

Restore from Restored: Video Restoration with Pseudo Clean Video

Seunghwan Lee¹, Donghyeon Cho², Jiwon Kim³, Tae Hyun Kim^{*1}

seunghwanlee@hanyang.ac.kr, cdh12242@gmail.com, ai.kim@sk.com, taehyunkim@hanyang.ac.kr

¹Dept. of Computer Science, Hanyang University, Seoul, Korea

²Dept. of Electronic Engineering, Chungnam National University, Daejeon, Korea

³SKT Vision AI Labs/T-Brain X, Seoul, Korea

Abstract

In this study, we propose a self-supervised video denoising method called “restore-from-restored.” This method fine-tunes a pre-trained network by using a pseudo clean video during the test phase. The pseudo clean video is obtained by applying a noisy video to the baseline network. By adopting a fully convolutional neural network (FCN) as the baseline, we can improve video denoising performance without accurate optical flow estimation and registration steps, in contrast to many conventional video restoration methods, due to the translation equivariant property of the FCN. Specifically, the proposed method can take advantage of plentiful similar patches existing across multiple consecutive frames (i.e., patch-recurrence); these patches can boost the performance of the baseline network by a large margin. We analyze the restoration performance of the fine-tuned video denoising networks with the proposed self-supervision-based learning algorithm, and demonstrate that the FCN can utilize recurring patches without requiring accurate registration among adjacent frames. In our experiments, we apply the proposed method to state-of-the-art denoisers and show that our fine-tuned networks achieve a considerable improvement in denoising performance.

1. Introduction

Video restoration, which aims to recover the high-quality video frames from the low-quality video, is one of the oldest research fields in video processing. Video denoising, which removes noise in the video frames, has been investigated considerably. However, estimating a clean image from a corrupted frame is a well-known inverse problem. To solve such an ill-posed problem, various types of approaches, in-

cluding prior model, likelihood model, optimization, and deep learning techniques, have been introduced.

A common natural image property used for image restoration is patch-recurrence in which similar patches exist within a single image. Particularly, patch-recurrence has been considerably studied in single-image super-resolution (SR) methods [9, 11, 12]. Although these patches can be deformed by camera and/or object motion in a video, patch-recurrence over neighboring video frames is much richer than that of a single image and can further improve the quality of the restored frames [19, 18]. Moreover, rich patch-recurrence information can greatly help in the fine-tuning of video restoration networks during the test stage, without using ground-truth clean images.

Lehtinen *et al.* [16] proposed single-image denoising method (noise-to-noise), which allows the training of the restoration network without ground-truth clean images. Ehret *et al.* [7] proposed a frame-to-frame training technique, which extends the noise-to-noise training algorithm for video restoration; the frame-to-frame training algorithm can also perform fine-tuning without using the ground-truth clean video by aligning noisy patches among consecutive frames using optical flow. However, estimating accurate optical flow under large displacements, occlusion, and severe degradation (*e.g.*, noise and blur) is a challenging task. Thus, in this work, we propose a new training algorithm called “restore-from-restored,” which allows fine-tuning of pre-trained networks without using the ground-truth clean video and accurate optical flow for registration.

Our proposed method updates the parameters of pre-trained networks using pseudo clean images, which are outputs of the pre-trained baseline networks from noisy input frames. Our algorithm is simple yet effective for removing noise in video frames and works particularly well with the existence of numerous recurring patches. That is, we generate pairs of training images, which are composed of the pseudo clean video and its noisy versions, to fine-tune the network. In practice, pixel locations of the same patches in different video frames vary due to motions, but with the aid

* Corresponding author.

Code is available at <https://github.com/shlee0/RFR-video-denoising>.

of the translation equivariant property of a fully convolutional network (FCN), our algorithm can update the network parameters without using optical flow to align the translated patches only if they are fully convolutional. We demonstrate the superiority of the proposed algorithm by applying it to state-of-the-art video denoising networks and providing improved denoising results. The contributions of this study are summarized as follows:

- We propose a novel self-supervised training algorithm to fine-tune fully pre-trained networks without using the clean ground-truth video.
- We explain why and how the proposed training scheme works with the patch-recurrence property.
- The proposed method can be easily integrated with state-of-the-art denoising networks and yields state-of-the-art denoising results on the benchmark datasets including not only synthetic but also real noise.

2. Related Works

In this section, we provide a brief overview of recent works that are related to the proposed restoration algorithm, in terms of training with and without using ground-truth clean data.

Training with ground-truth clean data. When a set of high-quality images is available, we can generate synthetic degraded images, and train deep neural networks with these images and restore them to their original high-quality state.

In the case of image denoising, Xie *et al.* [26] applied deep neural networks to model the mapping of clean images from noisy input images. They generated pairs of noisy and clean images to train the neural networks. Since then, numerous studies on the image denoising task have been conducted using deep CNN with train pairs of clean and synthetically noisy images [28, 29, 30, 32, 15, 17, 31, 10, 1]. Several studies have adopted residual learning schemes to allow deep neural networks and extend the receptive field [28, 32]. Moreover, to incorporate long-range dependencies among pixels, several studies utilized non-local networks [31, 17]. Recent efforts have attempted to deal with unknown noise in real photographs (blind restoration). Guo *et al.* [10] proposed a two-stage method that consists of noise estimation and non-blind denoising steps. Gao and Grauman [8] proposed an on-demand learning method to handle various corruption levels for each restoration task including denoising, inpainting, and deblurring. These research trends have also been applied to video restoration problems. For instance, Davy *et al.* [6] not only incorporated non-local information with a non-local patch search module but also adopted a residual learning scheme for video denoising. FastDVDnet [25] proposed a cascaded two-step architecture with a modified multi-scale U-Net.

