

Unidirectional Video Denoising by Mimicking Backward Recurrent Modules with Look-ahead Forward Ones

Junyi Li¹, Xiaohe Wu¹, Zhenxin Niu², and Wangmeng Zuo¹

¹ Harbin Institute of Technology

² Xidian University

{nagejacob, csxhwu, zhenxingniu}@gmail.com

wmzuo@hit.edu.cn




Fig. 1: Illustration of representative deep network architectures for video denoising. When handling the current frame, (a) CNN cannot exploit long-term temporal information, and (b) ForwardRNN cannot exploit any future frames. (c) BiRNN is effective in exploiting the information from all frames, but can only be performed in an *offline* manner. In comparison, (d) FloRNN can leverage both the historical information and the crucial near-future frames, thereby being very appealing for *unidirectional* video denoising.

Abstract. While significant progress has been made in deep video denoising, it remains very challenging for exploiting historical and future frames. Bidirectional recurrent networks (BiRNN) have exhibited appealing performance in several video restoration tasks. However, BiRNN is intrinsically offline because it uses backward recurrent modules to propagate from the last to current frames, which causes high latency and large memory consumption. To address the offline issue of BiRNN, we present a novel recurrent network consisting of forward and look-ahead recurrent modules for unidirectional video denoising. Particularly, look-ahead module is an elaborate forward module for leveraging information from near-future frames. When denoising the current frame, the hidden features by forward and look-ahead recurrent modules are combined, thereby making it feasible to exploit both historical and near-future frames. Due to the scene motion between non-neighboring frames, border pixels missing may occur when warping look-ahead feature from near-future frame to current frame, which can be largely alleviated by incorporating forward warping and proposed border enlargement. Experiments show that our method achieves state-of-the-art performance

with constant latency and memory consumption. Code is available at <https://github.com/nagejacob/FloRNN>.

Keywords: video denoising, recurrent neural networks, temporal alignment.

1 Introduction

Recent years have witnessed the great success of deep networks in video denoising [8, 11, 12, 31, 42–44, 48, 51]. Compare to image denoising, temporal information plays a pivotal role in video denoising, which is generally restricted by the spatial misalignment among consecutive frames. Self-similar spatial-temporal patch aggregation has been suggested for spatio-temporal modeling [12, 44], but usually results in heavy computational cost. Other spatio-temporal alignment methods, *e.g.*, optical flow [41], deformable convolution [10], and kernel prediction network (KPN) [32], have also been studied and applied for video denoising [32, 42, 46, 48, 51]. Instead of explicit motion estimation and compensation, cascaded U-Net [43] and multi-stage recurrent network [31] are further suggested for the efficiency issue. Albeit the progress in deep video denoising, it remains a challenging issue to exploit historical and future frames for video denoising.

Bidirectional recurrent networks (BiRNN) provide a convenient way for temporal modeling and have been very appealing in several video restoration tasks [5, 6, 23, 39, 40, 50]. In video super-resolution (VSR), BasicVSR [5] and its extension [6] have outperformed most state-of-the-art methods in terms of VSR performance and efficiency. Benefited from the forward and backward recurrent modules, BiRNN is effective in exploiting the information from all frames to restore the current frame. In comparison, convolutional network (CNN) takes neighborhood frames as inputs and only exploits short-term temporal information. However, BiRNN is intrinsically an offline approach where the backward recurrent module is deployed to propagate information from the last frame to current frame, and the restoration result of current frame cannot be obtained unless all video frames are processed. Meanwhile, hidden features of all frames have to be maintained in the memory during inference, which causes high memory consumption. The high latency and large memory consumption limit the practicability of BiRNN.

To address the offline issue of BiRNN, we present a novel recurrent network consisting of **F**orward and **l**ook-ahead recurrent modules (*i.e.*, FloRNN) for unidirectional video denoising. From [5], the future frames are important to boost the denoising performance of current frame. Simply discarding backward modules from BiRNN (denote as ForwardRNN) hampers the use of future frames and results in inferior performance. Moreover, among all future frames, the near-future frames are most important to enhance the denoising results of current frame. To leverage future information while addressing offline issue, we present a look-ahead recurrent module for exploiting near-future frames. As shown in Fig. 1(d), our look-ahead module F_l adopts a forward recurrent architecture, but

propagates k frames ahead of the forward module F_f . The look-ahead feature of the near-future frame (*i.e.*, $t + k$) is warped back to align with current frame (*i.e.*, t), and incorporated with the forward feature from F_f to produce current denoising results. Our FloRNN adopts a unidirectional propagation, yet extracts crucial near-future frames with F_l . Thus, it has the potential to approach the performance of BiRNN while maintaining constant latency and memory consumption. Furthermore, one can adjust k to meet the latency constraint, smaller k for real-time tasks, and larger k for better results.

Nonetheless, there remains a major issue to address for applying look-ahead module. As shown in Fig. 3(a), with an assumption of $k = 2$, the motion between the t -th and $(t+2)$ -th frames makes that an amount of border pixels are lost when warping the look-ahead feature from $(t + 2)$ -th frame back to t -th frame. Obviously, such insufficient utilization of near-future frames is harmful to denoising performance. To address this issue, we present to incorporate forward warping and border enlargement mechanism. As shown in Fig. 3(b), forward warping is adopted in the look-ahead module instead of backward warping. Meanwhile, we suggest to enlarge the border of look-ahead feature to a fixed ratio (*e.g.*, 10% in this work) during the forward warping. Due to that forward warping allows warped feature to splat out of border (*i.e.*, pixels out of the red box), we can save the missed border pixels when warping look-ahead feature to future frames, and recover them when aligning look-ahead feature back to the current frame. Our forward warping and border enlargement mechanism can largely reduce the amount of missed border pixels for better exploiting near-future frames and also benefiting denoising performance.

