

FastDVDnet: Towards Real-Time Deep Video Denoising Without Flow Estimation

Matias Tassano
GoPro France
mtassano@gopro.com

Julie Delon
MAP5, Université de Paris & IUF
julie.delon@parisdescartes.fr

Thomas Veit
GoPro France
tveit@gopro.com

Abstract

In this paper, we propose a state-of-the-art video denoising algorithm based on a convolutional neural network architecture. Until recently, video denoising with neural networks had been a largely under explored domain, and existing methods could not compete with the performance of the best patch-based methods. The approach we introduce in this paper, called FastDVDnet, shows similar or better performance than other state-of-the-art competitors with significantly lower computing times. In contrast to other existing neural network denoisers, our algorithm exhibits several desirable properties such as fast runtimes, and the ability to handle a wide range of noise levels with a single network model. The characteristics of its architecture make it possible to avoid using a costly motion compensation stage while achieving excellent performance. The combination between its denoising performance and lower computational load makes this algorithm attractive for practical denoising applications. We compare our method with different state-of-art algorithms, both visually and with respect to objective quality metrics.

1. Introduction

Despite the immense progress made in recent years in photographic sensors, noise reduction remains an essential step in video processing, especially when shooting conditions are challenging (low light, small sensors, etc.).

Although image denoising has remained a very active research field through the years, too little work has been devoted to the restoration of digital videos. It should be noted, however, that some crucial aspects differentiate these two problems. On the one hand, a video contains much more information than a still image, which could help in the restoration process. On the other hand, video restoration requires good temporal coherency, which makes the restoration process much more demanding. Furthermore, since all recent cameras produce videos in high definition—or even

larger—very fast and efficient algorithms are needed.

In this paper we introduce another network for deep video denoising: FastDVDnet. This algorithm builds on DVDnet [38], but at the same time introduces a number of important changes with respect to its predecessor. Most notably, instead of employing an explicit motion estimation stage, the algorithm is able to implicitly handle motion thanks to the traits of its architecture. This results in a state-of-the-art algorithm which outputs high quality denoised videos while featuring very fast running times—even thousands of times faster than other relevant methods.

1.1. Image denoising

Contrary to video denoising, image denoising has enjoyed consistent popularity in past years. A myriad of new image denoising methods based on deep learning techniques have drawn considerable attention due to their outstanding performance. Schmidt and Roth proposed in [31] the cascade of shrinkage fields method. The trainable nonlinear reaction diffusion model proposed by Chen and Pock in [8] builds on the former. In [5], a multi-layer perceptron was successfully applied for image denoising. Methods such as these achieve performances comparable to those of well-known patch-based algorithms such as BM3D [10] or non-local Bayes (NLB [21]). However, their limitations include performance restricted to specific forms of prior, or the fact that a different set of weights must be trained for each noise level.

Another widespread approach involves the use of convolutional neural networks (CNN), e.g. RBDN [30], MWCNN [24], DnCNN [43], and FFDNet [44]. Their performance compares favorably to other state-of-the-art image denoising algorithms, both quantitatively and visually. These methods are composed of a succession of convolutional layers with nonlinear activation functions in between them. A salient feature that these CNN-based methods present is the ability to denoise several levels of noise with only one trained model. Proposed by Zhang *et al.* in [43], DnCNN is an end-to-end trainable deep CNN for image denoising. One of its main features is that it implements resid-

ual learning [16], i.e. it estimates the noise existent in the input image rather than the denoised image. In a following paper [44], Zhang *et al.* proposed FFDNet, which builds upon the work done for DnCNN. More recently, the approaches proposed in [28, 23] combine neural architectures with non-local techniques.

1.2. Video denoising

Video denoising is much less explored in the literature. The majority of recent video denoising methods are patch-based. We note in particular an extension of the popular BM3D to video denoising, V-BM4D [25], and Video non-local Bayes (VNLB [2]). Neural network methods for video denoising have been even rarer than patch-based approaches. The algorithm in [7] by Chen *et al.* is one of the first to approach this problem with recurrent neural networks. However, their algorithm only works on grayscale images and it does not achieve satisfactory results, probably due to the difficulties associated with training recurring neural networks [26]. Vogels *et al.* proposed in [39] an architecture based on kernel-predicting neural networks able to denoise Monte Carlo rendered sequences. The Video Non-Local Network (VNLnet [11]) fuses a CNN with a self-similarity search strategy. For each patch, the network finds the most similar patches via its first non-trainable layer, and this information is later used by the CNN to predict the clean image. In [38], Tassano *et al.* proposed DVDnet, which splits the denoising of a given frame in two separate denoising stages. Like several other methods, it relies on the estimation of motion of neighboring frames. Other very recent blind denoising approaches include the work by Ehret *et al.* [13] and ViDeNN [9]. The latter shares with DVDnet the idea of performing denoising in two steps. However, contrary to DVDnet, ViDeNN does not employ motion estimation. Similarly to both DVDnet and ViDeNN, the use of spatio-temporal CNN blocks in restoration tasks has been also featured in [39, 6]. Nowadays, the state-of-the-art is defined by DVDnet, VNLnet and VNLB. VNLB and VNLnet show the best performances for small values of noise, while DVDnet yields better results for larger values of noise. Both DVDnet and VNLnet feature significantly faster inference times than VNLB. As we will see, the performance of the method we introduce in this paper compares to the performance of the state-of-the-art, while featuring even faster runtimes.