FastDVDnet exhibits fast runtimes by avoiding costly explicit motion compensation and handling motion implicitly due to the attributes of its dedicated architecture. Yue *et al.* [27] proposed RViDeNet to restore real noisy video frames in raw image spaces and achieved state-of-the-art denoising performance. RViDeNet adopts deformable convolution [5] to align consecutive frames and utilizes spatio-temporal fusion stages to reconstruct the result. Moreover, Yue *et al.* provided a new video denoising dataset, namely, Captured Raw Video Dataset (CRVD).

Regardless of the techniques used in previous works (*e.g.*, residual learning, non-local network, and model-blind approaches), these studies require the generation of synthetic images for network training. Thus, these supervised approaches hardly deal with datasets where clean images can be rarely obtained (*e.g.*, medical imaging system).

Training without ground-truth clean data. Several attempts have been made recently to learn restoration networks without using ground-truth clean data. Lehtinen *et al.* [16] trained a network with pairs of noisy patches under the assumption that the average of many differently corrupted pixels is close to clean data. Then, Krull *et al.* [14] and Baston and Royer [2] introduced self-supervised single image denoising methods without relying on clean data. Ehret *et al.* [7] introduced a frame-to-frame training method to learn video restoration networks without clean images by extending the strategy proposed in [16] to videos. For removing noise in a certain patch, their method searches corresponding patches among adjacent frames by using optical flow and then warps the patches to create pairs of aligned noisy patches for network training. Frame-to-frame training enables the exploitation of patch-recurrence property within input video frames and fine-tune the pre-trained networks using test inputs. This approach can boost the performance of the existing pre-trained networks because networks can be further optimized without the ground-truth targets at the test stage.

However, one disadvantage of this method [7] is that accurate optical flow, which is difficult to estimate under large displacements, occlusions and serious damages, is required to fine-tune network. In this study, we overcome the limitations of [7] by using a new training scheme called “restore-from-restored.” Technically, pseudo clean images are generated from a pre-trained video restoration network and then used as train targets for fine-tuning during the test phase. It produces a synergy effect with the patch-recurrence property that appears repeatedly over consecutive video frames. In the following sections, we provide detailed analysis on the proposed method and show the proposed method can boost the performance of the fully pre-trained denoising networks with the help of the patch-recurrence property.

To the best of our knowledge, the proposed method is the first neural approach to boost the performance of pre-trained

convolutional video restoration networks without using accurate registration or non-local operation while using recurring patches in the test-phase.

3. Self-Supervised Video Restoration

Patch-recurrence within the same scale is rich in natural images [24, 22] and becomes more redundant when multiple neighboring video frames are available [23]. To utilize this space-time recurring information among given video frames, conventional restoration methods require accurate correspondences between adjacent frames and thus need to compute the optical flow to align the neighboring frames to the reference frames [3, 21, 13].

In this work, we present a novel yet simple training algorithm (test-time fine-tuning) that can be applied to video restoration networks. Our fine-tuning algorithm is based on self-supervision and does not require ground-truth clean images. Moreover, the proposed algorithm allows restoration networks to exploit patch-recurrence without accurate optical flow estimation and registration steps while improving performance by a large margin. Many convolutional video restoration networks, including state-of-the-art methods, can be easily fine-tuned using our self-supervised training algorithm without changing their original network architecture if they are fully convolutional.

3.1. Restore-from-restored

In this section, we explain how we can fine-tune and improve the performance of the pre-trained video restoration networks without using the clean video frames during the test stage.

In general, conventional video restoration networks are trained with labeled ground-truth clean images; these networks learn a function f_θ , which maps a corrupted input frame \mathbf{Y} to a clean target frame \mathbf{X} , where θ denotes the function parameters. Specifically, the network parameter θ is trained by minimizing the loss function L between the network outcome and the training target as

$$Loss(\theta) = L(f_\theta(\mathbf{Y}), \mathbf{X}). \quad (1)$$

For the loss function L , common choices are L1 and L2 losses in many denoising approaches [2, 16]. Although image restoration networks trained by minimizing the distance between the network output and the training target can produce highly satisfactory results, these networks can be further upgraded by utilizing redundant spatio-temporal information (e.g., patch-recurrence) over neighboring video frames [13, 6]. However, optical flow estimation networks or non-local operation modules to exploit the recurring patches among different frames are expensive and require additional resources to extract the temporal information [3, 6, 21, 17, 31].