Extensive experiments are conducted to evaluate our FloRNN on commonly used video denoising benchmarks. Both synthetic additive white Gaussian noise (AWGN) and real-world noise are considered in our experiments. For AWGN, our method outperforms the second best competing method PaCNet [44] by a large margin (*i.g.*, 0.76dB by PSNR) on DAVIS. On real-world video datasets such as CRVD [51] and IOCV [27], FloRNN also achieves the best quantitative results than the competing methods. We also show that our look-ahead module could be applied to state-of-the-art BiRNN methods (*e.g.*, BasicVSR++ [6]), which forms a unidirectional counterpart and benefits from the advances in BiRNNs in the supplementary material.

To sum up, the main contribution of this work includes:

- A novel recurrent network (*i.e.*, FloRNN) is presented for unidirectional video denoising by incorporating forward and look-ahead recurrent modules.
- Forward warping and border enlargement are equipped in look-ahead recurrent module for better exploiting the near-future frames and also benefiting denoising performance.
- Experiments show that our method performs favorably against state-of-the-arts on various video denoising datasets.

2 Related Work

2.1 Image Denoising.

With the observation that an image patch usually has many similar counterparts within the same image, many traditional methods are developed for joint modeling of a stack of similar patches to remove noise, such as non-local means (NLM) [3], BM3D [9], and WNNM [16]. Other methods use handcrafted priors on image patches, *e.g.*, sparsity [14] and Gaussian Mixture [56]. Recently, CNN based methods have achieved favorable performance. DnCNN [52] incorporates residual learning [18] and batch normalization [20] for better convergence. FFDNet [53] takes noise level map as input, and trains a single model for handling various noise levels. With the collected real-world denoising datasets [1, 35], several methods [17, 24, 55] have also been suggested to handle real noise.

2.2 Video Denoising.

For exploiting temporal information, traditional patch-based methods [2, 4, 30] search similar patches on volumetric data, and are very time consuming. By leveraging deep image denoising, several methods adopt two-stage scheme including spatial denoising and temporal fusion. ViDeNN [8] uses a plain CNN as fusion network, while DVDNet [42] assists the fusion with optical flow and warping neighbor frames to current frame. Temporal aggregation mechanisms have also been investigated and integrated into network design, *e.g.*, non-local search [12, 44], cascaded U-Net [43], kernel prediction network (KPN) [32, 46, 48], deformable convolution [51], and channel shifting [37]. However, these aggregation mechanisms either are complicated and inefficient, or cannot achieve state-of-the-art performance. EMVD [31] proposes a multi-stage recurrent architecture for mobile devices, but performs inferior when scaling to larger model on GPUs.

2.3 RNN in Video Restoration.

RNN provides a convenient way for temporal modeling and thus can be readily applied to video restoration [7, 19]. Taking VSR as an example, FRVSR [38] warps the $(t-1)$ -th output to t -th frame, and use it as additional input to super-resolve the t -th frame. RLSP [15] introduces high-dimensional latent states for efficient propagation. RSDN [21] divides the input into structure and detail components to effectively exploit temporal correlation. BasicVSR [5] adopts bidirectional propagation and optical flow based alignment, achieving state-of-the-art VSR performance. BasicVSR++ [6] further improves BasicVSR with flow-guided deformable alignment and second order propagation. However, bidirectional propagation in BasicVSR [5] and its extension [6] makes them unable to be used for online video restoration. RNN has also been investigated for video deblurring [33, 40, 54], video inpainting [22, 23], video frame interpolation [39, 47], and video deraining [29, 49, 50].

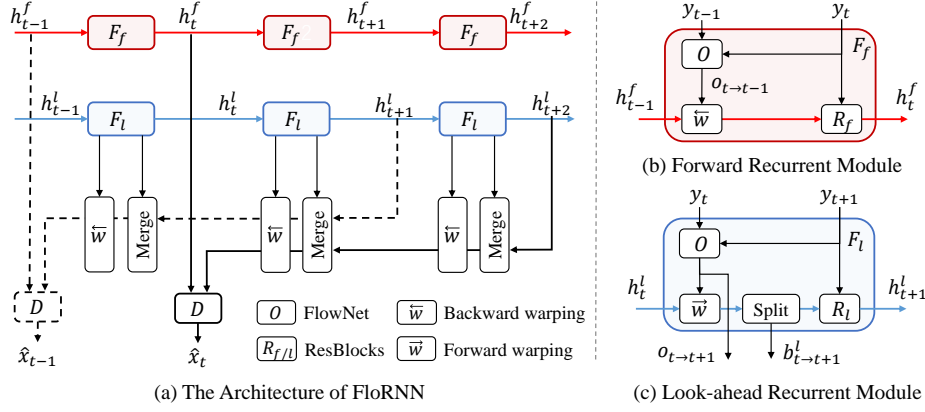


Fig. 2: **Illustration of FloRNN.** (a) FloRNN is mainly built with three components, *i.e.*, forward recurrent module F_f , look-ahead recurrent module F_l and decoder D . (b) F_f propagates from first to current frames which exploits all previous frames, (c) F_l propagates similarly to F_f in a forward manner, but k (*e.g.*, $k = 2$) frames ahead of F_f . The look-ahead feature is aligned back to current frame (details of Split and Merge operations are provided in Fig. 3 (b)) for exploiting crucial near-future frames. With the propagated historical and near-future information, FloRNN achieves compelling results while has constant latency and memory occupation.

