needed for color transformation and correction from the color reference frame. It is built using ARDBs as shown in Fig. 7. Each ARDB consists of one convolutional layer and two residual dense blocks. We follow the default RDB settings in [42] except that the growth rate is set to 64. The generated intermediate feature maps are fused under the guidance of the channel-wise attention. To this end, we adopt the SENet in [50] to produce attention weights to enhance/suppress the contributions of the associated feature maps in accordance with their relevance. In addition to the local fusion in each ARDB, we also employ a global fusion with channel-wise attention to further strengthen the learning ability of the reconstruction module. To strike a balance between system performance and computational complexity, totally four ARDBs are used in our design.

D. On the Practicability of Raw VSR

Although the present work mainly focuses on the technical feasibility of raw VSR, here we would like to briefly comment on its practical prospect. In particular, we shall address the following two questions: 1) the availability of camera raw data, and 2) the application scenarios of raw VSR. As raw data record untouched radiance information that can be fruitfully exploited for many different purposes, they receive considerable attention from professional photographers. To facilitate their use, the major camera manufacturers (*e.g.*, Canon and Nikon) have provided the official interface to access raw data. This removes a major hurdle for the wide adoption of raw VSR. Moreover, in most scenarios where conventional VSR is currently being considered, raw VSR can potentially be a more favorable choice to its superior performance and consistency with the real imaging pipeline. It is worth mentioning that some recent works [51], [52] attempt to generate raw-like data based on processed data. However, such raw-like data inevitably suffer from the information loss issue due to the non-invertibility of ISP operations, and great caution should be exercised when using them in lieu of authentic raw data such as those in our RawVD. More discussions can be found in Section V-D.

V. EXPERIMENTAL RESULTS

Extensive experiments are conducted to demonstrate the competitive performance of the proposed RawVSR. They also provide solid justifications for our design (especially, the SDI module and the reconstruction module) and convincing evidence regarding the value of raw data.

A. Implementation

The proposed RawVSR is end-to-end trainable without any additional supervision or pre-training for sub-modules. We adopt RawVD for training and testing as it is the only available video dataset for raw-data-based VSR methods. The scale ratio S is set to 4; the size of the defocus blur is randomly chosen from $\{3, 5, 7, 9, 11\}$; the two parameters σ_1^2 and σ_2^2 of heteroscedastic Gaussian noise are randomly sampled from $[0, 0.1]$ and $[0, 0.02]$, respectively. We use patches of size 64×64 cropped from LR raw videos as the training inputs, and augment the training data through inverting the frame order, horizontal flips and rotations [53]. Following [54], batch normalization is not used in RawVSR. Without bells and whistles, we adopt a weighted version of Mean Square Error (MSE) and Structural SIMilarity index (SSIM) [55] as our loss function to measure the difference between the final reconstructed SR frame and the ground truth. Furthermore, we provide an additional supervision to the provisional SR frame to enforce the learning of texture detail using an extra SSIM loss. The total loss function L can be formulated as:

$$L = L_F^{MSE} + \lambda_F L_F^{SSIM} + \lambda_P L_P^{SSIM}, \quad (9)$$

where L_F^{MSE} , L_F^{SSIM} , and L_P^{SSIM} are respectively the MSE loss, the SSIM loss from the final SR frame, and the SSIM loss from the provisional SR frame. The trade-off parameters λ_F and λ_P are both set to 0.005 to balance the value ranges of these sub-level losses. Since the SSIM loss is designed to measure the structural similarity between two images and is relatively insensitive to color distortions, two images photographed in the same scene under different light conditions might have a high SSIM value but a low PSNR value. This fact inspires us to employ SSIM for assessing the quality of the provisional SR frame. The textural consistency between the provisional SR frame and the ground truth then enables the spatial-specific transformation to focus on color correction. In this way, the texture restoration branch and the color correction branch are effectively functionally disentangled. It will be seen in Section V-C that the additional supervision on the provisional SR frame improves the interpretability of our network and makes it more compatible with the real imaging pipeline, where the super-resolution process is carried out before color correction. We would point out that the linear measurement X_{lin} provided in RawVD is not employed as the supervision reference of the provisional SR frame. This is because the proposed RawVSR is intended to be device-independent, whereas the use of X_{lin} might jeopardize its generalization ability due to the fact that X_{lin} is produced by a specific ISP (which is Rawpy in the current setting).

To accelerate network training, the Adam optimizer [56] is used with a batch size of 30 and the default parameter values $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We set the initial learning rate to $2e-4$, which is reduced by half every 20 epochs until a total of 100 epochs is reached. The training process is carried out on a PC with two NVIDIA GTX 1080Ti. All PSNR/SSIM values are evaluated using the test data in RawVD.

TABLE I
QUANTITATIVE COMPARISONS ON THE RAWVD TEST DATA FOR $2\times$ AND $4\times$ VSR. **RED** AND **BLUE** INDICATE THE BEST AND THE SECOND BEST PERFORMANCE, RESPECTIVELY