Therefore, we develop a simple and effective fine-tuning algorithm that can exploit patch-recurrence in space-time without explicitly searching similar patches/features through optical flow estimation or non-local operation. To achieve this goal, we assume that we have a fully pre-trained network f_{θ_0} and use initially denoised video frames $\{\tilde{\mathbf{X}}_1, \dots, \tilde{\mathbf{X}}_T\}$ as train targets for the fine-tuning, where $\tilde{\mathbf{X}}_t = f_{\theta_0}(\mathbf{Y}_t)$ and t denotes the frame index. Although these denoised images are not clean ground-truth images and may include some artifacts, they can be used as pseudo clean targets to fine-tune the network parameter in our study. Using the pseudo clean images, we can synthesize pseudo noisy images by adding random noise \mathbf{N} to the pseudo clean images, and the pairs of pseudo clean and pseudo noisy images can be used to fine-tune the denoising networks by minimizing the loss as follows:

$$Loss(\theta) = \sum_{t=1}^T L(f_\theta(\tilde{\mathbf{X}}_t + \mathbf{N}), \tilde{\mathbf{X}}_t), \quad (2)$$

Notably, the ground-truth frames and motion flows are not used during our fine-tuning process. Nevertheless, under an assumption that the distributions of the original noisy input and the corresponding pseudo noisy images are similar, we can update the network parameter θ by minimizing the proposed loss with the initially denoised frames and their synthetically corrupted counterparts if the networks are fully convolutional. Assume that we have self-similar noisy patches \mathbf{y}_a and \mathbf{y}_b where \mathbf{y}_b is a translated version of \mathbf{y}_a by a translational motion A (i.e., $A(\mathbf{y}_a) = \mathbf{y}_b$). Then, we can generate pseudo clean patches $\tilde{\mathbf{x}}_a$ and $\tilde{\mathbf{x}}_b$ with a pre-trained network, and we see that $A(\tilde{\mathbf{x}}_a) = \tilde{\mathbf{x}}_b$ when the network is an FCN, and it yields,

$$\begin{aligned} L(f_\theta(\tilde{\mathbf{x}}_a + \mathbf{N}), \tilde{\mathbf{x}}_a) &\approx L(f_\theta(A(\tilde{\mathbf{x}}_a + \mathbf{N})), A(\tilde{\mathbf{x}}_a)) \\ &\approx L(f_\theta(\tilde{\mathbf{x}}_b + \mathbf{N}), \tilde{\mathbf{x}}_b). \end{aligned} \quad (3)$$

Therefore, overall loss in (2) does not depend on locations of self-similar patches over multiple frames during the back-propagation with FCNs, and our fine-tuned network predicts the averaged version of denoised self-similar patches with L2 loss when one of the self-similar noisy patches is given as input. We call this process “restore-from-restored” training, since the proposed method is relying on the initially restored (denoised) frames. We repeat our training algorithm several times until convergence and achieve considerable improvement over the fully pre-trained initial network.

3.2. Frame-to-frame vs. Restore-from-restored

The notion of our “restore-from-restored” algorithm is based on recent noise-to-noise training mechanism by Lehtinen *et al.* [16]. Noise-to-noise demonstrated that image restoration networks can be trained without using

ground-truth clean data for certain types of noise (e.g., zero-mean noise, such as Gaussian noise and Bernoulli noise) and can be extended into the frame-to-frame approach by Ehret *et al.* [7] to process a video.

Exploiting space-time patch-recurrence. The frame-to-frame training algorithm [7] allows the network to train with self-supervision during the test-stage. Specifically, the network is fine-tuned with two aligned noisy frames by minimizing the loss as

$$Loss(\theta) = \sum_{t=1}^T L(f_{\theta}(\mathbf{Y}_t), \mathbf{Y}_{t-1}^w), \quad (4)$$

where \mathbf{Y}_{t-1}^w denotes the warped version of the noisy frame \mathbf{Y}_{t-1} and is aligned to the reference frame \mathbf{Y}_t . Thus, the calculation of optical flow for registration is necessary in (4). However, the accurate estimation of optical flow is difficult to achieve in some conditions, such as severe degradation, large displacement between frames.

By contrast, our proposed loss in (2) does not need warping and alignment. If the denoising network f_{θ} is an FCN, then our method can maintain the performance without accurate registration due to the translation equivariant nature of FCN [4]. That is, our “restore-from-restored” approach is not disturbed by the existence of large translational motions compared with optical flow-based methods.

Noise reduction. Assume that a set of perfectly aligned images $\{\mathbf{Y}_1, \dots, \mathbf{Y}_T\}$ (e.g., burst mode images from a camera on a tripod) is given, and these images are corrupted by zero-mean Gaussian random noise whose standard deviation is σ . Using the frame-to-frame training algorithm [7], the denoised frames become an averaged version of noisy inputs ($=\frac{1}{T} \sum_{t=1}^T \mathbf{Y}_t$) with the fine-tuned parameter, and the noise variance of the denoised frame is reduced to $\frac{1}{T}\sigma^2$ (refer to the Appendix in [16] for details).

By contrast, the latent frame predicted from a fine-tuned parameter by minimizing the L2 version of the proposed loss in (2) is $\frac{1}{T} \sum_{t=1}^T \tilde{\mathbf{X}}_t$. Thus, the noise variance of our denoising result becomes $\frac{1}{T}\sigma_{\theta_0}^2$, where σ_{θ_0} denotes the mean standard deviation of the remaining noise in $\tilde{\mathbf{X}}_t$. In general, as f_{θ_0} is a fully pre-trained network, the noise level of the residual noise σ_{θ_0} is much lower than the original noise level σ (i.e., $\sigma_{\theta_0} \ll \sigma$). Therefore, the noise variance of the latent frame from our “restore-from-restored” algorithm is much lower than that from the frame-to-frame training. Notably, our algorithm can achieve better results when there is a rich patch-recurrence, and our algorithm can show space-time varying denoising performance. For example, noisy regions where structures are highly repeating can be restored much better than regions with slightly repeating or unique patterns.

