

ARVo: Learning All-Range Volumetric Correspondence for Video Deblurring

Dongxu Li^{*,1,2}, Chenchen Xu^{*,1,2}, Kaihao Zhang^{*,†,1},
Xin Yu³, Yiran Zhong¹, Wenqi Ren⁴, Hanna Suominen^{1,5}, Hongdong Li¹
¹ANU, ²DATA61-CSIRO, ³UTS, ⁴IIE-CAS, ⁵University of Turku




Figure 1: Qualitative comparisons of different models on the newly introduced HFR-DVD dataset. The dataset features sharp frames captured with a high frame-rate at 1,000 fps, providing more realistic motion blurs and less artefacts. Our model matches pixel-pairs across all the spatial range in the feature space, thus having a better capacity to build correspondence with large pixel displacements. This makes our model more favorable to handle blurs, especially due to fast motions.

Abstract

Video deblurring models exploit consecutive frames to remove blurs from camera shakes and object motions. In order to utilize neighboring sharp patches, typical methods rely mainly on homography or optical flows to spatially align neighboring blurry frames. However, such explicit approaches are less effective in the presence of fast motions with large pixel displacements. In this work, we propose a novel implicit method to learn spatial correspondence among blurry frames in the feature space. To construct distant pixel correspondences, our model builds a correlation volume pyramid among all the pixel-pairs between neighboring frames. To enhance the features of the reference frame, we design a correlative aggregation module that maximizes the pixel-pair correlations with its neighbors based on the volume pyramid. Finally, we feed the aggregated features into a reconstruction module to obtain the restored frame. We design a generative adversarial paradigm to optimize the model progressively. Our proposed method is evaluated on the widely-adopted DVD dataset, along with

a newly collected High-Frame-Rate (1000 fps) Dataset for Video Deblurring (HFR-DVD). Quantitative and qualitative experiments show that our model performs favorably on both datasets against previous state-of-the-art methods, confirming the benefit of modeling all-range spatial correspondence for video deblurring.

1. Introduction

Hand-held devices are popular in capturing videos of dynamic scenes, where prevalent high-speed object movements and abrupt camera shakes lead to undesirable blurs in videos. To this end, the task of video deblurring aims to improve the video quality by restoring sharp frames from blurry video sequences and benefits a wide range of vision applications [18, 14, 13, 20, 19, 21, 53, 57].

Video deblurring methods remove blurs by taking advantage of sharper scene patches from neighboring frames [23, 3]. However, since neighboring frames are usually not spatially aligned with the reference frame due to motions, it is non-trivial to construct visual correspondence between frames. Existing works mainly tackle the misalignment between frames using homography or optical flows. Some

* Authors contributed equally.

† Corresponding author: kaihao.zhang@anu.edu.au.

previous works [36, 7, 5, 38, 45, 26] first estimate the relative motions using optical flows, then warp neighboring frames to the reference frame. However, in occurrence of fast object motions with large displacements, the constraint on the velocity smoothness may not always hold. As a result, estimating optical flow from blurry images remains a challenging research problem by itself [33, 39]. In addition, these methods are also notoriously less effective in the presence of occlusions and severe depth variations. Recent works [9, 56] use dynamic filters to restore videos in the feature domain. However, their ability to address large pixel displacement is restricted by the local receptive fields. Therefore, it remains a challenge to design a flexible yet effective visual correspondence method with abrupt motions.

To tackle this problem, as a complement to the existing pixel based warping methods, we propose an *implicit* approach that estimates the pairwise image correspondence in the feature space. In particular, we first extract visual features from a sequence of warped consecutive frames. Then, we pair up the reference frame and its neighboring frame, and compute their pixel-wise correlations by matching their features in the embedding space. Specifically, in order to account for distant pixel correspondence, we compute correlations for pixel pairs in all the spatial range. In this way, we enable the model to better capture the visual dependencies at various lengths, thus being more effective in handling blurs to fast motions.

To further increase the receptive fields while maintaining the fine-grained visual details, we propose to build a pyramid of correlation volumes in order to achieve feature matching at different spatial scales. Particularly, rather than subsampling all the frames to the same scale as in typical feature pyramid modules [6], we keep the reference frame feature maps while subsampling only those of neighboring frames. By maintaining the spatial scale of the reference frame, we keep high resolution information, allowing the model to learn correspondence taking into account the fine-grained visual details.

We optimize our model progressively in a generative adversarial paradigm with a new temporal consistency loss. The training proceeds in stages, where later stages take as input the restored frames from the previous stage. In this way, we ease the optimization by allowing the model to restore details gradually. We propose an adversarial loss to encourage the temporal consistency between the restored frames. Distinct from the previous adversarial video deblurring method [51] that takes as input a single restored frame, we use a discriminator to distinguish between restored frame sequence and the ground-truth sharp frame sequence. We show the two training strategies combined improve the restoration quality quantitatively and visually.

To validate the effectiveness of the proposed method, we first conduct evaluations on the widely-adopted DVD

dataset [36]. In addition, since the DVD dataset uses automatic frame interpolation to increase the frame rate, it exhibits artefacts in the synthetic blurry frames. To this end, we contribute a new large-scale dataset for video deblurring research, called *HFR-DVD*. The HFR-DVD dataset is captured using a high-speed camera in 1,000 fps, featuring sharper frames and more realistic motion blurs. Our experiments show that the proposed feature correlation methods help to improve the restoration on both datasets.