3 Method

Given a video sequence consisting of T noisy frames $\{\mathbf{y}_t\}_{t=1}^T$, video denoising aims to produce the prediction $\{\hat{\mathbf{x}}_t\}_{t=1}^T$ for approximating its clean video $\{\mathbf{x}_t\}_{t=1}^T$. For better denoising of current frame \mathbf{y}_t , one favorable solution is to exploit all frames $\{\mathbf{y}_t\}_{t=1}^T$ to predict $\hat{\mathbf{x}}_t$. BasicVSR [5] originally suggested for VSR provides a bidirectional recurrent network (BiRNN, as shown in Fig. 1(c)) to leverage the information from all frames, and we empirically find that it is also very appealing for video denoising. However, due to the use of backward recurrent module, all the succeeding frames are required to produce backward hidden feature, making that BiRNN can only be performed in an offline manner.

To address the *offline* issue of BiRNN, we propose a novel recurrent network for *unidirectional* video denoising. As shown in Fig. 1(d), it consists of a **F**orward recurrent module as well as a **l**ook-ahead recurrent module, named as FloRNN. Analogous to BiRNN, FloRNN adopts the same forward recurrent module F_f for history frames propagation and decoder D for producing denoising results. The difference is that FloRNN substitutes the backward recurrent module of BiRNN with an elaborate look-ahead recurrent module F_l for exploiting near-feature frames. The look-ahead feature from F_l is then aligned and incorporated with forward feature from F_f to enhance the denoising results of current frame. To address the border missing issue during the alignment, we further

introduce a border enlargement mechanism with forward warping. The details of our framework are illustrated in Fig. 2. In the following, we will describe the forward module, look-ahead module, forward warping and border enlargement, respectively.

3.1 Forward Recurrent Module

Analogous to BiRNN, we first introduce the forward recurrent module F_f to propagate information from the first to the current frames. As shown in Fig. 2(b), it adopts a recurrent manner by propagating the forward hidden feature \mathbf{h}_{t-1}^f and combining it with the current frame \mathbf{y}_t to obtain the hidden feature \mathbf{h}_t^f at frame t .

$$\mathbf{h}_t^f = F_f(\mathbf{y}_t, \mathbf{y}_{t-1}, \mathbf{h}_{t-1}^f). \quad (1)$$

Nonetheless, \mathbf{h}_{t-1}^f is aligned with \mathbf{y}_{t-1} instead of \mathbf{y}_t . Following [5], we estimate the optical flow [41] from frame t to $t-1$, which is used to align and aggregate \mathbf{h}_{t-1}^f with \mathbf{y}_t . Thus, the forward recurrent module in Eqn. (1) can be further written as,

$$\begin{aligned} \mathbf{o}_{t \rightarrow t-1} &= O(\mathbf{y}_t, \mathbf{y}_{t-1}), \\ \mathbf{h}_{t-1 \rightarrow t}^f &= \overleftarrow{w}(\mathbf{h}_{t-1}^f, \mathbf{o}_{t \rightarrow t-1}), \\ \mathbf{h}_t^f &= R_f(\mathbf{y}_t, \mathbf{h}_{t-1 \rightarrow t}^f). \end{aligned} \quad (2)$$

where $O(\cdot, \cdot)$ denotes an optical flow network, and $\mathbf{o}_{t \rightarrow t-1}$ denotes the estimated optical flow from t -th frame to $(t-1)$ -th frame. $\overleftarrow{w}(\mathbf{h}_{t-1}^f, \mathbf{o}_{t \rightarrow t-1})$ stands for aligning \mathbf{h}_{t-1}^f with \mathbf{y}_t using backward warping [5] to obtain the warped hidden feature $\mathbf{h}_{t-1 \rightarrow t}^f$. R_f aggregates the warped hidden feature and current frame with multiple residual blocks to obtain the hidden feature \mathbf{h}_t^f of t -th frame.

3.2 Look-ahead Recurrent Module

To leverage the future information while addressing the offline issue, we propose the look-ahead recurrent module for exploiting near-future frames. As shown in Fig. 1(d), our look-ahead module adopts a specifically designed forward recurrent mechanism,

$$\mathcal{H}_{t+1}^l = F_l(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathcal{H}_t^l). \quad (3)$$

Here we use \mathcal{H}_{t+1}^l to indicate that the output of look-ahead module may contain hidden feature \mathbf{h}_{t+1}^l and other variables. As the look-ahead module propagates k frames ahead of the forward module, \mathbf{h}_{t+k}^l captures temporal information of k near-future frames. For restoring t -th frame, we align the look-ahead feature \mathbf{h}_{t+k}^l back to t -th frame to produce the warped look-ahead feature $\mathbf{h}_{t+k \rightarrow t}^l$, which can be generally written as,

$$\mathbf{h}_{t+k \rightarrow t}^l = \text{Align}(\mathbf{h}_{t+k}^l, \dots). \quad (4)$$




Fig. 3: Illustration of forward warping and border enlargement for our look-ahead recurrent module. (a) Aligning h_{t+2}^l to t -th frame with backward warping \overleftarrow{w} suffers from border pixels missing problem. (b) Forward warping \overrightarrow{w} allows the out of border pixels to splat to the enlarged border. We save the enlarged border when warping to future frames and recover it when warping back, which largely mitigates the border information missing problem.

With the forward feature \mathbf{h}_t^f and the aligned look-ahead feature $\mathbf{h}_{t+k \rightarrow t}^l$, the denoising result at frame t can then be obtained by,

$$\hat{\mathbf{x}}_t = D(\mathbf{h}_t^f, \mathbf{h}_{t+k \rightarrow t}^l). \quad (5)$$

where the decoder D contains two convolution layers.