4. Proposed Method

Algorithm 1: Offline video denoising algorithm

Input: degraded video frames $\{\mathbf{Y}_1, \dots, \mathbf{Y}_T\}$
Output: denoised video frames $\tilde{\mathbf{X}}^K$
Require: pre-trained network f_{θ_0} , iteration number K , learning rate α

- 1 $i \leftarrow 0$
- while** $i \leq K$ **do**
- 2 **foreach** t **do**
- 3 Restore: $\tilde{\mathbf{X}}_t^i \leftarrow f_{\theta_i}(\mathbf{Y}_t)$
- 4 **end**
- 5 $\tilde{\mathbf{X}}^i \leftarrow \{\tilde{\mathbf{X}}_1^i, \dots, \tilde{\mathbf{X}}_T^i\}$
- 6 $Loss(\theta_i) = \sum_{t=1}^T L(f_{\theta_i}(\tilde{\mathbf{X}}_t^i + \mathbf{N}), \tilde{\mathbf{X}}_t^i) + \sum_{t=1}^T L(f_{\theta_i}(\tilde{\mathbf{X}}_t^0 + \mathbf{N}), \tilde{\mathbf{X}}_t^0)$
- 7 $\theta_{i+1} \leftarrow \theta_i - \alpha \nabla_{\theta_i} Loss(\theta_i)$
- 8 $i \leftarrow i+1$
- end**
- 9 **Return:** $\tilde{\mathbf{X}}^K$

Algorithm 2: Online video denoising algorithm

Input: degraded frame \mathbf{Y}_t at time-step t
Output: restored frame $\tilde{\mathbf{X}}_t$, fine-tuned network f_{θ_t}
Require: fine-tuned network $f_{\theta_{t-1}}$, restored frame at the previous time-step $\tilde{\mathbf{X}}_{t-1}$, learning rate α

- 1 $Loss(\theta_{t-1}) = L(f_{\theta_{t-1}}(\tilde{\mathbf{X}}_{t-1} + \mathbf{N}), \tilde{\mathbf{X}}_{t-1})$
- 2 $\theta_t \leftarrow \theta_{t-1} - \alpha \nabla_{\theta_{t-1}} Loss(\theta_{t-1})$
- 3 Restore: $\tilde{\mathbf{X}}_t \leftarrow f_{\theta_t}(\mathbf{Y}_t)$
- Return:** $\tilde{\mathbf{X}}_t, f_{\theta_t}$ // Will be reused in the next time step

By minimizing the proposed loss function in (2), we can restore clean images from degraded ones in offline and online manners, similar to [7]. In the offline video denoising mode, we can take advantage of all the given frames for fine-tuning the network parameter. In the online denoising mode, fine-tuning and denoising are conducted successively in a sequential fashion (frame-by-frame).

Offline denoising. The offline video denoising algorithm is elaborated in Algorithm 1. For offline fine-tuning at each train step i , we remove the noise in the input noisy images $\{\mathbf{Y}_1, \dots, \mathbf{Y}_T\}$ using denoising network f_{θ_i} and obtain denoised images $\{\tilde{\mathbf{X}}_1^i, \dots, \tilde{\mathbf{X}}_T^i\}$. Next, we use pairs of training

images $\{(\tilde{\mathbf{X}}_1 + \mathbf{N}, \tilde{\mathbf{X}}_1^i), \dots, (\tilde{\mathbf{X}}_T + \mathbf{N}, \tilde{\mathbf{X}}_T^i)\}$ to fine-tune the network parameter by minimizing the loss function in (2). To avoid over-fitting and generating over-smoothed results, we use the initially denoised frames $\{\tilde{\mathbf{X}}_1^0, \dots, \tilde{\mathbf{X}}_T^0\}$ as additional train targets during the fine-tuning stages. In our supplementary material, we demonstrate the effect of the second loss term in Algorithm 1 by comparing results with and without using the initial frames during offline fine-tuning when K is large.

Online denoising. We can also adapt the network parameter in a sequential manner. Our online video denoising algorithm is given in Algorithm 2. In contrast to the offline denoising mode, our online denoising algorithm uses only a previous frame for the update at each time step. To do so, we slightly modify the proposed loss function in (2) for the time step t , to take a single pair of images $(\tilde{\mathbf{X}}_{t-1} + \mathbf{N}, \tilde{\mathbf{X}}_{t-1})$ which is previous pseudo clean and noisy frames. In Algorithm 2, temporally varying parameter θ_t denotes the fine-tuned network parameter at time t ; thus, each frame is denoised with a different network parameter.

5. Experiments

In our experiments, we apply our offline and online denoising algorithms to conventional denoising networks, including state-of-the-art video denoising methods, and evaluate the denoising performance quantitatively and qualitatively. Please refer to our supplementary material for additional experimental results, including video clips.