Our contributions are summarized as follows:

- we introduce a novel video deblurring method by constructing spatial correspondence between pixel pairs in the feature space. To account for distant pixel displacements, we match pixel pairs in all the spatial range between the reference frame and neighboring frames;
- we propose a correlative aggregation module to enhance the reference frame feature based on the correlation volume pyramid;
- in order to encourage the temporal consistency in the restored frames, we develop a adversarial loss for optimizing the model;
- we benchmark existing video deblurring methods on a new large-scale high-frame-rate dataset HFR-DVD with sharper frames and more realistic blurs;
- our video deblurring model ARVo outperforms previous methods quantitatively and qualitatively on the DVD and HFR-DVD datasets, establishing a new state-of-the-art for video deblurring.

2. Related Works

We approach the video deblurring task by constructing dense correspondence between features. Our work is thus broadly related to image deblurring video deblurring, and dense correspondence matching.

Image deblurring. Early works on image deblurring assume a uniform blur kernel and mainly focus on designing effective image priors [35, 15, 27, 46, 31, 28] to regularize the deblurring process into a well-posed problem. Recent deep learning models rely on paired sharp and blurry images for supervised training. In particular, the methods [38, 5, 50] learn to model the non-uniform blur kernels. Tao *et al.* [40] design a recurrent structure to iteratively refine the deblurring results. Zhang *et al.* [49] propose a spatially invariant recurrent network to remove blurs in the feature space. Gao *et al.* [4] propose a nested skip connection structure with selective parameter sharing. These models show promising deblurring results on the synthetic blurry images. The work [52] attempts to model the realistic blurs from real-life blurred images. Most learning-based works are trained with the typical reconstruction loss, *e.g.* L_1 or L_2 losses, while [16] uses an additional adversarial loss.




Figure 2: Overview of the proposed video deblurring pipeline. Our method first warps the neighboring frame to the reference frame for feature extraction. Then it computes all-range volumetric correspondence between the reference and neighboring frames as a correlation volume pyramid. Next, it relies on a correlative aggregation module to enhance the reference frame feature maps. Finally, we obtain the restored frame from a reconstruction network. Note that we only display one neighboring frame for illustration purposes, while our model in facts computes one intra-frame and two inter-frame correlation volumes.

Video deblurring. Early works on video deblurring mostly attempt to restore images by estimating blur kernels [22, 48, 47, 7, 29, 30, 50]. Recently, Su *et al.* [36] align neighboring frames using homography and optical flows and use a CNN to restore frames. Kim *et al.* [8] employ a recurrent neural network architecture and propose a dynamic temporal blending technique to enforce temporal consistency among frames. Instead of explicit alignment, Zhou *et al.* [56] use dynamic filters to adaptively align frame features. However, their method is restricted by the limited local receptive field and is less effective for fast motions. The works [45, 26] estimate sharpness as a prior and utilize sharp images as exemplars to aid the learning of deblurring networks. Zhang *et al.* [51] use adversarial loss to improve visual quality of estimated sharp images. However, their discriminator only considers a single output frame, thus neglecting temporal constraints. Our training paradigm encourages temporal consistency among multiple output frames, thus regularizing the model to produce temporally-consistent deblurred results.

Dense correspondence matching. Many vision tasks can be considered as a correspondence matching problem, *e.g.* stereo matching [54, 2, 1] and flow estimation [55, 42], where the essence is the correlation measures for paired images. In particular, Nikolaus *et al.* [24] directly regress disparity maps from an 1D correlation volume. Alex *et al.* [11] concatenate features of image frames to form the correlation volume. The work [34] learns to fuse multiple disparity maps to provide a more accurate correlation volume. The recent work on optical flow estimation [41] adopts the similar idea but constructs a multi-scale correlation volume. In their model, coarser correlation volumes are constructed by pooling operations over finer correlation volumes. In this way, their models capture distant pixel displacements while preserving fine-grained image details.

3. Approach

Our model takes as input three consecutive frames from a blurred video and restores the sharp frame in the middle. In Fig. 2, we provide an overview of our approach, which consists mainly of three phases: (1) feature extraction; (2) construction of visual correspondence; and (3) sharp frame restoration. All the phases are differentiable and jointly optimized in an end-to-end trainable architecture.

3.1. Frame Feature Extraction

Upon receiving the inputs, our model first applies explicit alignment by warping the neighboring frames towards the reference frame based on the optical flow estimation, as in previous works [36, 26]. In the following, we take the reference frame \mathbf{I}_i and its immediate following frame \mathbf{I}_{i+1} as an example for illustration. The optical flow estimation module \mathcal{O} takes \mathbf{I}_i and \mathbf{I}_{i+1} and predicts a optical flow field:

$$\mathbf{u}_{i,i+1}^x, \mathbf{u}_{i,i+1}^y = \mathcal{O}(\mathbf{I}_i, \mathbf{I}_{i+1}), \quad (1)$$

where $\mathbf{u}_{i,i+1}^x$ and $\mathbf{u}_{i,i+1}^y$ denote the x , y components of the estimated flow field, respectively.