RawVD $2\times$										
Name	Bicubic	VSRNet [11]	VESPCN [33]	SPMC [34]	DUF [14]	RBPN [15]	TDAN [39]	EDVR [16]	RawVSR [†]	RawVSR
Store	27.84/0.8299	28.99/0.8457	29.02/0.8494	29.41/0.8503	30.75/0.8710	-	31.27/0.8863	31.98/0.9060	33.04/0.9129	34.30/0.9263
Painting	28.80/0.8086	29.38/0.8225	29.45/0.8218	28.82/0.7998	31.22/0.8512	-	31.17/0.8556	32.55/0.8816	32.80/0.8773	33.79/0.8943
Train	28.11/0.7738	29.18/0.8081	29.32/0.8121	30.01/0.8298	31.25/0.8464	-	30.94/0.8396	31.42/0.8617	31.66/0.8522	32.79/0.8729
City	27.76/0.7527	28.99/0.8012	29.22/0.8047	29.44/0.8210	30.47/0.8190	-	30.26/0.8227	31.11/0.8517	31.41/0.8588	32.11/0.8637
Walk	26.43/0.7553	28.92/0.7871	28.78/0.7820	30.25/0.8439	30.13/0.8117	-	29.26/0.8207	30.58/0.8415	29.63/0.8282	31.20/0.8516
Average	27.79/0.7841	29.09/0.8129	29.16/0.8140	29.59/0.8290	30.76/0.8399	-	30.58/0.8450	31.53/0.8685	31.71/0.8659	32.84/0.8818
RawVD $4\times$										
Name	Bicubic	VSRNet [11]	VESPCN [33]	SPMC [34]	DUF [14]	RBPN [15]	TDAN [39]	EDVR [16]	RawVSR [†]	RawVSR
Store	22.59/0.6507	25.24/0.7135	25.46/0.7250	26.34/0.7351	27.84/0.8105	27.53/0.7647	26.67/0.7802	28.23/0.8202	28.29/0.8191	29.04/0.8400
Painting	24.78/0.6712	26.17/0.7119	26.02/0.7120	26.13/0.7199	27.71/0.7833	26.40/0.7209	26.75/0.7539	27.74/0.7908	28.51/0.7940	29.02/0.8104
Train	24.12/0.6166	25.80/0.6671	25.96/0.6724	26.70/0.6984	27.53/0.7377	27.64/0.7358	26.94/0.7205	27.81/0.7470	28.02/0.7443	28.59/0.7625
City	24.94/0.6330	26.59/0.6907	26.63/0.6927	27.03/0.6684	27.93/0.7508	28.39/0.7332	27.55/0.7336	28.24/0.7558	28.43/0.7640	29.08/0.7843
Walk	22.46/0.6229	25.52/0.6766	25.62/0.6726	25.30/0.6797	27.30/0.7512	26.46/0.7337	26.36/0.7229	27.02/0.7533	27.90/0.7607	28.06/0.7724
Average	23.78/0.6389	25.87/0.6919	25.94/0.6950	26.30/0.7002	27.66/0.7667	27.29/0.7377	26.85/0.7422	27.81/0.7734	28.23/0.7765	28.76/0.7939

B. Quantitative and Qualitative Comparisons

We compare the proposed RawVSR with several existing processed-data-based VSR methods, including VSRNet [11], VESPCN [33], SPMC [34], DUF [14], RBPN [15], TDAN [39], and EDVR [16]. In particular, EDVR is the champion of the NTIRE 2019 challenge on video deblurring and super-resolution [40] and can be considered as the current state-of-the-art. For fair comparisons, we laboriously retrain all the aforementioned methods using the same strategy as described in Section V-A until convergence. These methods are trained with (Y_{rgb}, X_{rgb}) pairs instead of (Y_{raw}, X_{rgb}) pairs (see Fig. 2) since they are incapable of handling raw data. In contrast, the proposed RawVSR can leverage both raw data and processed data for training. To distinguish it from the original version of RawVSR for which the training is based on (Y_{raw}, X_{rgb}) pairs, the version trained with (Y_{rgb}, X_{rgb}) pairs is denoted as RawVSR[†], where the color correction branch is removed since the color reference is not needed in this case. Quantitative comparisons on the RawVD test data for $2\times$ and $4\times$ VSR are shown in Table I (with the only exception of RBPN for which we are unable to implement the same training strategy on our PC for $2\times$ VSR due to its large model size). For qualitative comparisons, we only demonstrate in Fig. 8 $4\times$ VSR since it is more visually distinguishable than $2\times$ VSR. It can be seen that RawVSR[†] outperforms the other processed-data-based methods in terms of PSNR and SSIM metrics, and visually restores sharper edges and finer details in most cases (see, *e.g.*, the word *Country* in video “Painting”, and the balcony part of the building in video “City”), which provides solid justifications for our overall network design. As compared to RawVSR[†], the results of RawVSR are even more appealing, both quantitatively and qualitatively. This improvement provides convincing evidence regarding the benefits of raw data since it is clearly attributed to their richer information content. In addition, to validate that the proposed method is device-independent, three raw 4K videos (named “Tree”, “Door”, and “Flower”) filmed by Adobe Lightroom software³ on an iPhone 8p are used to complement our RawVD test data for $4\times$ VSR. The quantitative

and qualitative comparisons are illustrated in Table II and Fig. 9, respectively. The superior VSR results indicate that the proposed RawVSR, trained only on RawVD, continues to perform well on the data collected by other devices without fine-tuning, and thus possesses good generalization abilities. More evidence can be found in Section V-E.

C. Visualization of the Color Correction Process

To show that the texture restoration branch and the color correction branch indeed fulfill their designated roles, Fig. 10 visualizes how a provisional SR frame \hat{X}_{sr}^t undergoes the color correction process to produce the corresponding final SR frame X_{sr}^t . In particular, we highlight three representative pixels (*i.e.*, $\hat{X}_{sr}^t(800, 1650)$, $\hat{X}_{sr}^t(200, 1350)$, and $\hat{X}_{sr}^t(450, 350)$) extracted from a provisional SR frame in video “Painting” and their color-corrected counterparts (*i.e.*, $X_{sr}^t(800, 1650)$, $X_{sr}^t(200, 1350)$, and $X_{sr}^t(450, 350)$) together with the associated spatial-specific color transformation matrices. Clearly, although the provisional SR frame \hat{X}_{sr}^t suffers severe color distortion, it exhibits correct details and is texturally consistent with the final SR frame X_{sr}^t . This validates the effectiveness of additional SSIM-based supervision in guiding the texture restoration branch to accomplish its desired purpose. The color-corrected pixels are produced by applying the spatial-specific transformation matrices on their respective provisional pixels according to Equ. (2). It is worth noting that even though the pixels extracted from \hat{X}_{sr}^t are color-wise quite similar, their associated color transformation matrices are distinctively different, resulting in visually more distinguishable color-corrected pixels. This provides strong supporting evidence for the usefulness of the color correction branch.