5.1. Implementation Details

We use officially available fully pre-trained network parameters for our baseline denoisers. During the fine-tuning period, we minimize the L2 loss for the proposed offline and online denoising algorithms. We also use Adam for the updates, with a learning rate of $1e-5$. We measure the performance in terms of PSNR and SSIM on RGB color channels for the objective evaluation.

First, for Gaussian noise removal, we use VNLnet [6] and FastDVDnet [25]. VNLnet is a state-of-the-art video denoising method and shows the best denoising performance. Although VNLnet is integrated with the non-local operation module, we show that our algorithm can be easily applied to networks with the non-local module and can further improve the performance of the baseline networks. For the evaluation, we use 7 video clips consisting of 100 sequences each on the Derf database¹ and additional 7 video clips in the DAVIS video segmentation challenge dataset [20]. We use down-scaled video frames (960×540), and generate noisy input videos by adding Gaussian random noise with different noise levels ($\sigma=15, 25, 40$).

¹<https://media.xiph.org/video/derf/>

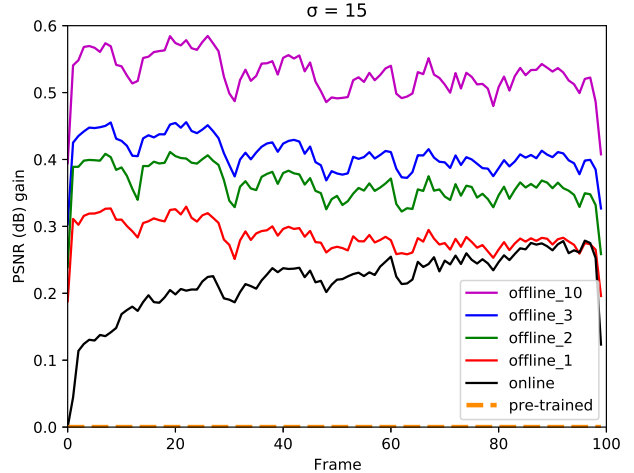


Figure 1: Performance gains from our online and offline denoising algorithms. Baseline network is FastDVDnet [25]. Differences of PSNR values before and after fine-tuning are measured on Derf datasets on $\sigma = 15$. Number i in “offline- i ” denotes the number of steps (*i.e.*, K). Note that the rapid rise and drop of performance at the very first and last time step are due to a usability of adjacent frames.

Next, for real noise removal, we use RViDeNet [27] and evaluate on the CRVD dataset [27]. We improve the baseline performance with the proposed fine-tuning algorithm in the sRGB space. To handle real noise, we synthesize random noise \mathbf{N} with the noise model used in the pre-training process of RViDeNet. Refer to our supplementary material for detailed settings.

5.2. Denoising Performance

Quantitative results. First, we fine-tune FastDVDnet [25] on the Derf dataset with the proposed offline and online denoising algorithms to evaluate the Gaussian denoising performance; the performance gains are depicted in Fig. 1. We achieve consistently improved denoising performance over the baseline as the number of iterations (*i.e.*, K) increases with the offline denoising algorithm. Meanwhile, we obtain improved results as the frame number increases with our online denoising algorithm because the proposed online method enables sequential parameter update. In Tab. 1 and Tab. 2, we provide the PSNR and SSIM values of the denoising results from our algorithms on the Derf and DAVIS testsets. These results show that the online and offline denoising algorithms can produce steadily better results than the baseline models, and the offline denoising algorithm with 10 updates (*i.e.*, $K = 10$) achieves the best performance. In Tab. 3, we add quantitative denoising results on the real noise dataset, CRVD [27]. As CRVD includes only short sequences ($T = 7$), we only evaluate the performance of the offline restoration algorithm; the results




Figure 2: Denoising results with FastDVDnet [25] on Derf testsets corrupted by Gaussian noise ($\sigma = 40$). Visual comparisons with our online and offline ($K = 10$ in Algorithm 1) update procedures.

show that the proposed learning algorithm can effectively remove real noise remarkably and enhance the denoising results.

Visual results. In Fig. 2, we provide qualitative comparison results. The input images are corrupted with high-level Gaussian noise ($\sigma = 40$), and FastDVDnet is fine-tuned by our offline and online restoration algorithms. We also provide real noise denoising results with fine-tuned RViDeNet in Fig. 3. Our methods can produce much better visual results and restore tiny details compared with the initially pre-trained network.

Run-time. We report the run-time for a single update step on the NVIDIA Tesla V100 graphics unit. FastDVDnet and VNLnet take approximately 0.16 and 0.4 second to handle a 960×540 input frame with our online denoising algorithm

in Algorithm 2; RViDeNet takes approximately 3.0 second to process the 256×256 input frame with our offline denoising method in Algorithm 1.