The flow field maps each pixel (x, y) in \mathbf{I}_{i+1} to its corresponding coordinates $(x', y') = (x + \mathbf{u}_{i,i+1}^x(x), y + \mathbf{u}_{i,i+1}^y(y))$ in \mathbf{I}_i . We then obtain the aligned neighboring frame $\hat{\mathbf{I}}_{i+1}$ by a bilinear interpolation on \mathbf{I}_{i+1} using the mapped coordinates. We repeat the explicit alignment step for the preceding frame, and acquire the sequence of warped input frames $\{\hat{\mathbf{I}}_{i-1}, \mathbf{I}_i, \hat{\mathbf{I}}_{i+1}\}$.

In order to acquire the visual features of aligned frames, we apply an encoder network \mathcal{F}_θ with trainable parameters θ to $\{\hat{\mathbf{I}}_{i-1}, \mathbf{I}_i, \hat{\mathbf{I}}_{i+1}\}$, respectively, and obtain their feature maps with C channels at a reduced resolution. Namely, the encoder is a mapping $\mathcal{F}_\theta : \mathbb{R}^{H \times W \times 3} \rightarrow$

$\mathbb{R}^{H' \times W' \times C}$. In our implementation, the encoder \mathcal{F}_θ consists of two stride convolutional layers for downsampling the feature maps, and 15 residual blocks, where five are applied at the full resolution, five at the half resolution and five at the quarter resolution. The output of the frame feature extraction phase is a sequence of frame features $\{\mathcal{F}_\theta(\hat{\mathbf{I}}_{i-1}), \mathcal{F}_\theta(\mathbf{I}_i), \mathcal{F}_\theta(\hat{\mathbf{I}}_{i+1})\}$, where the feature has the dimension $H/4 \times W/4 \times C$.

3.2. All-Range Correlation Volume Pyramid

The explicit alignment method by optical flow builds spatial correspondence between neighboring frames. However, there remain two major drawbacks in such design that might adversely impact the alignment quality. Firstly, when computing optical flows, the constraint on the velocity smoothness may not be satisfied in the presence of abrupt motions. Secondly, each blurry image is an integration of multi-frame instance and sharp snapshots with a small motion, but flow fields describe only inter-frame motions while do not capture motions occurring during the exposure time. Consequently, the optical flow estimation may not be accurate enough to support effective alignment, particularly in dynamic scenes with fast object motions.

Motivated by this observation, we propose a method that is robust to identify visual correspondence with distant pixel displacements in order to mitigate the issue of fast motions. Such a deblurring method is also expected to handle intra-frame motion blurs happening during the exposure. In the following, we present the novel implicit method to construct spatial correspondence by measuring the visual similarity between each frame pair.

All-Range Correlation Volume. Convolution kernels operate in a local spatial neighborhood, which limits its ability to capture long-range correspondence. Inspired by the recent work in dense correspondence matching [41], we compute visual correspondence by considering pixel features in all the spatial range at a time. Given the features of a frame pair, e.g. $\mathcal{F}_\theta(\mathbf{I}_i), \mathcal{F}_\theta(\hat{\mathbf{I}}_{i+1}) \in \mathbb{R}^{H' \times W' \times C}$, a 4D correlation volume $\mathbf{C}(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1}) \in \mathbb{R}^{H' \times W' \times H' \times W'}$ at the full resolution of the feature maps is computed as follows:

$$\mathbf{C}(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})_{xyuv} = \exp\left(\sum_c \mathcal{F}_\theta(\mathbf{I}_i)_{xyc} \cdot \mathcal{F}_\theta(\hat{\mathbf{I}}_{i+1})_{uvc}\right), \quad (2)$$

where c is an index along the channel dimension of frame features. Eq. (2) pairs up all the pixel features in the neighboring feature maps, and computes their visual correlation using a radial basis function kernel. This kernel function is introduced to signify the magnitude of strong correlations while suppressing the weak ones.

For each reference frame \mathbf{I}_i , we compute two *inter-frame correlation volumes* with the two immediate neighboring frames, i.e. $\mathbf{C}(\mathbf{I}_i, \hat{\mathbf{I}}_{i-1})$ and $\mathbf{C}(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})$. To account for intra-frame motions within the exposure time, we also in-

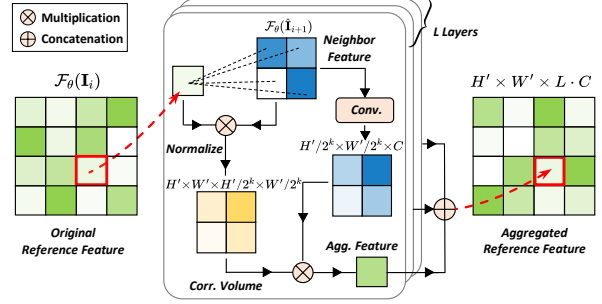


Figure 3: Illustration of the relative aggregation module, which aggregates the features for the reference frame based on the normalized correlation volume.

clude an *intra-frame correlation volume* $\mathbf{C}(\mathbf{I}_i, \mathbf{I}_i)$, to identify correspondence within the reference frame. As a result, we acquire three correlation volumes by comparing the reference frame with all the input frames on every pixel pairs across all the spatial range in the feature space.

Correlation Volume Pyramid. In order to enhance the network capacity to learn detailed visual information, we further construct a correlation volume pyramid by incorporating multi-scale features. Inspired by the work [37], we construct multi-scale representations by pooling over features, instead of input frames, to avoid the heavy computational overhead. Specifically, we augment our network with an L -layer correlation pyramid $\{\mathbf{C}^k\}_{k=1}^L$ by max-pooling over the features of neighboring frames with L kernels of growing sizes. Thereby, the dimension of the correlation volume \mathbf{C}^k is $H' \times W' \times H'/2^k \times W'/2^k$. Different from [37], we downsample only the feature maps of the neighboring frame while keeping the resolution of the reference frame features. As shown in [41], such design helps to model motions with multiple displacement scales while maintaining high-resolution image details.