D. Authentic Raw Data vs. ISP-Inverted Data

Several recent works [51], [52] have attempted to invert the camera ISP to produce raw-like data (which will be referred to as ISP-inverted data) from processed data and use them in lieu of authentic raw data for training. However, the resemblance between ISP-inverted data and raw data should not be overestimated. Indeed, such inversion can never be perfect due to the presence of non-invertible ISP components (*e.g.*, quantization); actually it may even cause additional information loss. To gain

³<https://www.adobe.com/ca/products/photoshop-lightroom.html>

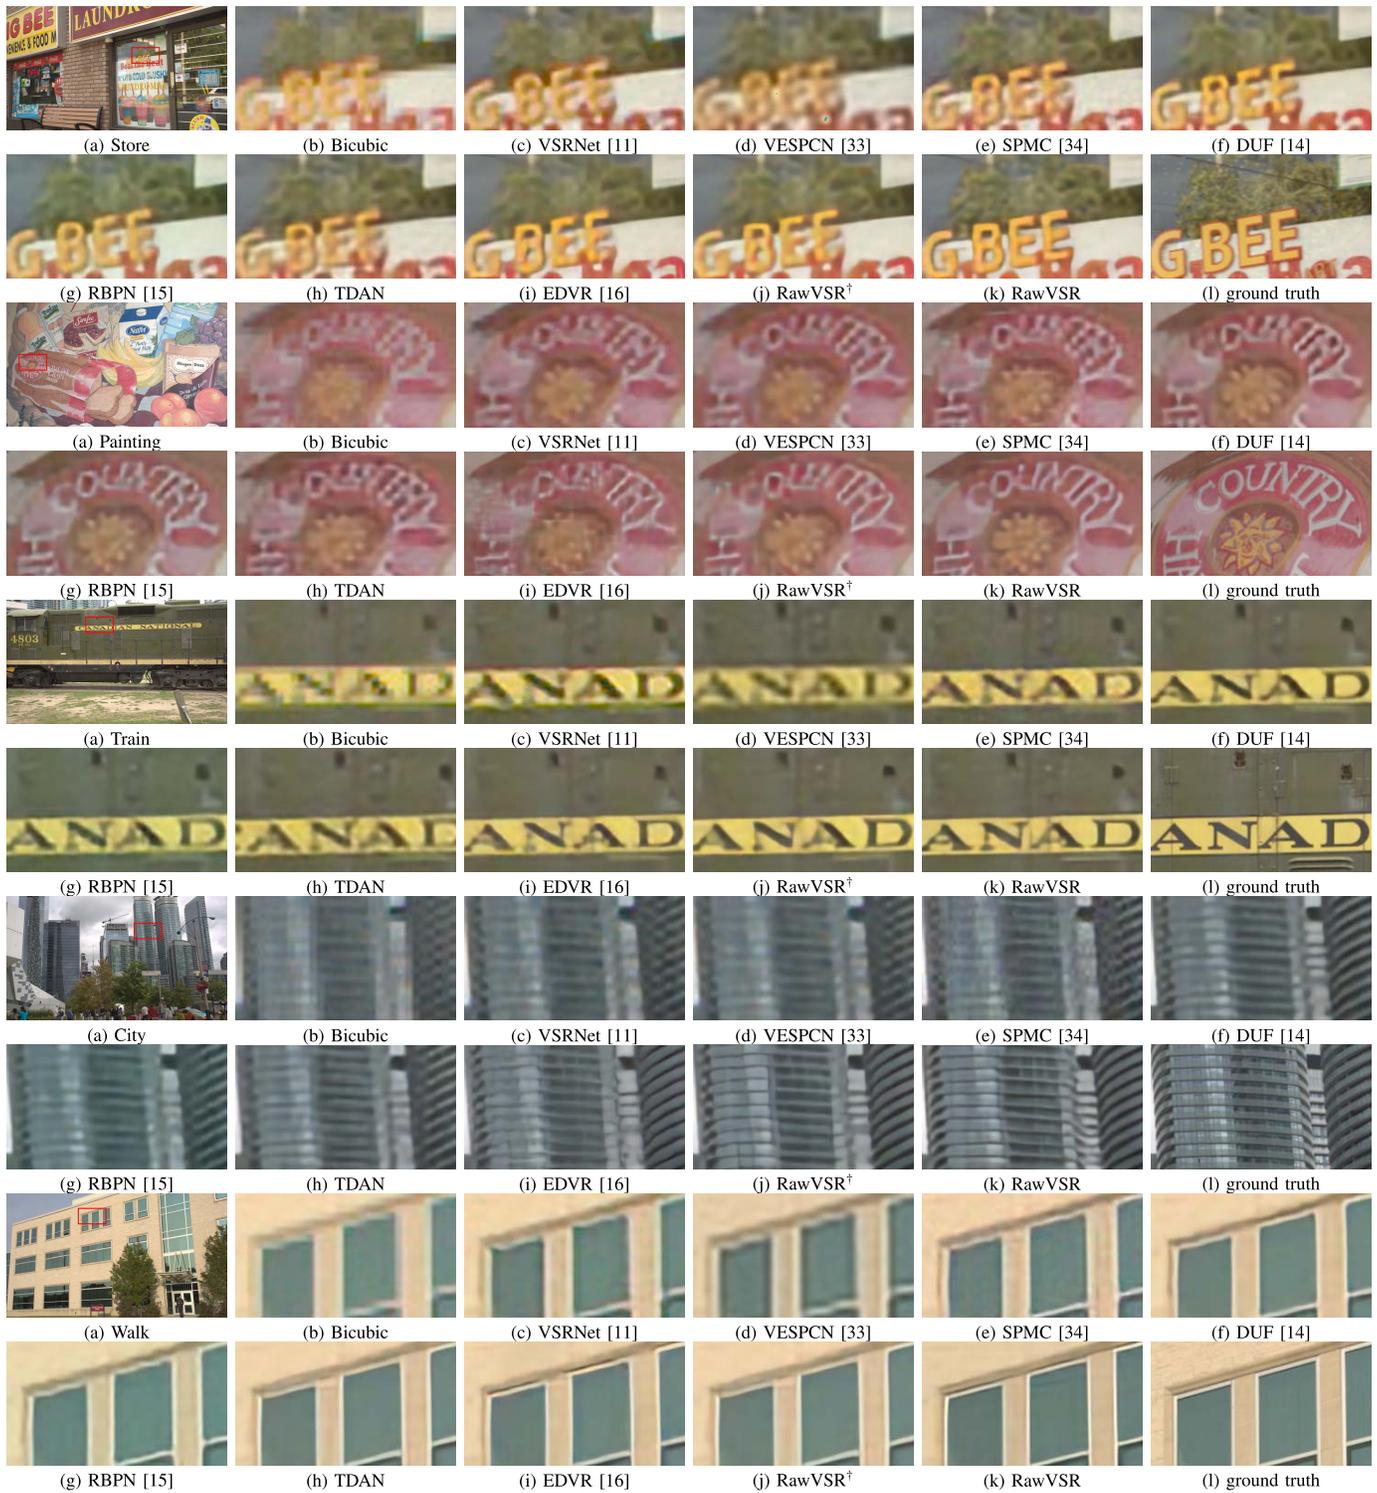


Fig. 8. Qualitative comparisons on the RawVD test data for 4 \times VSR. Zoom in for better visualization.