2. FastDVDnet

For video denoising algorithms, temporal coherence and flickering removal are crucial aspects in the perceived quality of the results [33, 32]. In order to achieve these, an algorithm must make use of the temporal information existent in neighboring frames when denoising a given frame of an image sequence. In general, most previous approaches based

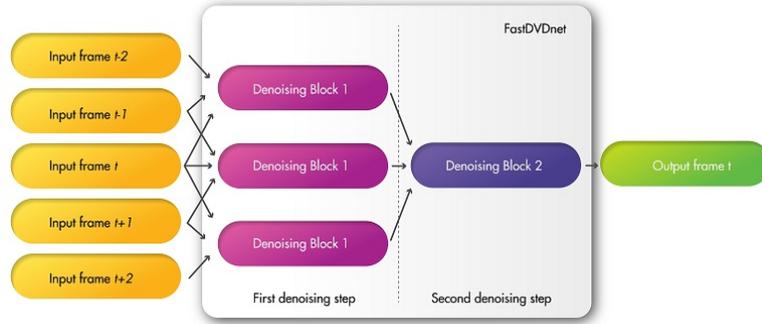
on deep learning have failed to employ this temporal information effectively. Successful state-of-the-art algorithms rely mainly on two factors to enforce temporal coherence in the results, namely the extension of search regions from spatial neighborhoods to volumetric neighborhoods, and the use of motion estimation.

The use of volumetric (i.e. spatio-temporal) neighborhoods implies that when denoising a given pixel (or patch), the algorithm is going to look for similar pixels (patches) not only in the reference frame, but also in adjacent frames of the sequence. The benefits of this are two-fold. First, the temporal neighbors provide additional information which can be used to denoise the reference frame. Second, using temporal neighbors helps to reduce flickering as the residual error in each frame will be correlated.

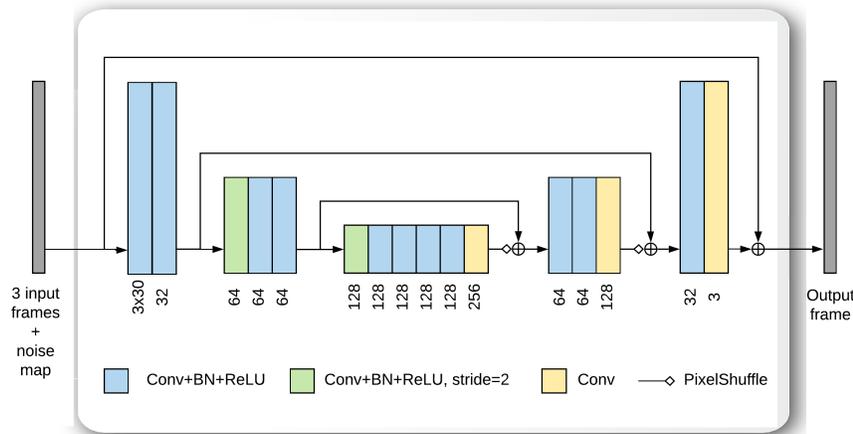
Videos feature a strong temporal redundancy along motion trajectories. This fact should facilitate denoising videos with respect to denoising images. Yet, this added information in the temporal dimension also creates an extra degree of complexity which could be difficult to tackle. In this context, motion estimation and/or compensation has been employed in a number of video denoising algorithms to help to improve denoising performance and temporal consistency [22, 38, 2, 25, 4].

We thus incorporated these two elements into our architecture. However, our algorithm does not include an explicit motion estimation/compensation stage. The capacity of handling the motion of objects is inherently embedded into the proposed architecture. Indeed, our architecture is composed of a number of modified U-Net [29] blocks (see Section 2.1 for more details about these blocks). Multi-scale, U-Net-like architectures have been shown to have the ability to learn misalignment [42, 12]. Our cascaded architecture increases this capacity of handling movement even further. In contrast to [38], our architecture is trained end-to-end without optical flow alignment, which avoids distortions and artifacts due to erroneous flow. As a result, we are able to eliminate a costly dedicated motion compensation stage without sacrificing performance. This leads to an important reduction of runtimes: our algorithm runs three orders of magnitude faster than VNLB, and an order of magnitude faster than DVDnet and VNLnet.

Figure 1a displays a diagram of the architecture of our method. When denoising a given frame at time t , $\tilde{\mathbf{I}}_t$, its $2T = 4$ neighboring frames are also taken as inputs. That is, the inputs of the algorithm will be $\{\tilde{\mathbf{I}}_{t-2}, \tilde{\mathbf{I}}_{t-1}, \tilde{\mathbf{I}}_t, \tilde{\mathbf{I}}_{t+1}, \tilde{\mathbf{I}}_{t+2}\}$. The model is composed of different spatio-temporal denoising blocks, assembled in a cascaded two-step architecture. These denoising blocks are all similar, and consist of a modified U-Net model which takes three frames as inputs. The three blocks in the first denoising step share the same weights, which leads to a reduction of memory requirements of the model and facili-



(a)



(b)

Figure 1. *Architecture used in FastDVDnet.* (a) A high-level diagram of the architecture. Five consecutive frames are used to denoise the middle frame. The frames are taken as triplets of consecutive frames and input to the *Denoising Blocks 1*. The instances of these blocks have all the same weights. The triplet composed by the outputs of these blocks are used as inputs for *Denoising Block 2*. The output of the latter is the estimate of the central input frame (*Input frame t*). Both *Denoising Block 1* and *Denoising Block 2* share the same architecture, which is shown in (b). The denoising blocks of FastDVDnet are composed of a modified multi-scale U-Net.

tates the training of the network. Similar to [44, 14], a noise map is also included as input, which allows the processing of spatially varying noise [37]. In particular, the noise map is a separate input which provides information to the network about the distribution of the noise at the input. This information is encoded as the expected per-pixel standard deviation of this noise. For instance, when denoising Gaussian noise, the noise map will be constant; when denoising Poisson noise, the noise map will depend on the intensity of the image. Indeed, the noise map can be used as a user-input parameter to control the trade-off between noise removal vs. detail preservation (see for example the online demo in [37]). In other cases, such as JPEG denoising, the noise map can be estimated by means of an additional CNN [15]. The use of a noise map has been shown to improve denoising performance, particularly when treating spatially vari-

ant noise [3]. Contrary to other denoising algorithms, our denoiser takes no other parameters as inputs apart from the image sequence and the estimation of the input noise.