3.3. Deblurring with Volumetric Correspondence

In order to restore a sharp reference frame, we develop a relative aggregation module, which outputs an aggregated feature map of the reference frame by looking up the correlated spatial positions in the neighboring frames. We then concatenate the aggregated feature maps of the reference frame and finally feed them into a reconstruction network to obtain the final deblurring result.

Correlative Aggregation Module. Given a pair of frame features $\{\mathcal{F}_\theta(\mathbf{I}_i), \mathcal{F}_\theta(\hat{\mathbf{I}}_{i+1}) \in \mathbb{R}^{H' \times W' \times C}\}$, and the correlation volume pyramid $\{\mathbf{C}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})\}$, where \mathbf{C}^k has the dimension $\mathbb{R}^{H' \times W' \times H'/2^k \times W'/2^k}$, the correlative aggregation module outputs an aggregated feature map $\rho(\mathbf{I}_{i \leftarrow i+1})$ at the pyramid scale k . The aggregated feature map $\rho(\mathbf{I}_{i \leftarrow i+1})$ has the same dimension as the reference frame feature $\mathcal{F}_\theta(\mathbf{I}_i)$.

We note that the last two dimensions of the correla-

tion volume $\mathbf{C}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})_{xy} \in \mathbb{R}^{H'/2^k \times W'/2^k}$ represents a correlation matrix, which describes the correspondence between the particular feature $\mathcal{F}_\theta(\mathbf{I}_i)_{xy}$ and the downsampled neighboring feature map in all the spatial positions. In order to aggregate the features based on the correspondence, we first normalize the correlation volume along these two dimensions by their sum. Namely, we compute the normalized correlation volume $\tilde{\mathbf{C}}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})$ as follows,

$$\tilde{\mathbf{C}}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})_{xyuv} = \frac{\mathbf{C}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})_{xyuv}}{\sum_u \sum_v \mathbf{C}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1})_{xyuv}}. \quad (3)$$

Since the correlation volume only provides information on the spatial correspondence and does not consider channel-wise information, we apply another convolution operation $\phi(\cdot)$ to the neighbor feature map. In this way, we selectively aggregate the channel information and obtain a refined feature map $\phi(\hat{\mathbf{I}}_{i+1})_k \in \mathbb{R}^{H/2^k \times W/2^k \times C}$.

Next, we aggregate the pixels on the downsampled feature map based on the correlations. Specifically, we compute the aggregated features $\rho(\mathbf{I}_{i \leftarrow i+1}) \in \mathbb{R}^{H \times W \times LC}$ as follows:

$$\begin{aligned} \rho(\mathbf{I}_{i \leftarrow i+1})_k &= \tilde{\mathbf{C}}^k(\mathbf{I}_i, \hat{\mathbf{I}}_{i+1}) \phi(\hat{\mathbf{I}}_{i+1})_k, \\ \rho(\mathbf{I}_{i \leftarrow i+1}) &= \parallel_{k=1}^L \rho(\mathbf{I}_{i \leftarrow i+1})_k, \end{aligned} \quad (4)$$

where the \parallel operator concatenates aggregated feature maps from each volume pyramid layer along the channel dimension. We also repeat the correlative aggregation once on the preceding neighboring frame and once on the reference frame itself as a self-correlative aggregation. In this way, we obtain aggregated feature maps $\rho(\mathbf{I}_{i \leftarrow i-1})$ and $\rho(\mathbf{I}_{i \leftarrow i})$, respectively.

We concatenate all the three aggregated feature maps of the reference frame along the feature dimension, each of dimension $H \times W \times L \cdot C$, thus obtaining the final feature map. Lastly, we follow an reconstruction network to restore the sharp reference frame. We implement the reconstruction network as in [40] considering its effectiveness. Detailed architecture will be provided in the supplementary materials.

3.4. Progressive Generative Adversarial Learning

Inspired by [26], we split the training procedure into stages and enable the model to progressively achieve fine-detailed restoration. To further elaborate, we use T_i^s to denote a single training step in s -th stage to restore \mathbf{I}_i . For the first stage, T_i^1 takes \mathbf{I}_{i-1} , \mathbf{I}_i and \mathbf{I}_{i+1} and outputs \mathbf{R}_i^1 as a restoration for \mathbf{I}_i . The later stages take the output of the previous stage as input. For example, T_i^2 takes \mathbf{R}_{i-1}^1 , \mathbf{R}_i^1 and \mathbf{R}_{i+1}^1 in order to restore \mathbf{I}_i . In this way, a two-staged learning requires five consecutive frames as input.