a concrete understanding, we follow the same procedure carried out in [51] to produce ISP-inverted data from processed data Y_{rgb} in RawVD, and input them to our network for testing. The results are then compared with those based on authentic raw data as shown in Fig. 11. It can be seen that ISP-inverted data incur two evident degradations: 1) significant color distortion (which is likely caused by the distribution mismatch between ISP-inverted data and raw data), 2) blurred texture detail (which is a consequence of information loss).

Therefore, authentic raw data as those collected in RawVD remain indispensable for raw VSR.

E. Validation of Device Independence

Due to the separation of the super-resolution process and the color correction process, the proposed RawVSR can adapt easily to different camera-ISPs. We have demonstrated this desirable characteristic in Table I and Fig. 9 by blindly testing RawVSR, trained with data from simulated Canon-ISP

TABLE II

QUANTITATIVE COMPARISONS ON THE IPHONE DATA FOR 4× VSR. **RED** AND **BLUE** INDICATE THE BEST AND THE SECOND BEST PERFORMANCE, RESPECTIVELY

iPhone 4x				
Name	DUF [14]	EDVR [16]	RawVSR [†]	RawVSR
Tree	27.15/0.6630	27.26/0.6504	<u>28.66/0.7073</u>	29.54/0.7381
Door	26.11/0.6334	26.05/0.6231	<u>26.28/0.6694</u>	28.62/0.6957
Flower	25.90/0.6225	25.74/0.6062	<u>27.07/0.6771</u>	28.17/0.7307
Average	26.38/0.6396	26.35/0.6266	<u>27.34/0.6846</u>	28.78/0.7215

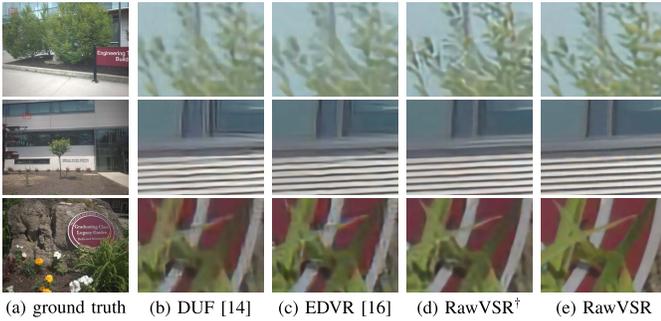


Fig. 9. Qualitative comparisons on the iPhone data for 4× VSR. The names of video sequences from the first row to the last row are “Tree”, “Door”, and “Flower”, respectively. Zoom in for better visualization.

(i.e., Rawpy), on a new iPhone dataset. Here we shall provide more evidence. Due to the lack of Digital Single-Lens Reflex (DSLR) cameras (other than Canon 5D3) at our disposal, we choose to generate LR color references by simulating ISPs of three major camera manufacturers (Nikon, Sony, and Fujifilm) based on the analysis of their color styles and preferences. According to our observation, the video frames produced by Nikon-ISP exhibit the lowest color temperature as compared to the other two, which leads to an overall cold tone; in contrast, those by Fujifilm-ISP show the highest color temperature, yielding a warm tone; Sony-ISP is somewhere in between in terms of color temperature, yet produces the highest luminance. Based on these insights, we can closely approximate the effects of these ISPs by suitably adjusting the exposure degree, contrast, and color temperature of LR color references in RawVD.

Specifically, we increase the color temperature from 6000K to 8000K to approximate the warm tone of Fujifilm-ISP, and decrease it to 5200K to approximate the cold tone of Nikon-ISP. Since the Sony-ISP has a unique preference, simply adjusting one color setting does not yield good approximation. Instead, we increase the exposure degree from the default value 0 to 0.7, decrease the contrast from 0 to -8 , and change the color temperature to 6500K in order to achieve the effect similar to that of Sony-ISP.

Fig. 12 shows the final SR frames with simulated Nikon, Sony, and Fujifilm-ISPs, respectively, and the corresponding provisional SR frames as well as their associated color histograms for a representative frame in video “Painting”. Here, the input to the texture restoration branch is fixed, and different color references (based on simulated ISPs) are fed into the color correction branch without additional training. Note that the choice of simulated ISP has no impact on the provisional

frame (see the 2nd row in Fig. 12) and its associated color histogram (see the 1st row in Fig. 12). As such, the texture restoration branch and the color correction branch are indeed functionally disentangled, which is beneficial to the interpretability of overall network design. On the other hand, the color of the final SR frame changes in accordance with the provided color reference (see the 3rd row in Fig. 12), which is also reflected in the color histograms (see the 4th row in Fig. 12). This shows convincingly that the proposed RawVSR can automatically adapt to the selected ISP, thus is device-independent.

F. Validation of Temporal Consistency

A lack of temporal consistency may lead to flickering artifacts (e.g., in the form of jagged edges) that can degrade the quality of reconstructed SR frames [35]. To compare the proposed RawVSR with the existing methods from this perspective, we extract some co-located columns from the reconstructed “City” sequence, rotate and then stack them vertically to produce a temporal profile [13]; we also generate another temporal profile based on co-located rows from the reconstructed “Walk” sequence. It can be seen from Fig. 13 that RawVSR yields the most consistent temporal profiles in terms of texture clarity, edge sharpness, and detail accuracy. For example, in the temporal profile of the “Walk” sequence, the proposed RawVSR successfully recovers most line patterns whereas the other methods fail to distinguish dashed lines from solid ones.