Note that F_l and $\text{Align}(\cdot, \dots)$ can be implemented with different forms. For example, one can adopt the forward recurrent module defined in Eqn. (2) to implement F_l . As for $\text{Align}(\cdot, \dots)$, one direct solution is to compute the optical flow from \mathbf{y}_{t+k} to \mathbf{y}_t , then warp \mathbf{h}_{t+k}^l back to \mathbf{y}_t . In the following subsection, forward warping and border enlargement are presented as a reasonable implementation.

3.3 Forward Warping and Border Enlargement

As stated above, one straightforward way to implement the look-ahead module is adopting the forward recurrent module defined in Eqn. (2) as F_l and aligning look-ahead feature to current frame by backward warping. However, as shown in Fig. 3(a), such a straightforward implementation suffers from border information missing issue when aligning look-ahead feature back to current frame (*i.e.*, black part exists in $\mathbf{h}_{t+2 \rightarrow t}^l$). This is caused by scene motion between consecutive frames and backward warping mechanism in look-ahead module. To demonstrate this, we first give a brief introduction of backward warping. As shown in Fig. 4(a), taking warping \mathbf{h}_t^l to $\mathbf{h}_{t \rightarrow t+1}^l$ as example, the value of $\mathbf{h}_{t \rightarrow t+1}^l$




Fig. 4: Border pixels (colored in orange) in \mathbf{h}_t^l are out of border in the $(t+1)$ -th frame. (a) Backward warping ignores these border pixels in the warping result $\mathbf{h}_{t \rightarrow t+1}^l$, (b) forward warping splats these pixels to the outside of border, which can be preserved with border enlargement mechanism.

at position (x, y) is sampled from \mathbf{h}_t^l at position (x', y') according to backward optical flow $(x', y') = (\mathbf{o}_{t+1 \rightarrow t})_{x, y}$. When scene motion exists, border pixels in \mathbf{h}_t^l may have no correspondence in \mathbf{h}_{t+1}^l , and are implicitly dropped in the warped result. Along with the increase of timestamps, the part of lost information of current frame in look-ahead feature becomes larger. Consequently, when aligning \mathbf{h}_{t+k}^l to current frame t , the lost information cannot be recovered (as shown in Fig. 3(a)). Such information missing leads to unexpected inferior performance.

To address the above issue, we replace backward warping in the look-ahead module with forward warping incorporated with border enlargement mechanism. As shown in Fig. 4(b), forward warping splats the value of \mathbf{h}_t^l at position (x, y) to $\mathbf{h}_{t \rightarrow t+1}^l$ at position (x', y') according to forward optical flow $(x', y') = (\mathbf{o}_{t \rightarrow t+1})_{x, y}$. When scene motion exists, border pixels in \mathbf{h}_t^l are splatted to positions out of the border. This offers an opportunity to save and reuse the out of border pixels by border enlargement mechanism as illustrated in Fig. 3(b). Specifically, forward flow is first calculated for aligning the \mathbf{h}_t^l to $(t+1)$ -th frame,

$$\mathbf{o}_{t \rightarrow t+1} = O(\mathbf{y}_t, \mathbf{y}_{t+1}). \quad (6)$$

To address the border missing issue, we first enlarge the border of the feature to a certain percent (*e.g.*, 10%) and then perform forward warping,

$$\tilde{\mathbf{h}}_{t \rightarrow t+1}^l = \vec{w}(\mathbf{h}_t^l, \mathbf{o}_{t \rightarrow t+1}). \quad (7)$$

The warped feature $\tilde{\mathbf{h}}_{t \rightarrow t+1}^l$ is further split into two separate parts,

$$\{\mathbf{b}_{t \rightarrow t+1}^l, \mathbf{h}_{t \rightarrow t+1}^l\} = \text{Split}(\tilde{\mathbf{h}}_{t \rightarrow t+1}^l). \quad (8)$$

$\mathbf{b}_{t \rightarrow t+1}^l$ denotes the enlarged border region that contains the pixels splatting out of border. $\mathbf{h}_{t \rightarrow t+1}^l$ denotes the within-frame region aligned with $(t+1)$ -th frame, and is further aggregated with \mathbf{y}_{t+1} to produce the hidden feature \mathbf{h}_{t+1}^l with




Fig. 5: Analysis of memory consumption (left) and latency (right) for three representative recurrent methods.

multiple residual blocks R_l ,

$$\mathbf{h}_{t+1}^l = R_l(\mathbf{y}_{t+1}, \mathbf{h}_{t \rightarrow t+1}^l). \quad (9)$$

In this manner, border pixels are preserved in the enlarged border $\mathbf{b}_{t \rightarrow t+1}^l$, which could be reused to recover the missed border region when aligning the look-ahead feature \mathbf{h}_{t+k}^l to current frame. In particular, \mathbf{h}_{t+k}^l is warped to t -th frame in a frame by frame manner. We first initialize $\mathbf{h}'_{t+k} = \mathbf{h}_{t+k}^l$. When warping from $(t+i)$ -th frame to $(t+i-1)$ -th frame, we first merge the border information $\mathbf{b}_{t+i-1 \rightarrow t+i}^l$ with \mathbf{h}'_{t+i} .

$$\bar{\mathbf{h}}_{t+i} = \text{Merge}(\mathbf{h}'_{t+i}, \mathbf{b}_{t+i-1 \rightarrow t+i}^l) \quad (10)$$

Then, $\bar{\mathbf{h}}_{t+i}$ is aligned to $(t+i-1)$ -th frame with backward warping,

$$\mathbf{h}'_{t+i-1} = \overleftarrow{w}(\bar{\mathbf{h}}_{t+i}, \mathbf{o}_{t+i-1 \rightarrow t+i}). \quad (11)$$