Adversarial training succeeds in generating and restoring visually realistic images [17, 10, 16]. On video deblurring, Zhang [51] uses a discriminator to distinguish between

the restored single frame and its ground-truth, which does not consider temporal constraints. In contrast, we exploit the proposed progressive training scheme and design an adversarial loss to encourage models to output temporally-consistent results. In this regard, we propose a temporal discriminator \mathcal{D}_ζ , parameterized by ζ , that distinguishes $\vec{\mathbf{R}}_i = \{\mathbf{R}_{i-1}^1, \mathbf{R}_i^2, \mathbf{R}_{i+1}^2\}$ from their ground-truth sharp images $\vec{\mathbf{S}}_i = \{\mathbf{S}_{i-1}, \mathbf{S}_i, \mathbf{S}_{i+1}\}$. In specific, \mathcal{D}_ζ consists of four 3D convolution layers followed by a global pooling to extract temporal features for the ground-truth and restored frame sequences. It then predicts whether the sequence is a ground-truth sequence or not using a binary classifier. The adversarial loss is defined as follows,

$$L_{adv} = \mathbb{E}_\xi[\log(D_\zeta(\vec{\mathbf{S}}))] + \mathbb{E}_\chi[\log(1 - D_\zeta(\vec{\mathbf{R}}))], \quad (6)$$

where ξ is a distribution of the ground-truth sharp sequences and χ is that of the input sequences. The discriminator \mathcal{D}_ζ and the deblurring model are optimized adversarially in the typical minimax fashion. In this way, we encourage the model to produce a consistent sequence that visually resemble the ground-truth. We combine the proposed adversarial loss with L_1 loss for reconstruction, *i.e.* $L_{total} = L_1 + \alpha L_{adv}$, where $\alpha = 0.1$.

4. Experiments

In this section, we first present experimental results by evaluating our proposed model ARVo on standard datasets. Then we conduct ablation studies to investigate the effects of each component in the proposed model.

4.1. Implementation Details

We implement ARVo in PyTorch [32]. We adopt a three-layer correlation pyramid to achieve a trade-off between video deblurring quality and computational efficiency. We use PWC-Net [37] for optical flow estimation considering its robustness in real-life domains. During the training, we feed five consecutive frame patches of 256×256 resolution to the model and optimize it using Adam optimizer [12]. We set the initial learning rate as 10^{-4} and reduce it by half every 200 epochs. We also use random flipping and rotations to augment the training dataset. The weights are shared across all the progressive stages for a compact model.

4.2. Datasets and Evaluation Metrics

We evaluate the proposed methods on two datasets DVD [36] and a newly introduced HFR-DVD dataset. We adopt peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) [44] for evaluating the deblurring results.

DVD [36]. This dataset has 71 videos with 6,708 blurry-sharp pairs, where 61 videos are used for training (5,708 pairs) and 10 videos (1,000 pairs) for testing. These videos are captured at 240 fps with mobile phones and DSLR. The




Figure 4: Example frames from our newly introduced HFR-DVD dataset, showing a diversity in scenes and motions. The left half of the frame is sharp while the right half shows the synthetic blurs.

DVD dataset is widely adopted in the community and has promoted the development of video deblurring models.

HFR-DVD. To simulate realistic blurs, one typical approach is capturing videos at a high frame rate and then averaging multiple consecutive sharp frames for blurs. In this way, the quality of the synthetic blurry images is affected by the exposure time of devices. In particular, a lower frame rate leads to longer idling time between adjacent exposures, yielding unnatural spikes in the blur trajectories. In this regard, previous datasets mainly tackle the issue by introducing virtual frames. For example, Su *et al.* [36] generate inter-frame images based on optical flow estimation. Nah *et al.* [25] proposed to increase the frame rate of sharp videos by interpolations with a pretrained neural network. In contrast, we construct a High-Frame-Rate Dataset for Video Deblurring (HFR-DVD) by capturing videos at 1,000 fps with a SONY DSC-RX10 IV camera, thus avoiding the artifacts caused by automatic frame interpolation.

When capturing videos, we pay attentions to the video quality and the diversity of scenes. A few example frames are shown in Fig. 4. We synthesize the blurs by averaging 41 neighboring sharp frames. To reduce the noises, we downsize each frame to 960×540 resolution. We then subsample the high-frame-rate video temporally by keeping the middle frame every N frames. We choose N randomly from the interval [38, 44]. In this way, we finally obtain blurry videos at 25 fps. In total, we acquire 120 videos for training and 30 videos for testing, each with 90 frames, amounting to 13,500 blurry-sharp pairs altogether. We will release the HFR-DVD dataset publicly soon.

4.3. Comparisons with the State-of-The-Art

Synthetic Blurry Videos. We first compare our model ARVo with previous deblurring methods [7, 36, 40, 51, 43, 56, 45, 26] on the DVD dataset. Among them, the work [7] is based on variational models whereas the rest works rely on deep neural networks.

Table 1 shows the quantitative comparisons on the DVD dataset. As shown, our model performs favorably against previous works in both metrics. In Fig. 5 we show the deblurred images from different models. The work [36] designs an end-to-end convolution model for deblurring video frames. However, they rely mainly on explicit alignment methods such as homography or optical flows only, therefore are less favorable to restore abrupt motion blurs. The method [56] uses only two consecutive frames as the contextual information, thus do not fully utilize the temporal information in the video. As can be seen in Figure 5, it fails to fully restore the structure of scene objects. We further note that the method [26] computes a temporal sharpness prior to guide the model to focus on areas with severe blurs. However, the prior only considers a small range of pixels, thereby, it is less effective when fast motions are present. In contrast, our model takes in multiple consecutive frames as input, providing stronger video contextual clues. By implicitly constructing the pixel-pair correspondence across all the spatial range, our model is able to effectively captures pixel correlations with large displacements. This is particularly effective for blurs due to fast motions. As a result, our proposed model restores better the detailed scene content.