G. Ablation Studies

To gain a better understanding of key constituents of the proposed network as well as the influence of the input sequence length, we conduct elaborately designed ablation studies by examining various alternative implementations. All such implementations are trained from scratch with the same training strategy described in Section V-A. The proposed RawVSR (trained on raw data) will be referred to as the baseline.

1) *Justification of the SDI Module*: Most existing VSR methods adopt early fusion, slow fusion, or 3D convolutional fusion to merge input frames [33]. For early fusion, all input frames are concatenated at the very beginning and fed jointly into the network. As an example, VSRNet [11] follows this fusion strategy. In contrast, the slow fusion strategy merges input frames progressively according to a prescribed rule. The scheme of the proposed SDI module falls into this general category; one of our main contributions is, in a certain sense, an optimal fusion rule deduced via HMM analysis. 3D convolutional fusion performs convolution on input frames in both temporal and spatial domains, and tends to be computationally expensive. To justify the design of the SDI module, three variants are considered, each associated with a different fusion scheme. The first variant, named w/ EF, simply follows the early fusion strategy. The second variant, named w/ improved EF, is similar to the first one except that all other frames are first aligned to the middle frame. The last one, named w/ SF, adopts a slow fusion scheme, where all other frames are first aligned and fused with the middle frame in a pairwise manner

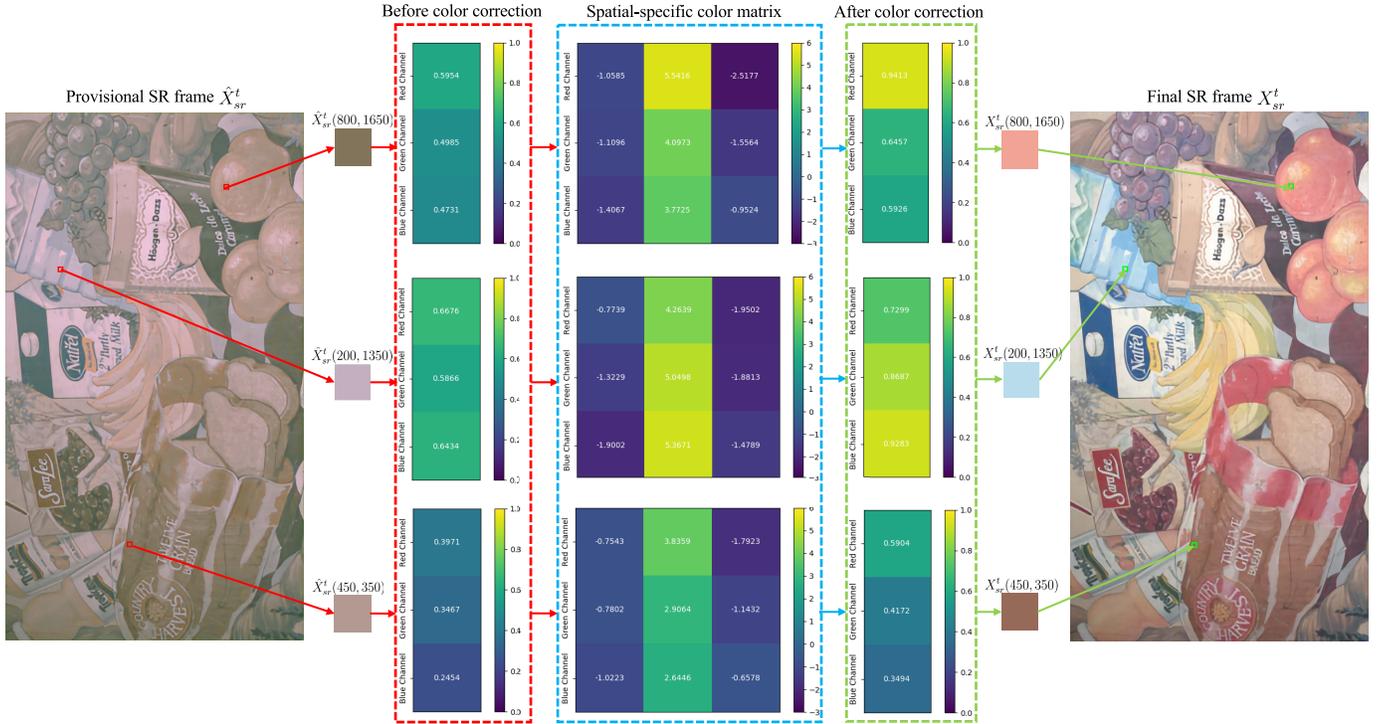
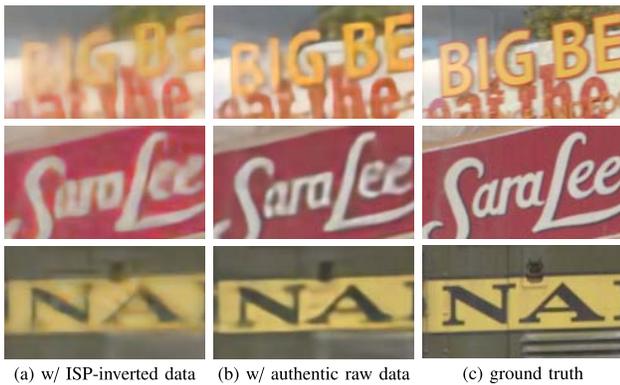


Fig. 10. Visualization of the color correction process (zoom-in for details).


 Fig. 11. The proposed RawVSR with ISP-inverted data and authentic raw data respectively on the RawVD test dataset for 4 \times VSR.

before joint fusion is performed. For fair comparisons, these alternative fusion modules are designed to have roughly the same size as that of the SDI module (no structural change is made to the rest of the network). Note that 3D convolutional fusion can not be implemented under this size constraint as it requires significantly more parameters. The comparison results are shown in Table III. It can be seen that the baseline performs favorably in terms of PSNR and SSIM metrics, which provides strong evidence in support of our design of the SDI module.