Observe that experiments presented in this paper focus on the case of additive white Gaussian noise (AWGN). Nevertheless, this algorithm can be extended to other types of noise, e.g. spatially varying noise (e.g. Poissonian). Let \mathbf{I} be a noiseless image, while $\tilde{\mathbf{I}}$ is its noisy version corrupted by a realization of zero-mean white Gaussian noise \mathbf{N} of standard deviation σ , then

$$\tilde{\mathbf{I}} = \mathbf{I} + \mathbf{N}. \quad (1)$$

2.1. Denoising blocks

Both denoising blocks displayed in Fig. 1a, *Denoising Block 1* and *Denoising Block 2*, consist of a modified U-Net architecture. All the instances of *Denoising Block 1*

share the same weights. U-Nets are essentially a multi-scale encoder-decoder architecture, with skip-connections [16] that forward the output of each one of the encoder layers directly to the input of the corresponding decoder layers. A more detailed diagram of these blocks is shown in Fig. 1b. Our denoising blocks present some differences with respect to the standard U-Net:

- The encoder has been adapted to take three frames and a noise map as inputs
- The upsampling in the decoder is performed with a *PixelShuffle* layer [34], which helps reducing gridding artifacts. Please see the supplementary materials for more information about this layer.
- The merging of the features of the encoder with those of the decoder is done with a pixel-wise addition operation instead of a channel-wise concatenation. This results in a reduction of memory requirements
- Blocks implement residual learning—with a residual connection between the central noisy input frame and the output—, which has been observed to ease the training process [37]

The design characteristics of the denoising blocks make a good compromise between performance and fast running times. These denoising blocks are composed of a total of $D = 16$ convolutional layers. In most layers, the outputs of its convolutional layers are followed by point-wise *ReLU* [20] activation functions $ReLU(\cdot) = \max(\cdot, 0)$, except for the last layer. Batch normalization layers (*BN* [17]) are placed between the convolutional and *ReLU* layers.

3. Discussion

Explicit flow estimation is avoided in FastDVDnet. However, in order to maintain performance, we needed to introduce a number of techniques to handle motion and to effectively employ temporal information. These techniques are discussed further in this section. Please see the supplementary materials for more details about ablation studies.

3.1. Two-step denoising

Similarly to DVDnet and ViDeNN, FastDVDnet features a two-step cascaded architecture. The motivation behind this is to effectively employ the information existent in the temporal neighbors, and to enforce the temporal correlation of the remaining noise in output frames. To prove that the two-step denoising is a necessary feature, we conducted the following experiment: we modified a *Denoising Block* of FastDVDnet (see Fig. 1b) to take five frames as inputs instead of three, which we will refer to as *Den_Block_5inputs*. In this way, the same amount of temporal neighboring

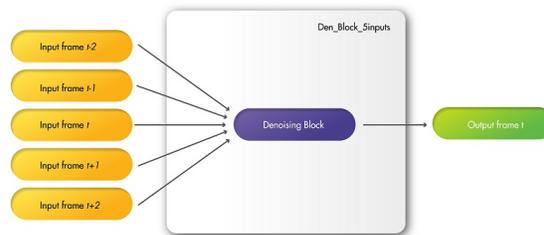


Figure 2. Architecture of the *Den_Block_5inputs* denoiser.

frames are considered and the same information as in FastDVDnet is processed by this new denoiser. A diagram of the architecture of this model is shown in Fig. 2. We then trained this new model and compared the results of denoising of sequences against the results of FastDVDnet (see Section 4 for more details about the training process).

It was observed that the cascaded architecture of FastDVDnet presents a clear advantage on *Den_Block_5inputs*, with differences in PSNR of up to 0.9dB. Please refer to the supplementary materials for more details. Additionally, results by *Den_Block_5inputs* present a sharp increase on temporal artifacts—flickering. Despite it being a multi-scale architecture, *Den_Block_5inputs* cannot handle the motion of objects in the sequences as well as the two-step architecture of FastDVDnet can. Overall, the two-step architecture shows superior performance with respect to the one-step architecture.

3.2. Multi-scale architecture and end-to-end training

In order to investigate the importance of using multi-scale denoising blocks in our architecture, we conducted the following experiment: we modified the FastDVDnet architecture by replacing its *Denoising Blocks* by the denoising blocks of DVDnet. This results in a two-step cascaded architecture, with single-scale denoising blocks, trained end-to-end, and with no compensation of motion in the scene. In our tests, it was observed that the usage of multi-scale denoising blocks improves denoising results considerably. Please refer to the supplementary materials for more details.

We also experimented with training the multi-scale denoising blocks in each step of FastDVDnet separately—as done in DVDnet. Although the results in this case certainly improved with respect to the case of the single-scale denoising blocks described above, a noticeable flickering remained in the outputs. Switching from this separate training to an end-to-end training helped to reduce temporal artifacts considerably.