We further evaluate our method quantitatively on the newly introduced HFR-DVD dataset. We remark that the HFR-DVD dataset exhibits more fast motions compared to DVD dataset, which further challenges the capacity of video deblurring models. We select five models [36, 51, 56, 26] among the best performers based on the quantitative evaluations on DVD as the baseline methods.

As Table 2 shows, our model performs favorably against baseline methods in both metrics. Also, we notice that it outperforms the previous best model by a slightly larger margin than on the DVD dataset. This validates our motivation for tackling fast motion blurs with all-range spatial correspondence. In Fig. 1, we show qualitatively the deblurring results on the HFR-DVD dataset. Due to the inter-frame correlation volume computed on all the pixel pairs, our model utilizes better the video context thus producing sharper restored frames with finer visual details.

Real Blurry Videos. In order to validate the generalization ability of our proposed model, we validate ARVo on the real-world blurry videos from [36]. As can be seen from Fig. 6, our model restores better structural information than the previous methods, which demonstrates a promising video deblurring outcome in real-life scenarios.

4.4. Model Analysis and Discussions

In this section, we take DVD dataset as an example to analyze the effects of multiple components in our proposed method. We also investigate how the proposed correlative aggregation module affects the feature learning to better understand the behavior of our models.

Table 1: Quantitative comparisons on the DVD dataset [36]. Following [26], we report PSNR and SSIM metrics considering all the restored video frames. As can be seen, our model performs favorably against previous approaches in both metrics.

Methods	Kim <i>et al.</i> [7]	EDVR [43]	Tao <i>et al.</i> [40]	Su <i>et al.</i> [36]	DBLRNet [51]	STFAN [56]	Xiang <i>et al.</i> [45]	TSP [26]	ARVo (Ours)
PSNR	26.94	28.51	29.98	30.01	30.08	31.15	31.68	32.13	32.80
SSIM	0.8158	0.8637	0.8842	0.8877	0.8845	0.9049	0.9157	0.9268	0.9352



Figure 5: Qualitative comparisons on the DVD [36] dataset. Our model benefits from constructing pixel pair correspondence in the feature space, thus being able to better utilize neighboring sharp patches to improve the restoration quality.




Figure 6: Qualitative comparisons on the real blurry frames. As shown, our model provides sharper restored frames while preserving more fine-grained visual details, which demonstrates the effectiveness of our model in real-life scenarios.

Table 2: Quantitative comparisons using the HFR-DVD dataset. We adhere to the same evaluation protocol as in Table 1 for a fair and consistent comparison.

Methods	[36]	[51]	[56]	[26]	Ours
PSNR	27.73	27.79	28.48	29.71	31.15
SSIM	0.8512	0.8505	0.8560	0.8822	0.9063

Effect of Correlation Pyramid We propose to construct a correlation volume pyramid in the feature space in order to handle large pixel displacements. To analyze the effect of correlation volumes, we experiment with different choices of pyramid layers with $L = 0, 1, 2, 3, 4$. When $L = 0$, we do not construct volume pyramid. Instead, we directly feed the frame features from the encoder to the reconstruction network. As shown in Table 3, learning features using correlation volumes evidently improves the deblurring results even with only one layer of pyramid. By adding pyramid layers at multiple spatial scales, we observe further im-

provement in the restoration quality. This is because the increased receptive fields help to capture better visual context and long-range motion dependencies. We notice no significant benefit by building more than three layers of pyramid. This shows that a pyramid layer with a downsampling rate of 16 is too coarse to provide useful information to enhance the deblurring model.

Effect of Alignment Methods While typical optical flow-based alignment methods provide useful spatial alignment, it is less effective to restore sharp images with fast motions. We complement the typical explicit alignment method by constructing correlation volumes in the feature space. In order to understand the contribution of different alignment methods, we conduct experiments to use (i) no alignment (no align); (ii) only optical flow-based alignment (O.F.); (iii) only correlation volumes (Vol.). Results are shown in Table 4. As can be seen, using no alignment method leads to less favorable results. This is because alignment methods, either explicit or implicit, help the model to identify pixel correspondence more easily. We also remark that the opti-

Table 3: Effects of using different layers of correlation pyramids on the video deblurring result. We highlight the setting in our final model in bold.

Num. Layers	0	1	2	3	4
PSNR	31.85	32.45	32.69	32.80	32.81
SSIM	0.9132	0.9310	0.9335	0.9352	0.9352

Table 4: Effects of using different alignment methods. The O.F. column shows results with only optical flow-based alignment; the Vol. column shows results when using only correlation volumes; the O.F. + Vol. shows results when using both methods together. The final setting is in bold.

Align Methods	no align	O.F.	Vol.	O.F. + Vol.
PSNR	31.19	31.85	32.01	32.80
SSIM	0.9055	0.9132	0.9192	0.9352

Table 5: Effects of progressive training using different number of stages and input frames. The final setting is in bold.

Num. Stages (Frames)	1 (3)	2 (5)	3 (7)
PSNR	32.28	32.80	32.83
SSIM	0.9254	0.9352	0.9349

cal flow-based methods and correlation volumes are complementary in their alignment effects. Though using correlation volumes solely provides slightly better results than explicit alignment, our model provides the best restoration quality when combining both alignment methods.

Effect of Progressive Learning In Table 5, we show the influence of progressively optimizing the deblurring model. As shown, our model gains an improvement using two stages while benefiting less from considering further frames. This is because the scenes change largely as the input video segment gets longer, thus making it harder to find related sharp patches. We decide on two stages to achieve a trade-off between deblurring efficiency and quality.