2) *Justification of the Reconstruction Module:* To justify the design of the reconstruction module, two variants are considered for comparison. The first one (w/o attention) has all channel-wise feature fusion removed in order to show the effectiveness of the attention mechanism, and the second one (w/ RDB) replaces the reconstruction module with a module of similar size — a cascade of 8 RDBs to verify the strength of

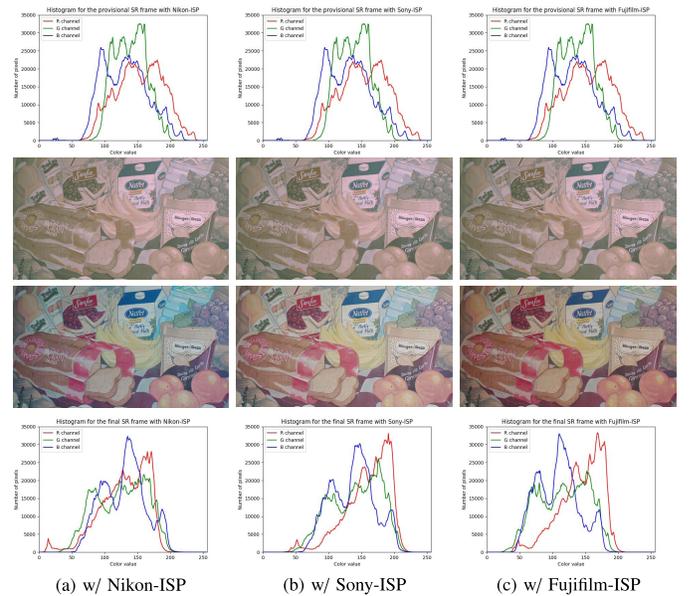


Fig. 12. Visualization of the provisional frames (in the 2nd row), the final SR frames (in the 3rd row), and their associated color histograms (in the 1st and 4th rows) with simulated Nikon, Sony and Fujifilm-ISPs, to validate the device-independence characteristic of RawVSR.

the reconstruction module in terms of dual-task performance. Quantitative comparisons in Table III provide strong evidence in favor of our original design.

3) *Influence of the Input Sequence Length:* Here we investigate how the overall VSR performance depends on the number of input frames. To this end, we consider the cases with 3 consecutive frames as inputs (w/ 3 frames), 5 consecutive frames as inputs (w/ 5 frames), and 7 consecutive frames as

TABLE III
ABLATION STUDIES FOR DIFFERENT VARIANTS FOR 4× VSR. **BOLD** INDICATES THE BEST PERFORMANCE

Ablation Variant	SDI			Reconstruction module		Input frame number		Baseline
	w/ EF	w/ improved EF	w/ SF	w/o attention	w/ RDB [42]	w/ 3 frames	w/ 5 frames	RawVSR
Store	28.00/0.8226	28.61/0.8304	28.82/0.8351	28.52/0.8303	28.48/0.8277	28.30/0.8176	28.75/0.8324	29.04/0.8400
Painting	28.05/0.7910	28.71/0.8007	28.91/0.8034	28.16/0.8006	28.00/0.7981	28.36/0.7927	28.59/0.8057	29.02/0.8104
Train	27.44/0.7481	28.08/0.7565	28.29/0.7556	28.07/0.7568	28.08/0.7555	28.01/0.7481	28.27/0.7554	28.59/0.7625
City	28.44/0.7699	28.84/0.7769	28.94/0.7775	28.73/0.7750	28.74/0.7743	28.64/0.7684	28.96/0.7776	29.08/0.7843
Walk	26.26/0.7525	27.78/0.7649	27.87/0.7628	26.69/0.7559	27.10/0.7567	27.95/0.7558	28.04/0.7695	28.06/0.7724
Average	27.64/0.7768	28.40/0.7859	28.57/0.7869	28.03/0.7837	28.08/0.7825	28.25/0.7765	28.52/0.7881	28.76/0.7939

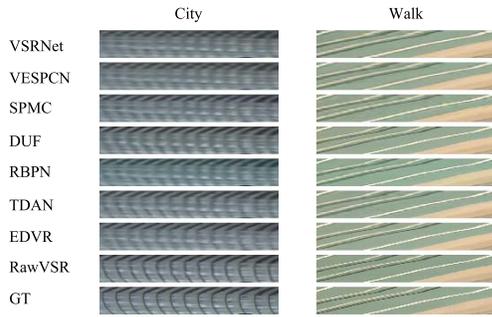


Fig. 13. Visualization of temporal profiles produced by the proposed RawVSR and some existing methods.

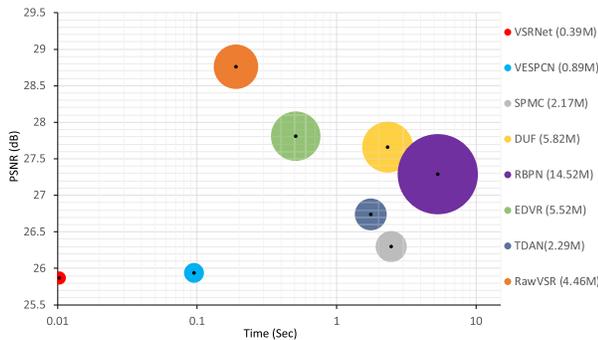


Fig. 14. Runtime and model size comparisons of different VSR methods for 4× VSR.

inputs (baseline). The corresponding quantitative results are shown in Table III. It is clear that the VSR performance improves as the number of input frames increases, which is expected since one can more likely find the missing information when there are more frames available for exploitation. It is also worth mentioning that due to the effective design of the SDI module, the proposed RawVSR can flexibly handle any input length without adjusting the network size (although increasing the number of input frames typically leads to longer processing time).

H. Runtime and Model Size Comparisons

Fig. 14 compares different VSR methods in terms of runtime and model size as well as PSNR values. Here runtime refers to the average time needed for producing one 1080p frame (without taking into account the data loading phase) for 4× VSR. It can be seen that the proposed RawVSR is able to achieve the highest PSNR value with moderate model size and near-real-time processing.