3.3. Handling of motion

Apart from the reduction of runtimes, avoiding the use of motion compensation by means of optical flow has an

additional benefit. Video denoising algorithms that depend explicitly on motion estimation techniques often present artifacts due to erroneous flow in challenging cases, such as occlusions or strong noise. The different techniques discussed in this section—namely a multi-scale of the denoising blocks, the cascaded two-step denoising architecture, and end-to-end training—not only provide FastDVDnet the ability to handle motion, but also help avoid artifacts related to erroneous flow estimation. Also, and similarly to [43, 38, 37], the denoising blocks of FastDVDnet implement residual learning, which helps to improve the quality of results a step further. Figure 3 shows an example on artifacts due to erroneous flow on three consecutive frames and of how the multi-scale architecture of FastDVDnet is able to avoid them.

4. Training details

The training dataset consists of input-output pairs

$$P_t^j = \left\{ \left((S_t^j, \mathbf{M}^j), \mathbf{I}_t^j \right) \right\}_{j=0}^{m_t},$$

where $S_t^j = (\tilde{\mathbf{I}}_{t-2}^j, \tilde{\mathbf{I}}_{t-1}^j, \tilde{\mathbf{I}}_t^j, \tilde{\mathbf{I}}_{t+1}^j, \tilde{\mathbf{I}}_{t+2}^j)$ is a collection of $2T + 1 = 5$ spatial patches cropped at the same location in contiguous frames, and \mathbf{I}^j is the clean central patch of the sequence. These are generated by adding AWGN of $\sigma \in [5, 50]$ to clean patches of a given sequence, and the corresponding noise map \mathbf{M}^j is built in this case constant with all its elements equal to σ . Spatio-temporal patches are randomly cropped from randomly sampled sequences of the training dataset.

A total of $m_t = 384000$ training samples are extracted from the training set of the DAVIS database [18]. The spatial size of the patches is 96×96 , while the temporal size is $2T + 1 = 5$. The spatial size of the patches was chosen such that the resulting patch size in the coarser scale of the *Denoising Blocks* is 32×32 . The loss function is

$$\mathcal{L}(\theta) = \frac{1}{2m_t} \sum_{j=1}^{m_t} \left\| \hat{\mathbf{I}}_t^j - \mathbf{I}_t^j \right\|^2, \quad (2)$$

where $\hat{\mathbf{I}}_t^j = \mathcal{F}((S_t^j, \mathbf{M}^j); \theta)$ is the output of the network, and θ is the set of all learnable parameters.

The architecture has been implemented in PyTorch [27], a popular machine learning library. The ADAM algorithm [19] is applied to minimize the loss function, with all its hyper-parameters set to their default values. The number of epochs is set to 80, and the mini-batch size is 96. The scheduling of the learning rate is also common to both cases. It starts at $1e-3$ for the first 50 epochs, then changes to $1e-4$ for the following 10 epochs, and finally switches to $1e-6$ for the remaining of the training. In other words, a learning rate step decay is used in conjunction

with ADAM. The mix of learning rate decay and adaptive rate methods has also been applied to other deep learning projects [36, 41], usually with positive results. Data is augmented by introducing rescaling by different scale factors and random flips. During the first 60 epochs, the orthogonalization of the convolutional kernels is applied as a means of regularization. It has been observed that initializing the training with orthogonalization may be beneficial to performance [44, 37].

5. Results

Two different testsets were used for benchmarking our method: the DAVIS-test testset, and Set8, which is composed of 4 color sequences from the *Derf's Test Media collection*¹ and 4 color sequences captured with a GoPro camera. The DAVIS set contains 30 color sequences of resolution 854×480 . The sequences of Set8 have been down-scaled to a resolution of 960×540 . In all cases, sequences were limited to a maximum of 85 frames. We used the DeepFlow algorithm [40] to compute flow maps for DVDnet and VNLB. VNLnet requires models trained for specific noise levels. As no model is provided for $\sigma = 30$, no results are shown for this noise level in either of the tables. We also compare our method to a commercial blind denoising software, Neat Video (NV [1]). For NV, its automatic noise profiling settings were used to manually denoise the sequences of Set8. Note that values shown are the average for all sequences in the testset, the PSNR of a sequence is computed as the average of the PSNRs of each frame.

In general, both DVDnet and FastDVDnet output sequences which feature remarkable temporal coherence. Flickering rendered by our methods is notably small, especially in flat areas, where patch-based algorithms often leave behind low-frequency residual noise. An example can be observed in Fig. 4 (which is best viewed in digital format). Temporally decorrelated low-frequency noise in flat areas appears as particularly bothersome for the viewer. More video examples can be found in the supplementary materials and on the website of the algorithm. The reader is encouraged to watch these examples to compare the visual quality of the results of our methods.

Patch-based methods are prone to surpassing DVDnet and FastDVDnet in sequences with a large portion of repetitive structures as these methods exploit the non-local similarity prior. On the other hand, our algorithms handle non-repetitive textures very well, see e.g. the clarity of the denoised text and vegetation in Fig. 5.