5.3. Comparison with the frame-to-frame

We compare our online denoising method with the frame-to-frame training algorithm [7], and their official code and parameters are used for the fine-tuning. As the baseline, a fully pre-trained DnCNN [28], which is trained with the Gaussian random noise ($\sigma = 25$) on the large external dataset is used, and then fine-tuned on the Derf dataset using frame-to-frame and our algorithms. For fair comparison, the same hyper-parameters (e.g., optimizer, learning rate, number of updates for each frame) are used to run the both frame-to-frame and our restoration algorithms; the

Method	σ	crowd	park joy	pedestrian	station	sunflower	touchdown	tractor	Average
FastDVDnet	15	31.17/0.9208	30.46/0.8959	38.00/0.9522	36.65/0.9323	37.95/0.9525	35.98/0.9109	34.03/0.9263	34.89/0.9273
		31.32/0.9233	30.72/0.9058	38.20/0.9537	36.83/0.9338	38.38/0.9554	36.06/0.9123	34.26/0.9283	35.11/0.9304
		31.57/0.9263	30.95/0.9108	38.44/0.9549	37.10/0.9363	39.03/0.9596	36.23/0.9138	34.63/0.9320	35.42/0.9334
	25	29.01/0.8813	28.34/0.8373	35.76/0.9328	34.96/0.9050	34.96/0.9306	33.97/0.8667	31.99/0.8946	32.71/0.8926
		29.32/0.8892	28.79/0.8610	36.27/0.9369	35.12/0.9076	36.14/0.9389	34.09/0.8678	32.32/0.8990	33.15/0.9000
		29.59/0.8943	29.06/0.8719	36.62/0.9393	35.38/0.9107	37.00/0.9455	34.33/0.8703	32.78/0.9051	33.54/0.9053
	40	26.43/0.8159	25.80/0.7441	32.44/0.9003	32.89/0.8634	30.66/0.8884	31.78/0.8022	29.69/0.8490	29.95/0.8376
		27.21/0.8375	26.82/0.7927	33.93/0.9118	33.34/0.8709	33.39/0.9129	32.18/0.8095	30.32/0.8594	31.03/0.8564
		27.67/0.8509	27.27/0.8192	34.74/0.9190	33.42/0.8717	34.93/0.9263	32.56/0.8140	30.98/0.8697	31.69/0.8679
VNLnet	15	32.68/0.9373	32.19/0.9200	38.85/0.9567	38.51/0.9501	39.58/0.9628	37.37/0.9347	35.12/0.9378	36.33/0.9428
		32.83/0.9386	32.47/0.9279	38.94/0.9570	38.60/0.9508	39.87/0.9640	37.41/0.9348	35.30/0.9393	36.49/0.9446
		33.00/0.9401	32.72/0.9328	39.05/0.9576	38.72/0.9515	40.22/0.9661	37.52/0.9350	35.54/0.9416	36.68/0.9464
	25	30.07/0.9013	29.48/0.8651	36.15/0.9353	36.57/0.9263	36.06/0.9430	35.21/0.8976	32.83/0.9072	33.77/0.9108
		30.34/0.9063	30.00/0.8849	36.58/0.9375	36.69/0.9280	37.37/0.9494	35.32/0.8991	33.11/0.9106	34.20/0.9165
		30.58/0.9105	30.33/0.8970	36.93/0.9397	36.83/0.9293	38.02/0.9539	35.49/0.8994	33.41/0.9142	34.51/0.9206
	40	27.09/0.8366	26.51/0.7724	32.48/0.8992	33.91/0.8817	31.01/0.8975	32.33/0.8122	30.09/0.8569	30.49/0.8509
		27.85/0.8536	27.61/0.8153	33.78/0.9078	34.36/0.8879	34.09/0.9231	32.71/0.8199	30.71/0.8672	31.57/0.8678
		28.35/0.8648	28.18/0.8425	34.71/0.9154	34.64/0.8910	35.62/0.9343	33.06/0.8220	31.33/0.8760	32.27/0.8780

Table 1: Denoising results with FastDVDnet [25] and VNLnet [6] on the Derf testset with different Gaussian noise levels ($\sigma = 15, 25, 40$). For each network architecture and each noise level, the PSNR and SSIM results of the baseline, online learning (Algorithm 2) and offline learning (Algorithm 1) are listed in each box from top to bottom. The best average results are written in bold letters.