After k steps, the warped look-ahead feature can be obtained as,

$$\mathbf{h}_{t+k \rightarrow t}^l = \mathbf{h}'_t. \quad (12)$$

From Fig. 3(b), with the proposed border enlargement mechanism, the aligned look-ahead feature $\mathbf{h}_{t+2 \rightarrow t}^l$ recovers the missed border information, which brings obvious performance gain. Moreover, we note that the optical flow $\mathbf{o}_{t+i-1 \rightarrow t+i}$ reuses the optical flow for forward warping in look-ahead module (*i.e.*, $\mathcal{H}_{t+1}^l = \{\mathbf{h}_{t+1}^l, \mathbf{b}_{t \rightarrow t+1}, \mathbf{o}_{t \rightarrow t+1}\}$), thereby being effective in improving consistency and efficiency in aligning look-ahead feature back.

4 Experiments

4.1 Experimental Settings

Datasets. To evaluate our method on both synthetic and real-world noisy videos, we conduct experiments on the following datasets,

Table 1: Quantitative comparison of three representative recurrent methods on Set8 [42], FloRNN performs close to BiRNN while maintaining unidirectional.

Model	Unidirectional	Online	PSNR/SSIM
ForwardRNN	✓	✓	33.12/0.9089
FloRNN	✓	✗	33.55/0.9153
BiRNN	✗	✗	33.74/0.9192

Table 2: Ablation study for look-ahead recurrent module on Set8 dataset [42].

(a) Effects on different alignment mechanism.

Warping	Border Enlargement	PSNR/SSIM
Backward warping	✗	33.44/.9132
Forward warping	✗	33.45/.9134
Forward warping	✓	33.55/.9153

(b) Effects on the number of near-future frames k .

k	0	1	2	3	4	5
PSNR	33.12	33.47	33.53	33.55	33.53	33.51

- Set8 [42] and DAVIS [36] are two widely used synthetic Gaussian video denoising datasets.
- CRVD [51] is a real-world video denoising dataset captured in raw domain. It contains 6 indoor scenes for training, 5 indoor scenes for testing, and 10 dynamic outdoor scenes without ground-truth for visual evaluation.
- IOCV [27] is a real-world video denoising test set in sRGB domain. Each noisy video is captured by a smartphone multiple times, and the mean video is taken as ground-truth.

Implementation Details. We adopt a pretrained PWC-Net [41] as our optical flow network and fix the parameters of flow network. We find that PWC-Net generalizes well on noisy data, perhaps benefiting from additive Gaussian noise data augmentation during training [13]. For raw videos, we use the demosaiced frames as inputs to PWC-Net. Training sequences are cropped at random spatial-temporal locations, with spatial patch size 96×96 . The batch size is set to 16 and training length $T = 10$. We use ℓ_2 loss to train our network, and adopt Adam optimizer [25] with initial learning rate 10^{-4} . After 100k iterations, the learning rate is reduced to 10^{-5} until convergence. To train Gaussian denoising models, we synthesize noisy videos by adding AWGN of $\sigma \in [0, 55]$ to clean ones. For CRVD [51] dataset, as each training sequence only contains 7 frames, we mirror the training sequences to 14 frames to facilitate the model training, and crop patches with Bayer pattern preserving [28]. The evaluation is conducted on an RTX2080Ti.

Table 3: Quantitative comparison of PSNR/SSIM on the Set8 dataset [42] for Gaussian denoising. Hereinafter, **red** and **blue** indicate the best and the second best results, respectively.

Set8	VBM4D [30]	VNLB [2]	DVDNet [42]	FastDVD [43]	VNLNet [12]	PaCNet [44]	FloRNN
$\sigma = 10$	36.05/-	37.26/-	36.08/.9510	36.44/.9540	37.28/.9606	37.06/.9590	37.57/.9639
$\sigma = 20$	32.19/-	33.72/-	33.49/.9182	33.43/.9196	34.02/.9273	33.94/.9247	34.67/.9379
$\sigma = 30$	30.00/-	31.74/-	31.68/.8862	31.68/.8889	-	32.05/.8921	32.97/.9138
$\sigma = 40$	28.48/-	30.39/-	30.46/.8564	30.46/.8608	30.72/.8622	30.70/.8623	31.75/.8911
$\sigma = 50$	27.33/-	29.24/-	29.53/.8289	29.53/.8351	-	29.66/.8349	30.80/.8696
avg	30.81/-	32.47/-	32.29/.8881	32.31/.8917	-	32.68/.8946	33.55/.9153

Table 4: Quantitative comparison of PSNR/SSIM on the DAVIS dataset [36] for Gaussian denoising.

DAVIS	VBM4D [30]	VNLB [2]	DVDNet [42]	FastDVD [43]	VNLNet [12]	PaCNet [44]	FloRNN
$\sigma = 10$	37.58/-	38.85/-	38.13/.9657	38.71/.9672	39.56/.9707	39.97/.9713	40.16/.9755
$\sigma = 20$	33.88/-	35.68/-	35.70/.9422	35.77/.9405	36.53/.9464	37.10/.9470	37.52/.9564
$\sigma = 30$	31.65/-	33.73/-	34.08/.9188	34.04/.9167	-	35.07/.9211	35.89/.9440
$\sigma = 40$	30.05/-	32.32/-	32.86/.8962	32.82/.8949	33.32/.8996	33.57/.8969	34.66/.9286
$\sigma = 50$	28.80/-	31.13/-	31.85/.8745	31.86/.8747	-	32.39/.8743	33.67/.9131
avg	32.39/-	34.34/-	34.52/.9195	34.64/.9188	-	35.62/.9221	36.38/.9435


4.2 Ablation Study

We conduct the ablation study on the Set8 dataset [42] to assess the effectiveness of the proposed look-ahead recurrent module as well as its major components.