Table 1 shows a comparison of PSNR and ST-RRED on the Set8 and DAVIS dataset, respectively. The Spatio-Temporal Reduced Reference Entropic Differences (ST-RRED) is a high performing reduced-reference video qual-

¹<https://media.xiph.org/video/derf>

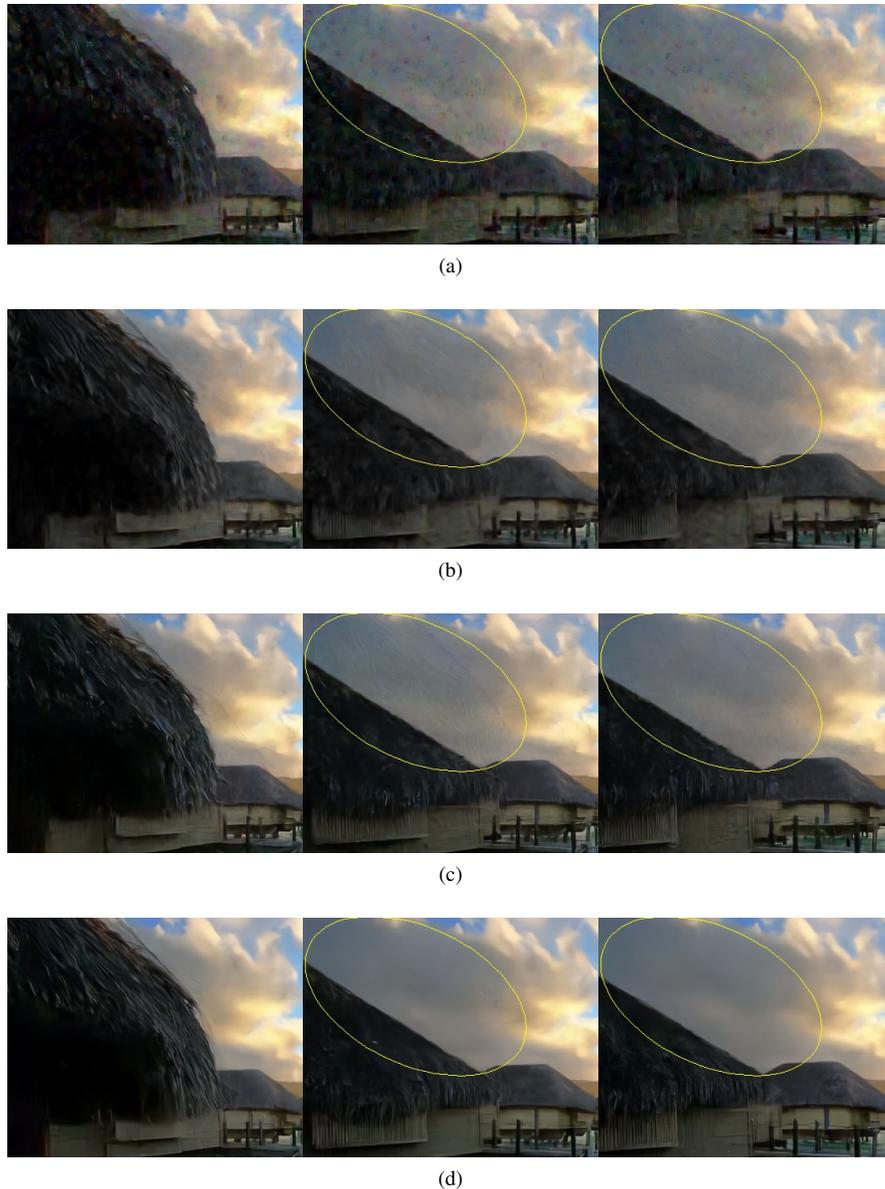


Figure 3. *Motion artifacts due to occlusion.* Three consecutive frames of the results of the ‘hypersmooth’ sequence, $\sigma = 50$ (a) V-BM4D. (b) VNLB. (c) DVDnet. (d) FastDVDnet. Video denoising algorithms that depend explicitly on motion estimation techniques often present artifacts due to erroneous flow in challenging cases. In the example above, the occlusion of the front building leads to motion artifacts in the results of V-BM4D, VNLB, and DVDnet. Explicit motion compensation is avoided in the architecture of FastDVDnet. Indeed, the network is able to implicitly handle motion due to its design characteristics. Best viewed in digital format.

ity assessment metric [35]. This metric not only takes into account image quality, but also temporal distortions in the video. We computed the ST-RRED scores with the implementation provided by the *scikit-video* library².

It can be observed that for smaller values of noise, VNLB performs better on Set8. Indeed, DVDnet tends to over denoise in some of these cases. FastDVDnet and VNLnet are the best performing algorithms on DAVIS for small sigmas

in terms of PSNR and ST-RRED, respectively. However, for larger values of noise DVDnet surpasses VNLB. FastDVDnet performs consistently well in all cases, which is a remarkable feat considering that it runs 80 times faster than DVDnet, 26 times faster than VNLnet, and more than 4000 times faster than VNLB (see Section 6). Contrary to other denoisers based on CNNs—e.g. VNLnet—, our algorithms are able to denoise different noise levels with only one trained model. On top of this, the use of methods involve no hand-tuned parameters, since they only take