Method	σ	chameleon	giant-slamom	girl-dog	hoverboard	monkeys-trees	salsa	subway	Average
FastDVDnet	15	36.65/0.9697	40.79/0.9685	34.25/0.9183	39.55/0.9613	31.59/0.9567	33.41/0.9440	37.56/0.9246	36.26/0.9490
		36.73/0.9702	41.08/0.9708	34.29/0.9202	39.63/0.9617	31.63/0.9569	33.62/0.9519	39.20/0.9601	36.60/0.9560
		36.92/0.9711	41.29/0.9717	34.35/0.9214	39.71/0.9616	31.75/0.9578	33.84/0.9626	40.03/0.9741	36.84/0.9601
	25	33.98/0.9534	38.67/0.9550	31.61/0.8590	37.36/0.9486	28.71/0.9161	30.17/0.9002	33.48/0.8737	33.43/0.9151
		34.22/0.9550	38.94/0.9577	31.73/0.8639	37.54/0.9494	28.74/0.9164	30.59/0.9143	36.32/0.9334	34.01/0.9272
		34.53/0.9567	39.16/0.9588	31.88/0.8679	37.82/0.9504	28.84/0.9184	30.93/0.9328	37.95/0.9637	34.44/0.9355
	40	30.97/0.9263	36.53/0.9384	29.01/0.7785	34.82/0.9315	26.29/0.8506	26.76/0.8162	28.83/0.8070	30.46/0.8640
		31.58/0.9311	36.84/0.9410	29.44/0.7907	35.32/0.9338	26.32/0.8529	27.65/0.8500	32.88/0.8885	31.43/0.8840
		32.30/0.9364	37.16/0.9428	29.84/0.7985	35.92/0.9361	26.38/0.8546	28.26/0.8819	35.85/0.9505	32.25/0.9001
VNLnet	15	37.30/0.9724	42.31/0.9751	35.68/0.9407	39.83/0.9626	34.87/0.9792	34.04/0.9461	37.60/0.9188	37.37/0.9564
		37.37/0.9725	42.32/0.9751	35.73/0.9418	39.82/0.9621	34.94/0.9794	34.21/0.9517	39.42/0.9522	37.69/0.9621
		37.50/0.9734	42.40/0.9753	35.81/0.9428	39.88/0.9621	35.02/0.9798	34.39/0.9616	40.41/0.9720	37.97/0.9667
	25	34.46/0.9567	39.75/0.9604	32.76/0.8887	37.58/0.9501	31.96/0.9591	30.60/0.9015	32.88/0.8642	34.28/0.9258
		34.66/0.9576	39.81/0.9606	32.93/0.8947	37.66/0.9497	32.02/0.9596	31.03/0.9134	35.66/0.9085	34.83/0.9349
		34.98/0.9595	39.91/0.9610	33.09/0.8973	37.89/0.9500	32.09/0.9598	31.40/0.9314	37.98/0.9598	35.33/0.9456
	40	31.26/0.9302	37.23/0.9425	29.71/0.8061	34.96/0.9335	29.13/0.9208	27.05/0.8198	28.55/0.8052	31.13/0.8797
		31.81/0.9338	37.41/0.9431	30.24/0.8224	35.31/0.9343	29.23/0.9236	27.99/0.8510	32.05/0.8639	32.01/0.8960
		32.57/0.9391	37.65/0.9438	30.71/0.8306	35.93/0.9366	29.30/0.9242	28.69/0.8798	35.65/0.9433	32.93/0.9139

Table 2: Denoising results with FastDVDnet [25] and VNLnet [6] on the DAVIS testset with different Gaussian noise levels ($\sigma = 15, 25, 40$). For each network architecture and each noise level, the PSNR and SSIM results of the baseline, online learning (Algorithm 2) and offline learning (Algorithm 1) are listed in each box from top to bottom. The best average results are written in bold letters.

comparison results are provided in Fig. 4.

First, we show that the frame-to-frame algorithm cannot outperform the baseline network when the noise distribution of the test video is identical to that in the large external train set (*i.e.*, $\sigma = 25$). By contrast, our online restoration algorithm can still elevate the denoising performance compared with the fully pre-trained baseline by adapting the network parameter to the specific input video (Fig. 4 (a)). Next, compared with the baseline, the frame-to-frame algorithm and our approach can improve denoising quality by a large mar-

gin when the noise distribution of the input video (*i.e.*, $\sigma = 40$) is different from that of the training dataset (Fig. 4 (b)). However, when the input video includes large motion displacement, the frame-to-frame algorithm fails in estimating accurate optical flow and thus can show worse performance than the baseline; by contrast, our algorithm predicts consistently better results because ours does not rely on the optical flow (solid blue lines in Fig. 4 (c)).

In the comparison with the frame-to-frame method, our algorithm uses additional information regarding the noise




Figure 3: Denoising results with RViDeNet on the CRVD [27] dataset which includes real noise. We use $K = 10$, for our offline denoising algorithm.

	Noisy	RViDeNet	Ours
PSNR	31.79	39.95	40.13
SSIM	0.7517	0.9792	0.9795

Table 3: Denoising results with RViDeNet on the CRVD [27] dataset with real noise. The proposed offline learning method (Algorithm 1) show the quantitatively better results than the baseline (RViDeNet) and the results are written in bold letters.

distribution of the test input but shows considerably better denoising results. By contrast, the frame-to-frame algorithm requires additional resources and longer run-time to compute optical flow among video frames.

6. Conclusion

In this work, we present a new training algorithm for video denoising; this algorithm is straightforward and easy to train and produces state-of-the-art denoising results. Our training approach is based on the self-supervision and thus allows the network to adapt its pre-trained parameter for the given specific input video without using ground-truth clean frames. As we use the restored version of the input noisy frames rendered by the pre-trained denoiser as our fine-tuning target (pseudo clean images), we call the proposed algorithm “restore-from-restored.” Moreover, in contrast to conventional video restoration approaches, we restore the clean images without using accurate optical flow. We describe how the proposed training algorithm can exploit recurring patches among input video frames and improve the denoising performance. We also demonstrate the superiority of the proposed algorithm and show considerable improvements on the various benchmark datasets.