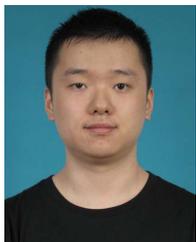
VI. CONCLUSION

We have proposed a new deep-learning-based method that can effectively exploit raw data for VSR. The associated neural network has a dual-branch structure that disentangles the super-resolution process from the color correction process; as a consequence, it better fits the real imaging pipeline and has the desirable property of being device-independent. Its SDI module is designed according to the architectural principle obtained via the analysis of a suitable probabilistic graphical model; this design philosophy is likely applicable to a wide range of problems, especially those involving data fusion. Its reconstruction module employs an attention mechanism to enhance the learning ability. Moreover, weight sharing is adopted to reduce the model size and to ensure the latent consistency of intermediate features in the two branches. Finally, it is also expected that the new raw video dataset collected in this paper can potentially benefit other video tasks including, among others, video interpolation and visual enhancement via exploiting raw data.

REFERENCES

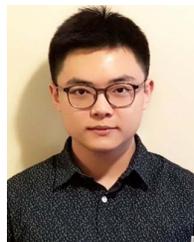
- [1] H. Greenspan, "Super-resolution in medical imaging," *Comput. J.*, vol. 52, no. 1, pp. 43–63, 2008.
- [2] D.-H. Trinh, M. Luong, F. Dibos, J.-M. Rocchisani, C.-D. Pham, and T. Q. Nguyen, "Novel example-based method for super-resolution and denoising of medical images," *IEEE Trans. Image Process.*, vol. 23, no. 4, pp. 1882–1895, Apr. 2014.
- [3] H. Demirel and G. Anbarjafari, "Discrete wavelet transform-based satellite image resolution enhancement," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 6, pp. 1997–2004, Jun. 2011.
- [4] F. Li, X. Jia, D. Fraser, and A. Lambert, "Super resolution for remote sensing images based on a universal hidden Markov tree model," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 3, pp. 1270–1278, Mar. 2010.
- [5] K. Jia and S. Gong, "Generalized face super-resolution," *IEEE Trans. Image Process.*, vol. 17, no. 6, pp. 873–886, Jun. 2008.
- [6] L. Zhang, H. Zhang, H. Shen, and P. Li, "A super-resolution reconstruction algorithm for surveillance images," *Signal Process.*, vol. 90, no. 3, pp. 848–859, Mar. 2010.
- [7] S. Hao, W. Wang, Y. Ye, E. Li, and L. Bruzzone, "A deep network architecture for Super-Resolution-Aided hyperspectral image classification with classwise loss," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 8, pp. 4650–4663, Aug. 2018.
- [8] P. H. Hennings-Yeomans, S. Baker, and B. V. K. V. Kumar, "Simultaneous super-resolution and feature extraction for recognition of low-resolution faces," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [9] B. Na and G. C. Fox, "Object detection by a super-resolution method and a convolutional neural networks," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2018, pp. 2263–2269.
- [10] R. Liao, X. Tao, R. Li, Z. Ma, and J. Jia, "Video super-resolution via deep draft-ensemble learning," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 531–539.
- [11] A. Kappeler, S. Yoo, Q. Dai, and A. K. Katsaggelos, "Video super-resolution with convolutional neural networks," *IEEE Trans. Comput. Imag.*, vol. 2, no. 2, pp. 109–122, Jun. 2016.