²<http://www.scikit-video.org>

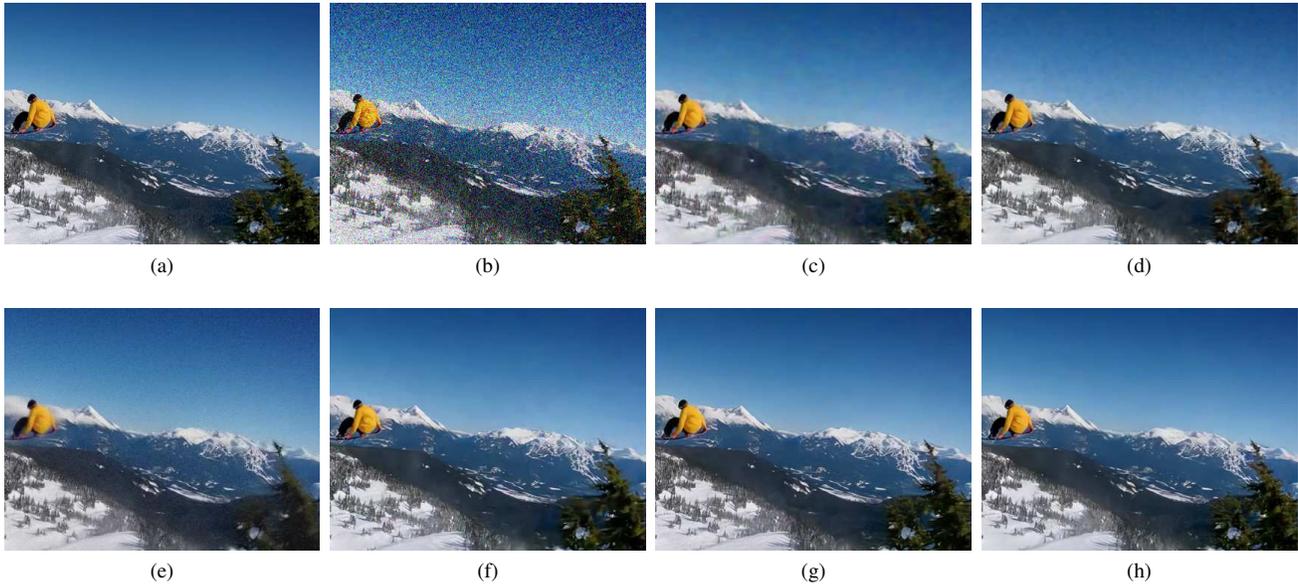


Figure 4. Comparison of results of the 'snowboarding' sequence. (a) Clean frame. (b) Noisy frame $\sigma = 40$. (c) V-BM4D. (d) VNLB. (e) NV. (f) VNLnet. (g) DVDnet. (h) FastDVDnet. Patch-based methods (V-BM4D, VNLB, and even VNLnet) struggle with noise in flat areas, such as the sky, and leave behind medium-to-low-frequency noise. This leads to results with noticeable flickering, as the remaining noise is temporally decorrelated. On the other hand, DVDnet and FastDVDnet output very convincing and visually pleasant results. Best viewed in digital format.

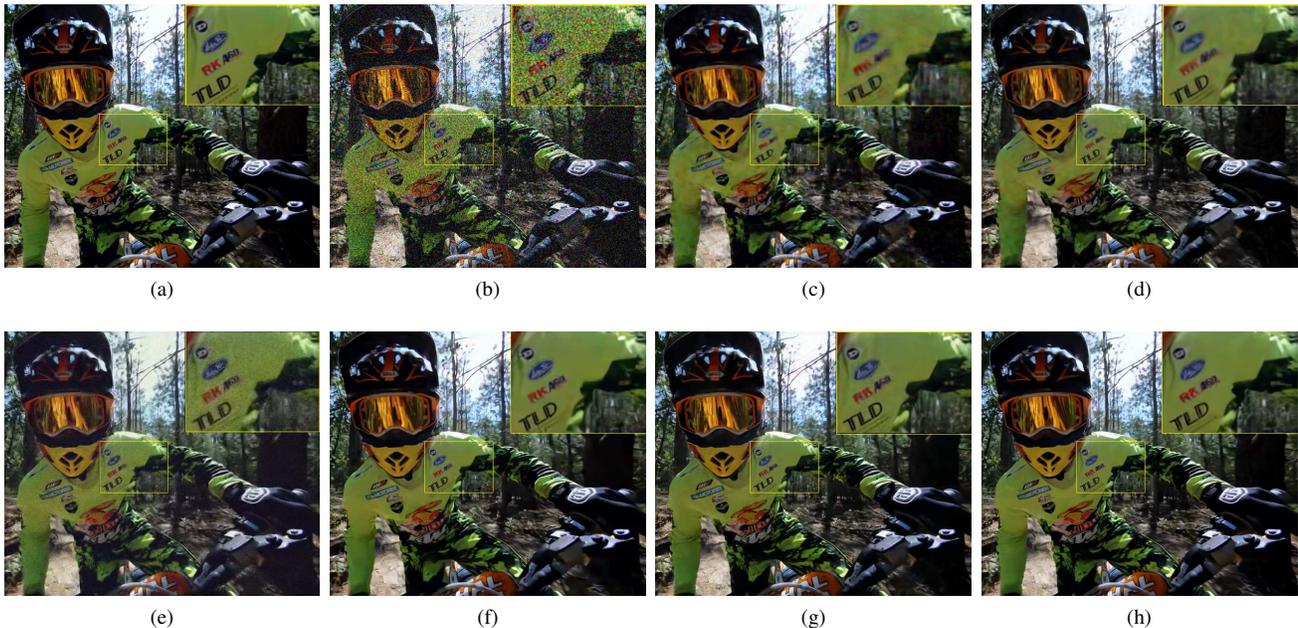


Figure 5. Comparison of results of the 'motorbike' sequence. (a) Clean frame. (b) Noisy frame $\sigma = 50$. (c) V-BM4D. (d) VNLB. (e) NV. (f) VNLnet. (g) DVDnet. (h) FastDVDnet. Note the clarity of the denoised text, and the lack of chroma noise for FastDVDnet, DVDnet, and VNLnet. Best viewed in digital format.

the image sequence and the estimation of the input noise as inputs. Table 2 displays a comparison with ViDeNN. This algorithm has not actually been trained for AWGN, but for clipped AWGN. Then, a FastDVDnet model to de-

noise clipped AWGN was trained for this case, which we call *FastDVDnet_clipped*. It can be observed that the performance of *FastDVDnet_clipped* is superior to the performance of ViDeNN by a wide margin.

Table 1. Comparison of PSNR / ST-RRED on the Set8 and DAVIS testset. For PSNR: larger is better; best results are shown in blue, second best in red. For ST-RRED: smaller is better; best results are shown bold.