To demonstrate the effectiveness of FloRNN, we first compare it with two representative recurrent networks, *i.e.*, ForwardRNN and BiRNN. The main difference between the above three methods is how future frames are leveraged. ForwardRNN takes the forward recurrent module to propagate information from the first to the current frames, without leveraging any future information. BiRNN deploys a backward recurrent module to propagate all future frames from the last frame to current frame. Our FloRNN uses a look-ahead recurrent module for exploiting near-future frames. Table 1 and Fig. 5 show the quantitative comparison among them. From Table 1, FloRNN and BiRNN achieve better performance than ForwardRNN, which demonstrates the importance of future information in video denoising. Although BiRNN exhibits better quantitative results (0.19dB better than ours), it has a crucial issue that it can only be performed in an *offline* manner. As shown in Fig. 5, the memory consumption and latency of BiRNN grow linearly with respect to the sequence length. For example, denoising a 1000 frame sequence ($\sim 42s$ at $24fps$), BiRNN consumes several hundred GB of memory, which limits its practicality for denoising long videos on common hardware. In comparison, both ForwardRNN and our FloRNN are unidirectional algorithms with constant memory consumption and latency. With the performance close to BiRNN, our FloRNN achieves much lower computation complexity and memory occupation, which demonstrates the effectiveness of our proposed look ahead recurrent module.

Table 5: Model complexity comparison on Set8 dataset [43].

	VBM4D [30]	VNLB [2]	DVDNet [42]	FastDVD [43]	VNLNet [12]	PaCNet [44]	FloRNN
PSNR	30.81	32.47	32.29	32.31	-	32.68	33.55
#.Frame	13	13	5	5	15	15	T/2 + 3
#.Param(M)	-	-	1.33	2.48	4.52	2.87	11.82
FLOPS(G)	-	-	1231.8	661.8	2.02×10^5	6.13×10^4	3002.8
Time(s)	420.0	156.0	2.51	0.08	1.65	35.24	0.31

Fig. 6: Visual comparison for Gaussian denoising ($\sigma=40$) on ‘tractor’ sequence of the DAVIS dataset [36].

Moreover, we explore the effects of two components in the look-ahead module, *i.e.*, (i) alignment mechanism, (ii) number of near-future frames k . Quantitative results are listed in the Table 2. Although backward warping is more popular in video restoration [5, 6, 42], from Table 2(a), forward warping and backward warping achieve similar performance. It can be explained that various splatting strategies [34] have been investigated to solve the conflict of target position in forward warping, and flow based feature warping in FloRNN may be robust to the hole problem of forward warping. Besides, incorporating forward warping and border enlargement significantly reduces the amount of lost border pixels during alignment, and achieves 0.1dB performance gain. Table 2(b) analyzes the effect of k . When k increases, the benefit from near-future information first increases then becomes steady. In our following experiments, we set $k = 3$ for best performance and acceptable latency.

4.3 Results

We further compare our FloRNN with the state-of-the-art methods [2, 12, 26, 30, 31, 42–45, 51] on both synthetic and real-world video denoising datasets. FloRNN shows appealing performance.

Table 6: Quantitative comparison of PSNR/SSIM on CRVD dataset [51] for real-world raw video denoising.

Method	RViDeNet [51]	FastDVDNet [43]	EMVD [31]	EDVR [45]	FloRNN
raw	43.97/.9874	44.30/.9891	44.51/.9897	44.71/.9902	45.16/.9907
sRGB	39.95/.9792	39.91/.9812	-	40.89/.9838	41.01/.9843

Table 7: Quantitative comparison of PSNR/SSIM on IOCV dataset [27] for real-world video denoising in sRGB domain.

Method	Noisy	CVMSSt-SVD [26]	VBM4D [30]	FastDVD [43]	VNLNet [12]	FloRNN
HUAWEI_BC	38.11/.9593	40.80/.9834	41.19/.9830	41.26/.9857	41.35/.9868	42.28/.9880
HUAWEI_FC	38.58/.9413	38.71/.9780	38.76/.9785	38.03/.9776	38.79/.9809	39.57/.9828
OPPO_BC	32.06/.9071	33.44/.9508	33.26/.9456	33.05/.9476	33.56/.9544	33.75/.9545
OPPO_FC	36.90/.9447	39.66/.9791	39.56/.9785	39.06/.9751	40.11/.9823	40.31/.9821

Gaussian denoising on Set8 and DAVIS datasets. Table 3 and 4 lists the quantitative comparison on Set8 [42] and DAVIS [36] datasets. From the table, our FloRNN outperforms the state-of-the-art methods by a large margin on both datasets in terms of PSNR/SSIM, which demonstrates its superiority on video denoising applications. We also analyse the model complexity of each method in table 5, and found that FloRNN achieves good trade-off between performance and running time. Specifically, FloRNN surpasses non-local based PaCNet [44] by 0.87dB on Set8, while hundred times faster. Compare to FastDVDNet [43] which is especially suggested for efficiency, the performance gain of FloRNN is promoted to 1.24dB, with 4 times slower. Experiments on DAVIS dataset shows similar results, which demonstrates that FloRNN generates best results with good efficiency. Fig. 6 illustrates the qualitative comparison on a dynamic scene with noise level $\sigma=40$. Despite the large motion and severe noise, FloRNN could clearly recover the characters, which are hard to recognize in other results. More results are provided in the supplementary material.