- [12] O. Shahar, A. Faktor, and M. Irani, "Space-time super-resolution from a single video," in *Proc. CVPR*, Jun. 2011, pp. 3353–3360.
- [13] M. S. M. Sajjadi, R. Vemulapalli, and M. Brown, "Frame-recurrent video super-resolution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6626–6634.
- [14] Y. Jo, S. W. Oh, J. Kang, and S. J. Kim, "Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3224–3232.
- [15] M. Haris, G. Shakhnarovich, and N. Ukita, "Recurrent back-projection network for video super-resolution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 3897–3906.
- [16] X. Wang, K. C. K. Chan, K. Yu, C. Dong, and C. C. Loy, "EDVR: Video restoration with enhanced deformable convolutional networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2019, pp. 1954–1963.
- [17] H. Jiang, Q. Tian, J. Farrell, and B. A. Wandell, "Learning the image processing pipeline," *IEEE Trans. Image Process.*, vol. 26, no. 10, pp. 5032–5042, Oct. 2017.
- [18] E. Schwartz, R. Giryes, and A. M. Bronstein, "DeepISP: Toward learning an end-to-end image processing pipeline," *IEEE Trans. Image Process.*, vol. 28, no. 2, pp. 912–923, Feb. 2019.
- [19] X. Xu, Y. Ma, and W. Sun, "Towards real scene super-resolution with raw images," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 1723–1731.
- [20] V. Bychkovsky, S. Paris, E. Chan, and F. Durand, "Learning photographic global tonal adjustment with a database of input/output image pairs," in *Proc. CVPR*, Jun. 2011, pp. 97–104.
- [21] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super resolution," *IEEE Trans. Image Process.*, vol. 13, no. 10, pp. 1327–1344, Oct. 2004.
- [22] C. Liu and D. Sun, "On Bayesian adaptive video super resolution," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 2, pp. 346–360, Feb. 2014.
- [23] Z. Ma, R. Liao, X. Tao, L. Xu, J. Jia, and E. Wu, "Handling motion blur in multi-frame super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 5224–5232.
- [24] T. Kohler, X. Huang, F. Schebesch, A. Aichert, A. Maier, and J. Hornegger, "Robust multiframe super-resolution employing iteratively re-weighted minimization," *IEEE Trans. Comput. Imag.*, vol. 2, no. 1, pp. 42–58, Mar. 2016.
- [25] X. Liu and J. Zhao, "Robust multi-frame super-resolution with adaptive norm choice and difference curvature based BTV regularization," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Nov. 2017, pp. 388–392.
- [26] X. Liu, L. Chen, W. Wang, and J. Zhao, "Robust multi-frame super-resolution based on spatially weighted half-quadratic estimation and adaptive BTV regularization," *IEEE Trans. Image Process.*, vol. 27, no. 10, pp. 4971–4986, Oct. 2018.
- [27] C. Chen, Q. Chen, J. Xu, and V. Koltun, "Learning to see in the dark," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3291–3300.
- [28] X. Liu, Y. Ma, Z. Shi, and J. Chen, "GridDehazeNet: Attention-based multi-scale network for image dehazing," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 7314–7323.
- [29] Y. Shao, L. Li, W. Ren, C. Gao, and N. Sang, "Domain adaptation for image dehazing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 2808–2817.
- [30] X. Xu, L. Siyao, W. Sun, Q. Yin, and M.-H. Yang, "Quadratic video interpolation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 1647–1656.
- [31] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 184–199.
- [32] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-C. Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 802–810.
- [33] J. Caballero *et al.*, "Real-time video super-resolution with spatio-temporal networks and motion compensation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, vol. 1, no. 2, p. 7.
- [34] X. Tao, H. Gao, R. Liao, J. Wang, and J. Jia, "Detail-revealing deep video super-resolution," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 22–29.
- [35] X. Liu, L. Kong, Y. Zhou, J. Zhao, and J. Chen, "End-to-end trainable video super-resolution based on a new mechanism for implicit motion estimation and compensation," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2020, pp. 2416–2425.
- [36] M. Drulea and S. Nedevschi, "Total variation regularization of local-global optical flow," in *Proc. 14th Int. IEEE Conf. Intell. Transp. Syst. (ITS)*, Oct. 2011, pp. 318–323.
- [37] W. Shi *et al.*, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1874–1883.
- [38] M. Haris, G. Shakhnarovich, and N. Ukita, "Deep back-projection networks for super-resolution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1664–1673.
- [39] Y. Tian, Y. Zhang, Y. Fu, and C. Xu, "TDAN: Temporally-deformable alignment network for video super-resolution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 3360–3369.
- [40] S. Nah *et al.*, "NTIRE 2019 challenge on video deblurring and super-resolution: Dataset and study," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2019, pp. 1996–2005.
- [41] A. Ignatov, L. Van Gool, and R. Timofte, "Replacing mobile camera ISP with a single deep learning model," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 536–537.
- [42] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, "Residual dense network for image super-resolution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2472–2481.
- [43] T. Xue, B. Chen, J. Wu, D. Wei, and W. T. Freeman, "Video enhancement with task-oriented flow," 2017, *arXiv:1711.09078*. [Online]. Available: <http://arxiv.org/abs/1711.09078>
- [44] K. Hirakawa and T. W. Parks, "Adaptive homogeneity-directed demosaicing algorithm," *IEEE Trans. Image Process.*, vol. 14, no. 3, pp. 360–369, Mar. 2005.
- [45] T. Plotz and S. Roth, "Benchmarking denoising algorithms with real photographs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1586–1595.
- [46] J. Cai, H. Zeng, H. Yong, Z. Cao, and L. Zhang, "Toward real-world single image super-resolution: A new benchmark and a new model," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 3086–3095.
- [47] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. Cham, Switzerland: Springer*, Oct. 2015, pp. 234–241.
- [48] W. Bao, W.-S. Lai, X. Zhang, Z. Gao, and M.-H. Yang, "MEMC-net: Motion estimation and motion compensation driven neural network for video interpolation and enhancement," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Sep. 17, 2019, doi: [10.1109/TPAMI.2019.2941941](https://doi.org/10.1109/TPAMI.2019.2941941).
- [49] X. Zhu, H. Hu, S. Lin, and J. Dai, "Deformable ConvNets V2: More deformable, better results," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9308–9316.
- [50] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7132–7141.
- [51] R. Zhou, R. Achanta, and S. Süsstrunk, "Deep residual network for joint demosaicing and super-resolution," in *Proc. Color Imag. Conf.*, vol. 2018, no. 1. Springfield, VA, USA: Society for Imaging Science and Technology, 2018, pp. 75–80.
- [52] T. Brooks, B. Mildenhall, T. Xue, J. Chen, D. Sharlet, and J. T. Barron, "Unprocessing images for learned raw denoising," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 11036–11045.
- [53] R. Timofte, R. Rothe, and L. Van Gool, "Seven ways to improve example-based single image super resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1865–1873.
- [54] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 136–144.
- [55] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [56] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <http://arxiv.org/abs/1412.6980>



Xiaohong Liu (Graduate Student Member, IEEE) received the B.E. degree in communication engineering from Southwest Jiaotong University, China, in 2014, and the M.A.Sc. degree in electrical and computer engineering from the University of Ottawa, Canada, in 2016. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, McMaster University, Canada.

His research interests include video super-resolution, image dehazing/deraining, and image forgery detection. He received the Ontario Graduate Scholarship in 2019, the NSERC Postgraduate Scholarship-Doctoral, and Borealis AI Global Fellowship Award in 2020. He serves as a reviewer for several journals, including IEEE TRANSACTIONS ON IMAGE PROCESSING, IEEE TRANSACTIONS ON MULTIMEDIA, IEEE SIGNAL PROCESSING LETTERS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, and IEEE TRANSACTIONS ON BROADCASTING.



Zhe Wang is currently pursuing the B.Math. degree with the David R. Cheriton School of Computer Science, University of Waterloo. His research interests include object tracking, video super-resolution, and medical image analysis.



Jun Chen (Senior Member, IEEE) received the B.E. degree in communication engineering from Shanghai Jiao Tong University, Shanghai, China, in 2001, and the M.S. and Ph.D. degrees in electrical and computer engineering from Cornell University, Ithaca, NY, USA, in 2004 and 2006, respectively.

He was a Postdoctoral Research Associate with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL, USA, from September 2005 to July 2006, and a Postdoctoral Fellow with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA, from July 2006 to August 2007. Since September 2007, he has been with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, ON, Canada, where he is currently a Professor. His research interests include information theory, machine learning, wireless communications, and signal processing.

Dr. Chen was a recipient of the Josef Raviv Memorial Postdoctoral Fellowship in 2006, the Early Researcher Award from the Province of Ontario in 2010, the IBM Faculty Award in 2010, the ICC Best Paper Award in 2020, and the JSPS Invitational Fellowship in 2021. He held the title of the Barber-Gennum Chair in Information Technology from 2008 to 2013, and the Joseph Ip Distinguished Engineering Fellow from 2016 to 2018. He served as an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY from 2014 to 2016. He is currently an Editor of the IEEE TRANSACTIONS ON GREEN COMMUNICATIONS AND NETWORKING.



Kangdi Shi received the B.E. degree in electrical and computer engineering from the University of Manitoba, Winnipeg, Canada, in 2017. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, Canada. His research areas include image classification, video super-resolution, and video coding.