Set8	VNLB	V-BM4D	NV	VNLnet	DVDnet	FastDVDnet
$\sigma = 10$	37.26 / 2.86	36.05 / 3.87	35.67 / 3.42	37.10 / 3.43	36.08 / 4.16	36.44 / 3.00
$\sigma = 20$	33.72 / 6.28	32.19 / 9.89	31.69 / 12.48	33.88 / 6.88	33.49 / 7.54	33.43 / 6.65
$\sigma = 30$	31.74 / 11.53	30.00 / 19.58	28.84 / 33.19	-	31.79 / 12.61	31.68 / 11.85
$\sigma = 40$	30.39 / 18.57	28.48 / 32.82	26.36 / 47.09	30.55 / 19.71	30.55 / 19.05	30.46 / 18.45
$\sigma = 50$	29.24 / 27.39	27.33 / 49.20	25.46 / 57.44	29.47 / 29.78	29.56 / 27.97	29.53 / 26.75

DAVIS	VNLB	V-BM4D	VNLnet	DVDnet	FastDVDnet
$\sigma = 10$	38.85 / 3.22	37.58 / 4.26	35.83 / 2.81	38.13 / 4.28	38.71 / 3.49
$\sigma = 20$	35.68 / 6.77	33.88 / 11.02	34.49 / 6.11	35.70 / 7.54	35.77 / 7.46
$\sigma = 30$	33.73 / 12.08	31.65 / 21.91	-	34.08 / 12.19	34.04 / 13.08
$\sigma = 40$	32.32 / 19.33	30.05 / 36.60	32.32 / 18.63	32.86 / 18.16	32.82 / 20.39
$\sigma = 50$	31.13 / 28.21	28.80 / 54.82	31.43 / 28.67	31.85 / 25.63	31.86 / 28.89

Table 2. Comparison with ViDeNN for clipped AWGN. See the text for more details. For PSNR: larger is better; best results are shown in bold.

DAVIS	ViDeNN	FastDVDnet.clipped
$\sigma = 10$	37.13	38.45
$\sigma = 30$	32.24	33.52
$\sigma = 50$	29.77	31.23

6. Running times

Our method achieves fast inference times, thanks to its design characteristics and simple architecture. Our algorithm takes only 100ms to denoise a 960×540 color frame, which is more than 3 orders of magnitude faster than V-BM4D and VNLB, and more than an order of magnitude faster than other CNN algorithms which run on GPU, DVDnet and VNLnet. The algorithms were tested on a server with a Titan Xp NVIDIA GPU card. Figure 6 compares the running times of different state-of-the-art algorithms.

7. Conclusion

In this paper, we presented FastDVDnet, a state-of-the-art video denoising algorithm. Denoising results of FastDVDnet feature remarkable temporal coherence, very low flickering, and excellent detail preservation. This level of performance is achieved even without a flow estimation step. The algorithm runs between one and three orders of magnitude faster than other state-of-the-art competitors. In this sense, our approach proposes a major step forward towards high quality real-time deep video noise reduction. Although the results presented in this paper hold for Gaussian noise, our method could be extended to denoise other types of noise.

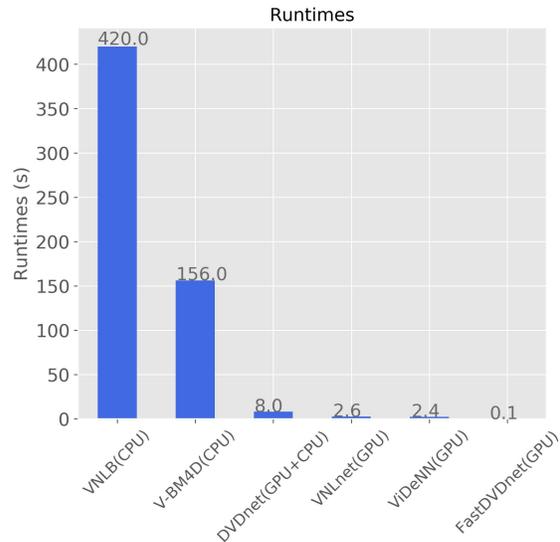


Figure 6. Comparison of running times. Time to denoise a color frame of resolution 960×540 . Note: values displayed for VNLB do not include the time required to estimate motion.

Acknowledgments

Julie Delon would like to thank the support of NVIDIA Corporation for providing us with the Titan Xp GPU used in this research. We thank Anna Murray and José Lezama for their valuable contribution. This work has been partially funded by the French National Research and Technology Agency (ANRT) and GoPro Technology France.