Raw domain video denoising on CRVD dataset. We further evaluate our FloRNN on real-world video denoising datasets. On CRVD dataset [51], EMVD [31] adopts forward recurrent mechanism designed for mobile devices, but performs inferior due to lack of future information when scaling up on GPUs. EDVR [45] utilizes pyramid deformable convolution for alignment and achieves good results, but its multi-frame convolutional backbone limits the long term propagation. Instead, from Table 6, our FloRNN benefits both long term propagation of historical frames and near future information, and outperforms EDVR by 0.45dB on raw domain. The performance on sRGB domain is calculated by rendering the raw results into sRGB ones with a pretrained ISP network. Our performance gain decreases on sRGB domain, which is possibly due to the inaccuracy of the ISP module.

sRGB domain video denoising on IOCV dataset. As IOCV [27] does not provide training sets, we apply our Gaussian denoising models for IOCV videos, and tune the noise level for each subset to get best results the same




Fig. 7: Visual comparison on the real-world IOCV dataset [27].




Fig. 8: Visual comparison of an outdoor scene on the CRVD dataset [51], we render the results in raw domain to sRGB domain with a pretrained ISP.

as [27]. Nonetheless, FloRNN achieves favorable results and surpasses competing methods (see Table 7). Visual comparison in Fig. 7 and Fig. 8 show that our FloRNN is able to remove the real-world video noise and recover fine-grained details.

5 Conclusion

In this paper, we propose a novel recurrent network for unidirectional video denoising. Look-ahead recurrent module is sufficient to exploit near-future frames in a forward manner. Combining with forward module, FloRNN achieves near BiRNN performance with constant memory occupation and latency. By analyzing backward and forward warping mechanisms, we found incorporating forward warping and border enlargement is favorable to address the border information missing problem during alignment. Experimental results demonstrate the superiority of the proposed method in removing both synthetic and real world noise. Our unidirectional video denoising algorithm is beneficial to various video applications, *e.g.*, video conference, live streaming, and could dynamically adjust the look-ahead step k for balance denoising results and latency.

Acknowledgement. This work is partially supported by the National Natural Science Foundation of China (NSFC) under Grant No.s 62006064 and U19A2073.

References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1692–1700 (2018) 4
2. Arias, P., Morel, J.M.: Video denoising via empirical bayesian estimation of space-time patches. *Journal of Mathematical Imaging and Vision* **60**(1), 70–93 (2018) 4, 11, 12
3. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05). vol. 2, pp. 60–65. IEEE (2005) 4
4. Buades, A., Lisani, J.L., Miladinović, M.: Patch-based video denoising with optical flow estimation. *IEEE Transactions on Image Processing* **25**(6), 2573–2586 (2016) 4
5. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: Basicvsr: The search for essential components in video super-resolution and beyond. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4947–4956 (2021) 2, 4, 5, 6, 12
6. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5972–5981 (2022) 2, 3, 4, 12
7. Chen, X., Song, L., Yang, X.: Deep rnns for video denoising. In: Applications of digital image processing XXXIX. vol. 9971, p. 99711T. international Society for optics and Photonics (2016) 4
8. Claus, M., van Gemert, J.: Videnn: Deep blind video denoising. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019) 2, 4
9. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing* **16**(8), 2080–2095 (2007) 4
10. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 764–773 (2017) 2
11. Davy, A., Ehret, T., Morel, J.M., Arias, P., Facciolo, G.: A non-local cnn for video denoising. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 2409–2413. IEEE (2019) 2
12. Davy, A., Ehret, T., Morel, J.M., Arias, P., Facciolo, G.: Video denoising by combining patch search and cnns. *Journal of Mathematical Imaging and Vision* **63**(1), 73–88 (2021) 2, 4, 11, 12, 13
13. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2758–2766 (2015) 10
14. Elad, M., Aharon, M.: Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing* **15**(12), 3736–3745 (2006) 4
15. Fuoli, D., Gu, S., Timofte, R.: Efficient video super-resolution through recurrent latent space propagation. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). pp. 3476–3485. IEEE (2019) 4

16. Gu, S., Zhang, L., Zuo, W., Feng, X.: Weighted nuclear norm minimization with application to image denoising. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2862–2869 (2014) [4](#)
17. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1712–1722 (2019) [4](#)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) [4](#)
19. Huang, Y., Wang, W., Wang, L.: Video super-resolution via bidirectional recurrent convolutional networks. IEEE transactions on pattern analysis and machine intelligence **40**(4), 1015–1028 (2017) [4](#)
20. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. PMLR (2015) [4](#)
21. Isobe, T., Jia, X., Gu, S., Li, S., Wang, S., Tian, Q.: Video super-resolution with recurrent structure-detail network. In: European Conference on Computer Vision. pp. 645–660. Springer (2020) [4](#)
22. Kim, D., Woo, S., Lee, J.Y., Kweon, I.S.: Deep video inpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5792–5801 (2019) [4](#)
23. Kim, D., Woo, S., Lee, J.Y., Kweon, I.S.: Recurrent temporal aggregation framework for deep video inpainting. IEEE transactions on pattern analysis and machine intelligence **42**(5), 1038–1052 (2019) [2](#), [4](#)
24. Kim, Y., Soh, J.W., Park, G.Y., Cho, N.I.: Transfer learning from synthetic to real-noise denoising with adaptive instance normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3482–3492 (2020) [4](#)
25. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
26. Kong, Z., Yang, X.: Color image and multispectral image denoising using block diagonal representation. IEEE Transactions on Image Processing **28**(9), 4247–4259 (2019) [12](#), [13](#)
27. Kong, Z., Yang, X., He, L.: A comprehensive comparison of multi-dimensional image denoising methods. arXiv preprint arXiv:2011.03462 (2020) [3](#), [10](#), [13](#), [14](#)
28. Liu, J., Wu, C.H., Wang, Y., Xu, Q., Zhou, Y., Huang, H., Wang, C., Cai, S., Ding, Y., Fan, H., et al.: Learning raw image denoising with bayer pattern unification and bayer preserving augmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019) [10](#)
29. Liu, J., Yang, W., Yang, S., Guo, Z.: Erase or fill? deep joint recurrent rain removal and reconstruction in videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3233–3242 (2018) [4](#)
30. Maggioni, M., Boracchi, G., Foi, A., Egiazarian, K.: Video denoising, deblocking, and enhancement through separable 4-d nonlocal spatiotemporal transforms. IEEE Transactions on image processing **21**(9), 3952–3966 (2012) [4](#), [11](#), [12](#), [13](#)
31. Maggioni, M., Huang, Y., Li, C., Xiao, S., Fu, Z., Song, F.: Efficient multi-stage video denoising with recurrent spatio-temporal fusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3466–3475 (2021) [2](#), [4](#), [12](#), [13](#)

32. Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R., Carroll, R.: Burst denoising with kernel prediction networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2502–2510 (2018) [2](#), [4](#)
33. Nah, S., Son, S., Lee, K.M.: Recurrent neural networks with intra-frame iterations for video deblurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8102–8111 (2019) [4](#)
34. Niklaus, S., Liu, F.: Softmax splatting for video frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5437–5446 (2020) [12](#)
35. Plotz, T., Roth, S.: Benchmarking denoising algorithms with real photographs. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1586–1595 (2017) [4](#)
36. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675 (2017) [10](#), [11](#), [12](#), [13](#)
37. Rong, X., Demandolx, D., Matzen, K., Chatterjee, P., Tian, Y.: Burst denoising via temporally shifted wavelet transforms. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16. pp. 240–256. Springer (2020) [4](#)
38. Sajjadi, M.S., Vemulapalli, R., Brown, M.: Frame-recurrent video super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6626–6634 (2018) [4](#)
39. Shen, W., Bao, W., Zhai, G., Chen, L., Min, X., Gao, Z.: Blurry video frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5114–5123 (2020) [2](#), [4](#)
40. Son, H., Lee, J., Lee, J., Cho, S., Lee, S.: Recurrent video deblurring with blur-invariant motion estimation and pixel volumes. ACM Transactions on Graphics (TOG) **40**(5), 1–18 (2021) [2](#), [4](#)
41. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018) [2](#), [6](#), [10](#)
42. Tassano, M., Delon, J., Veit, T.: Dvdnet: A fast network for deep video denoising. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 1805–1809. IEEE (2019) [2](#), [4](#), [10](#), [11](#), [12](#), [13](#)
43. Tassano, M., Delon, J., Veit, T.: Fastdvdnet: Towards real-time deep video denoising without flow estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) [2](#), [4](#), [11](#), [12](#), [13](#)
44. Vaksman, G., Elad, M., Milanfar, P.: Patch craft: Video denoising by deep modeling and patch matching. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2157–2166 (October 2021) [2](#), [3](#), [4](#), [11](#), [12](#), [13](#)
45. Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: Edvr: Video restoration with enhanced deformable convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019) [12](#), [13](#)
46. Xia, Z., Perazzi, F., Gharbi, M., Sunkavalli, K., Chakrabarti, A.: Basis prediction networks for effective burst denoising with large kernels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11844–11853 (2020) [2](#), [4](#)
47. Xiang, X., Tian, Y., Zhang, Y., Fu, Y., Allebach, J.P., Xu, C.: Zooming slow-mo: Fast and accurate one-stage space-time video super-resolution. In: Proceedings of

- the IEEE/CVF conference on computer vision and pattern recognition. pp. 3370–3379 (2020) [4](#)
48. Xu, X., Li, M., Sun, W., Yang, M.H.: Learning spatial and spatio-temporal pixel aggregations for image and video denoising. *IEEE Transactions on Image Processing* **29**, 7153–7165 (2020) [2](#), [4](#)
 49. Yang, W., Liu, J., Feng, J.: Frame-consistent recurrent video deraining with dual-level flow. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1661–1670 (2019) [4](#)
 50. Yang, W., Tan, R.T., Feng, J., Wang, S., Cheng, B., Liu, J.: Recurrent multi-frame deraining: Combining physics guidance and adversarial learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021) [2](#), [4](#)
 51. Yue, H., Cao, C., Liao, L., Chu, R., Yang, J.: Supervised raw video denoising with a benchmark dataset on dynamic scenes. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020) [2](#), [3](#), [4](#), [10](#), [12](#), [13](#), [14](#)
 52. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing* **26**(7), 3142–3155 (2017) [4](#)
 53. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing* **27**(9), 4608–4622 (2018) [4](#)
 54. Zhong, Z., Gao, Y., Zheng, Y., Zheng, B.: Efficient spatio-temporal recurrent neural network for video deblurring. In: *European Conference on Computer Vision*. pp. 191–207. Springer (2020) [4](#)
 55. Zhuo, S., Jin, Z., Zou, W., Li, X.: Ridnet: Recursive information distillation network for color image denoising. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. pp. 0–0 (2019) [4](#)
 56. Zoran, D., Weiss, Y.: From learning models of natural image patches to whole image restoration. In: *2011 International Conference on Computer Vision*. pp. 479–486. IEEE (2011) [4](#)