Effect of Generative Adversarial Training We observe generative adversarial training does not significantly impact the metrics, obtaining 32.79 (-0.01) for PSNR and 0.9354 (+0.02) for SSIM. This might attribute to that the adversarial loss does not optimize towards the pixel-values in regard to the ground-truth. However, we do observe a perceivable difference on the restored images. In Fig. 7, we demonstrate the effect of adversarial loss. As can be seen, the generative adversarial training helps to maintain the visual consistency across frames.

Feature Visualization for Correlative Aggregation In order to gain further insights into the proposed model, we compare the feature of the reference frame before and after the correlative aggregation and show the changes on the fea-



Figure 7: The middle image is from stage 2 while the rest are from stage 1. **Top:** without adversarial loss. **Bottom:** with adversarial loss. Adversarial training helps to maintain the visual quality in consecutive output frames.




Figure 8: **Left:** reference frame. **Right:** Visualization of the feature update after applying the correlative aggregation. Lighter positions indicate more significant changes. The green box shows a blurry patch that is updated evidently. The red box shows that the network updates less on a shaded road patch with less blurs.

ture map in Fig. 8. We observe that the network effectively learns to refine areas with severe blurs, while applying less updates on patches with less blurs.

5. Conclusion

In this work, we propose a novel implicit method to construct spatial correspondence for video deblurring. The method builds a correlation volume pyramid by matching pixel-pairs between the reference frame and the neighboring frame in all the spatial range. Based on such a correlation volume pyramid, we develop a correlative aggregation module to enhance the features of the reference frame. Sharp frames are then restored from the enhanced feature maps via a reconstruction network. We design a generative adversarial training scheme to optimize the model progressively. We also present a newly-collected high-frame-rate dataset for video deblurring (HFR-DVD), featuring sharper frames and more realistic blurs. Our model achieves favorable performance on both datasets quantitatively and qualitatively, establishing a new state-of-the-art for video deblurring.

Acknowledgement This research is funded in part by the ARC Centre of Excellence for Robotics Vision (CE140100016) and ARC-Discovery (DP 190102261).

References

- [1] Jean-Charles Bazin, Hongdong Li, In So Kweon, Cédric Demonceaux, Pascal Vasseur, and Katsushi Ikeuchi. A branch-and-bound approach to correspondence and grouping problems. *IEEE transactions on pattern analysis and machine intelligence*, 35(7):1565–1576, 2012.
- [2] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Tom Drummond, Hongdong Li, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *arXiv preprint arXiv:2010.13501*, 2020.
- [3] Sunghyun Cho, Jue Wang, and Seungyong Lee. Video deblurring for hand-held cameras using patch-based synthesis. *ACM Transactions on Graphics (TOG)*, 31(4):1–9, 2012.
- [4] Hongyun Gao, Xin Tao, Xiaoyong Shen, and Jiaya Jia. Dynamic scene deblurring with parameter selective sharing and nested skip connections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3848–3856, 2019.
- [5] Dong Gong, Jie Yang, Lingqiao Liu, Yanning Zhang, Ian Reid, Chunhua Shen, Anton Van Den Hengel, and Qinfeng Shi. From motion blur to motion flow: a deep learning solution for removing heterogeneous motion blur. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2319–2328, 2017.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015.
- [7] Tae Hyun Kim and Kyoung Mu Lee. Generalized video deblurring for dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5426–5434, 2015.
- [8] Tae Hyun Kim, Kyoung Mu Lee, Bernhard Scholkopf, and Michael Hirsch. Online video deblurring via dynamic temporal blending network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4038–4047, 2017.
- [9] Younghyun Jo, Seung Wug Oh, Jaeyeon Kang, and Seon Joo Kim. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3224–3232, 2018.
- [10] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [11] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 66–75, 2017.
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] Piotr Koniusz, Lei Wang, and Anoop Cherian. Tensor representations for action recognition. 2020.
- [14] Piotr Koniusz and Hongguang Zhang. Power normalizations in fine-grained image, few-shot image and graph classification. 2020.
- [15] Dilip Krishnan and Rob Fergus. Fast image deconvolution using hyper-laplacian priors. In *Advances in neural information processing systems*, pages 1033–1041, 2009.
- [16] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8183–8192, 2018.
- [17] 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.
- [18] Dongxu Li, Cristian Rodriguez, Xin Yu, and Hongdong Li. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1459–1469, 2020.
- [19] Dongxu Li, Chenchen Xu, Xin Yu, Kaihao Zhang, Ben Swift, Hanna Suominen, and Hongdong Li. Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. *arXiv preprint arXiv:2010.05468*, 2020.
- [20] Dongxu Li, Xin Yu, Chenchen Xu, Lars Petersson, and Hongdong Li. Transferring cross-domain knowledge for video sign language recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6205–6214, 2020.
- [21] Guangrui Li, Guoliang Kang, Wu Liu, Yunchao Wei, and Yi Yang. Content-consistent matching for domain adaptive semantic segmentation. In *European Conference on Computer Vision*, pages 440–456. Springer, 2020.
- [22] Yunpeng Li, Sing Bing Kang, Neel Joshi, Steve M Seitz, and Daniel P Huttenlocher. Generating sharp panoramas from motion-blurred videos. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2424–2431. IEEE, 2010.
- [23] Yasuyuki Matsushita, Eyal Ofek, Weina Ge, Xiaoou Tang, and Heung-Yeung Shum. Full-frame video stabilization with motion inpainting. *IEEE Transactions on pattern analysis and Machine Intelligence*, 28(7):1150–1163, 2006.
- [24] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016.
- [25] Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.