References

- [1] ABSOft. Neat Video. <https://www.neatvideo.com>, 1999–2019. 5
- [2] Pablo Arias and Jean-Michel Morel. Video denoising via empirical bayesian estimation of space-time patches. *Journal of Mathematical Imaging and Vision*, 60(1):70–93, Jan 2018. 2
- [3] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T. Barron. Unprocessing Images for Learned Raw Denoising. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [4] Antoni Buades, Jose-Luis Lisani, and Marko Miladinovic. Patch-based video denoising with optical flow estimation. *IEEE Transactions on Image Processing*, 25(6):2573–2586, Jun 2016. 2
- [5] H.C. Burger, C.J. Schuler, and S. Harmeling. Image denoising: Can plain neural networks compete with BM3D? In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2392–2399, 2012. 1
- [6] Jose Caballero, Christian Ledig, Andrew Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4778–4787, 2017. 2
- [7] Xinyuan Chen, Li Song, and Xiaokang Yang. Deep rnns for video denoising. volume 9971 of *SPIE Proceedings*, page 99711T. SPIE, Sep 2016. 2
- [8] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1256–1272, Jun 2017. 1
- [9] Michele Claus and Jan van Gemert. Videnn: Deep blind video denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 2
- [10] K Dabov, A Foi, and V Katkovnik. Image denoising by sparse 3D transformation-domain collaborative filtering. *IEEE Transactions on Image Processing (TIP)*, 16(8):1–16, 2007. 1
- [11] Axel Davy, Thibaud Ehret, Gabriele Facciolo, Jean-Michel Morel, and Pablo Arias. Non-local video denoising by cnn. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [12] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. pages 2758–2766. IEEE, Dec 2015. 2
- [13] Thibaud Ehret, Axel Davy, Jean-Michel Morel, Gabriele Facciolo, and Pablo Arias. Model-blind video denoising via frame-to-frame training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11369–11378, 2019. 2
- [14] Michaël Gharbi, Gaurav Chaurasia, Sylvain Paris, and Frédéric Durand. Deep joint demosaicking and denoising. *ACM Transactions on Graphics*, 35(6):1–12, Nov 2016. 3
- [15] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward Convolutional Blind Denoising of Real Photographs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1712–1722, jul 2019. 3
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 2, 4
- [17] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning (ICML)*, pages 448–456. JMLR.org, 2015. 4
- [18] Anna Khoreva, Anna Rohrbach, and Bernt Schiele. Video object segmentation with language referring expressions. In *ACCV*, 2018. 5
- [19] D.P. Kingma and J.L. Ba. ADAM: a Method for Stochastic Optimization. *Proc. ICLR*, pages 1–15, 2015. 5
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems (NIPS)*, pages 1–9, 2012. 4
- [21] M. Lebrun, A. Buades, and J. M. Morel. A nonlocal bayesian image denoising algorithm. *SIAM Journal on Imaging Sciences*, 6(3):1665–1688, Jan 2013. 1
- [22] Ce Liu and William Freeman. A high-quality video denoising algorithm based on reliable motion estimation. In *European Conference on Computer Vision (ECCV)*, pages 706–719. Springer, 2015. 2
- [23] Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, and Thomas S Huang. Non-local recurrent network for image restoration. In *Advances in Neural Information Processing Systems*, pages 1680–1689, 2018. 2
- [24] Pengju Liu, Hongzhi Zhang, Kai Zhang, Liang Lin, and Wangmeng Zuo. Multi-level wavelet-CNN for image restoration. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2018. 1
- [25] Matteo Maggioni, Giacomo Boracchi, Alessandro Foi, and Karen Egiazarian. Video denoising, deblocking, and enhancement through separable 4-d nonlocal spatiotemporal transforms. *IEEE Transactions on Image Processing*, 21(9):3952–3966, Sep 2012. 2
- [26] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML*, pages 1310–1318, 2013. 2
- [27] Adam Paszke, Gregory Chanan, Zeming Lin, Sam Gross, Edward Yang, Luca Antiga, and Zachary Devito. Automatic differentiation in PyTorch. *Advances in Neural Information Processing Systems 30*, pages 1–4, 2017. 5
- [28] Tobias Plötz and Stefan Roth. Neural nearest neighbors networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1087–1098, 2018. 2
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. *U-Net: Convolutional Networks for Biomedical Image Segmentation*, volume 9351 of *Lecture Notes in Computer Science*,

- chapter chapter 28, pages 234–241. Springer International Publishing, 2015. 2
- [30] V. Santhanam, V.I. Morariu, and L.S. Davis. Generalized Deep Image to Image Regression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [31] U. Schmidt and S. Roth. Shrinkage fields for effective image restoration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, number 8, pages 2774–2781, 2014. 1
- [32] K. Seshadrinathan and A.C. Bovik. Motion tuned spatio-temporal quality assessment of natural videos. *IEEE Transactions on Image Processing*, 19(2):335–350, Feb 2010. 2
- [33] Tamara Seybold. *Noise Characteristics and Noise Perception*, pages 235–265. Springer International Publishing, 2018. 2
- [34] Wenzhe Shi, Jose Caballero, Ferenc Huszar, Johannes Totz, Andrew P. Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. pages 1874–1883. IEEE, Jun 2016. 4
- [35] Rajiv Soundararajan and Alan C. Bovik. Video quality assessment by reduced reference spatio-temporal entropic differencing. *IEEE Transactions on Circuits and Systems for Video Technology*, 2013. 6
- [36] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, dec 2015. 5
- [37] Matias Tassano, Julie Delon, and Thomas Veit. An analysis and implementation of the ffdnet image denoising method. *Image Processing On Line*, 9:1–25, Jan 2019. 3, 4, 5
- [38] Matias Tassano, Julie Delon, and Thomas Veit. DVDnet: A fast network for deep video denoising. In *IEEE International Conference on Image Processing*, Sep 2019. 1, 2, 5
- [39] Thijs Vogels, Fabrice Rousselle, Brian Mcwilliams, Gerhards Röthlin, Alex Harvill, David Adler, Mark Meyer, and Jan Novák. Denoising with kernel prediction and asymmetric loss functions. *ACM Transactions on Graphics*, 37(4):1–15, Jul 2018. 2
- [40] Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui, and Cordelia Schmid. DeepFlow: Large displacement optical flow with deep matching. In *IEEE International Conference on Computer Vision (ICCV)*, Sydney, Australia, Dec. 2013. 5
- [41] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4148–4158, 2017. 5
- [42] Shangzhe Wu, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Deep High Dynamic Range Imaging with Large Foreground Motions. In *European Conference on Computer Vision (ECCV)*, pages 117–132, 2018. 2
- [43] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, Jul 2017. 1, 5
- [44] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, Sep 2018. 1, 2, 3, 5