Acknowledgement

This work was supported by the research fund of the National Research Foundation of Korea (NRF)

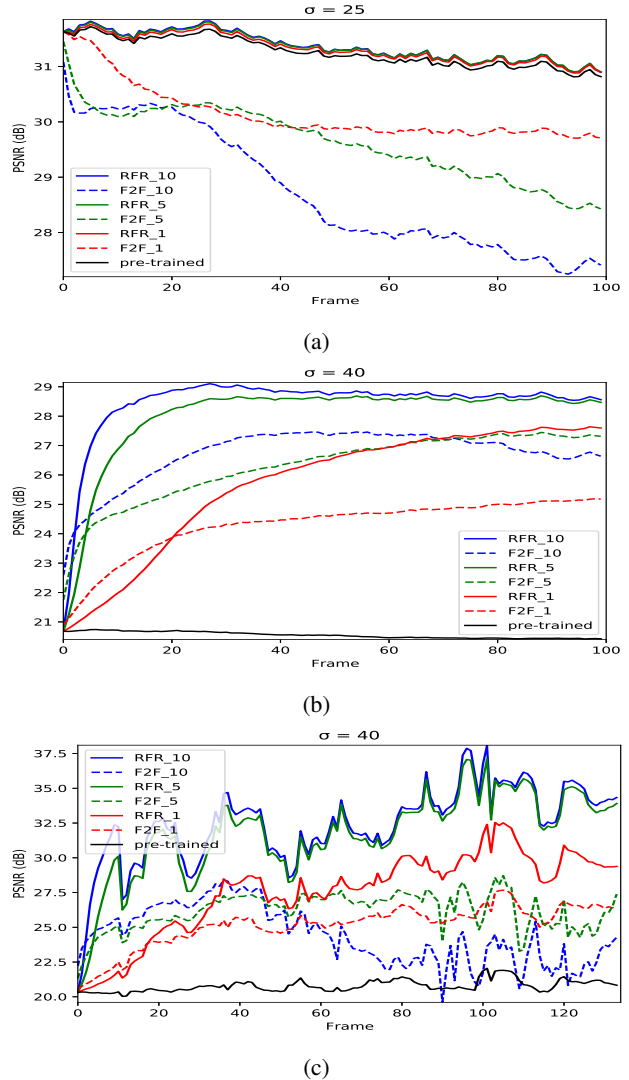


Figure 4: Comparisons of restore-from-restored (RFR) and frame-to-frame (F2F) online fine-tuning methods. The number i in “F2F $_i$ ” or “RFR $_i$ ” indicates the number of updates for each frame [7]. (a) Denoising results when the test input noise level is higher ($\sigma = 40$) than that in training ($\sigma = 25$). (b) Denoising results when noise level during the test and training stages is identical. (c) Denoising results on a video with large motions when noise level during the test and training stages is identical.

grant funded by the Korea government (MSIT) (NRF-2019R1A4A1029800), SKT Vision AI Labs/T-Brain X, Samsung Research Funding Center of Samsung Electronics under Project Number SRFCIT1901-06, and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2020-0-01373, Artificial Intelligence Graduate School Program (Hanyang University)).

References

- [1] Saeed Anwar and Nick Barnes. Real image denoising with feature attention. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [2] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *International Conference on Machine Learning (ICML)*, 2019.
- [3] 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 (CVPR)*, 2017.
- [4] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International Conference on Machine Learning (ICML)*, 2016.
- [5] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [6] Axel Davy, Thibaud Ehret, Jean-Michel Morel, Pablo Arias, and Gabriele Facciolo. Non-local video denoising by cnn. In *IEEE International Conference on Image Processing (ICIP)*, 2019.
- [7] 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 (CVPR)*, 2019.
- [8] Ruohan Gao and Kristen Grauman. On-demand learning for deep image restoration. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [9] Daniel Glasner, Shai Bagon, and Michal Irani. Super-resolution from a single image. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2009.
- [10] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [12] Jun-Jie Huang, Tianrui Liu, Pier Luigi Dragotti, and Tania Stathaki. Srhrf+: Self-example enhanced single image super-resolution using hierarchical random forests. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 71–79, 2017.
- [13] Tae Hyun Kim, Mehdi SM Sajjadi, Michael Hirsch, and Bernhard Schölkopf. Spatio-temporal transformer network for video restoration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [14] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] Stamatios Lefkimmiatis. Non-local color image denoising with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5882–5891, 2016.
- [16] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In *International Conference on Machine Learning (ICML)*, volume 80, 2018.
- [17] 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 (NIPS)*, pages 1680–1689, 2018.
- [18] 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, 2012.
- [19] Mona Mahmoudi and Guillermo Sapiro. Fast image and video denoising via nonlocal means of similar neighborhoods. *IEEE signal processing letters*, 12(12):839–842, 2005.
- [20] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017.
- [21] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6626–6634, 2018.
- [22] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [23] Oded Shahar, Alon Faktor, and Michal Irani. Space-time super-resolution from a single video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [24] Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [25] Matias Tassano, Julie Delon, and Thomas Veit. Fastdvdnet: Towards real-time deep video denoising without flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [26] Junyuan Xie, Linli Xu, and Enhong Chen. Image denoising and inpainting with deep neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 341–349, 2012.
- [27] Huanjing Yue, Cong Cao, Lei Liao, Ronghe Chu, and Jingyu Yang. Supervised raw video denoising with a benchmark dataset on dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [28] 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*, 26:3142–3155, 2017.
- [29] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2808–2817, 2017.
- [30] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn based image denoising. *IEEE Transactions on Image Processing*, 27:4608–4622, 2018.
- [31] Yulun Zhang, Kunpeng Li, Kai Li, Bineng Zhong, and Yun Fu. Residual non-local attention networks for image restoration. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [32] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2020.