- [26] Jinshan Pan, Haoran Bai, and Jinhui Tang. Cascaded deep video deblurring using temporal sharpness prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3043–3051, 2020.
- [27] Jinshan Pan, Deqing Sun, Hanspeter Pfister, and Ming-Hsuan Yang. Blind image deblurring using dark channel prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1628–1636, 2016.
- [28] Liyuan Pan, Yuchao Dai, and Miaomiao Liu. Single image deblurring and camera motion estimation with depth map. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2116–2125. IEEE, 2019.
- [29] Liyuan Pan, Yuchao Dai, Miaomiao Liu, and Fatih Porikli. Simultaneous stereo video deblurring and scene flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4382–4391, 2017.
- [30] Liyuan Pan, Yuchao Dai, Miaomiao Liu, Fatih Porikli, and Quan Pan. Joint stereo video deblurring, scene flow estimation and moving object segmentation. *IEEE Transactions on Image Processing*, 29:1748–1761, 2019.
- [31] Liyuan Pan, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Phase-only image based kernel estimation for single image blind deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6034–6043, 2019.
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019.
- [33] Travis Portz, Li Zhang, and Hongrui Jiang. Optical flow in the presence of spatially-varying motion blur. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1752–1759. IEEE, 2012.
- [34] Johannes L Schonberger, Sudeipta N Sinha, and Marc Pollefeys. Learning to fuse proposals from multiple scanline optimizations in semi-global matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 739–755, 2018.
- [35] Qi Shan, Jiaya Jia, and Aseem Agarwala. High-quality motion deblurring from a single image. *Acm transactions on graphics (tog)*, 27(3):1–10, 2008.
- [36] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1279–1288, 2017.
- [37] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018.
- [38] Jian Sun, Wenfei Cao, Zongben Xu, and Jean Ponce. Learning a convolutional neural network for non-uniform motion blur removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 769–777, 2015.
- [39] Huixuan Tang, Scott Cohen, Brian Price, Stephen Schiller, and Kiriakos N Kutulakos. Depth from defocus in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2740–2748, 2017.
- [40] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Jiaya Jia. Scale-recurrent network for deep image deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8174–8182, 2018.
- [41] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. *arXiv preprint arXiv:2003.12039*, 2020.
- [42] Jianyuan Wang, Yiran Zhong, Yuchao Dai, Kaihao Zhang, Pan Ji, and Hongdong Li. Displacement-invariant matching cost learning for accurate optical flow estimation. *arXiv preprint arXiv:2010.14851*, 2020.
- [43] 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.
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [45] Xinguang Xiang, Hao Wei, and Jinshan Pan. Deep video deblurring using sharpness features from exemplars. *IEEE Transactions on Image Processing*, 29:8976–8987, 2020.
- [46] Xin Yu, Feng Xu, Shunli Zhang, and Li Zhang. Efficient patch-wise non-uniform deblurring for a single image. *IEEE Transactions on Multimedia*, 16(6):1510–1524, 2014.
- [47] Haichao Zhang and Lawrence Carin. Multi-shot imaging: joint alignment, deblurring and resolution-enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2925–2932, 2014.
- [48] Haichao Zhang, David Wipf, and Yanning Zhang. Multi-image blind deblurring using a coupled adaptive sparse prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1051–1058, 2013.
- [49] Jiawei Zhang, Jinshan Pan, Jimmy Ren, Yibing Song, Linchao Bao, Rynson WH Lau, and Ming-Hsuan Yang. Dynamic scene deblurring using spatially variant recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2521–2529, 2018.
- [50] Kaihao Zhang, Wenhan Luo, Björn Stenger, Wenqi Ren, Lin Ma, and Hongdong Li. Every moment matters: Detail-aware networks to bring a blurry image alive. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 384–392, 2020.
- [51] Kaihao Zhang, Wenhan Luo, Yiran Zhong, Lin Ma, Wei Liu, and Hongdong Li. Adversarial spatio-temporal learning for video deblurring. *IEEE Transactions on Image Processing*, 28(1):291–301, 2018.
- [52] Kaihao Zhang, Wenhan Luo, Yiran Zhong, Lin Ma, Bjorn Stenger, Wei Liu, and Hongdong Li. Deblurring by realistic blurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2737–2746, 2020.

- [53] Shan Zhang, Dawei Luo, Lei Wang, and Piotr Koniusz. Few-shot object detection by second-order pooling. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020.
- [54] Yiran Zhong, Yuchao Dai, and Hongdong Li. Stereo computation for a single mixture image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 435–450, 2018.
- [55] Yiran Zhong, Pan Ji, Jianyuan Wang, Yuchao Dai, and Hongdong Li. Unsupervised deep epipolar flow for stationary or dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12095–12104, 2019.
- [56] Shangchen Zhou, Jiawei Zhang, Jinshan Pan, Haozhe Xie, Wangmeng Zuo, and Jimmy Ren. Spatio-temporal filter adaptive network for video deblurring. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2482–2491, 2019.
- [57] Hao Zhu and Piotr Koniusz. Simple spectral graph convolution. In *International Conference on Learning Representations*, 2021.