

Neural Compression-Based Feature Learning for Video Restoration

Cong Huang^{1*} Jiahao Li² Bin Li² Dong Liu¹ Yan Lu²
¹ University of Science and Technology of China ² Microsoft Research Asia

hcy96@mail.ustc.edu.cn, dongeliu@ustc.edu.cn, {li.jiahao, libin, yanlu}@microsoft.com

Abstract

How to efficiently utilize the temporal features is crucial, yet challenging, for video restoration. The temporal features usually contain various noisy and uncorrelated information, and they may interfere with the restoration of the current frame. This paper proposes learning noise-robust feature representations to help video restoration. We are inspired by that the neural codec is a natural denoiser. In neural codec, the noisy and uncorrelated contents which are hard to predict but cost lots of bits are more inclined to be discarded for bitrate saving. Therefore, we design a neural compression module to filter the noise and keep the most useful information in features for video restoration. To achieve robustness to noise, our compression module adopts a spatial-channel-wise quantization mechanism to adaptively determine the quantization step size for each position in the latent. Experiments show that our method can significantly boost the performance on video denoising, where we obtain 0.13 dB improvement over BasicVSR++ with only 0.23x FLOPs. Meanwhile, our method also obtains SOTA results on video deraining and dehazing.

1. Introduction

Video restoration aims to recover the high-quality video from the degraded input. Typical degradation includes various noises, rain, haze, etc. It has a wide range of applications, but this problem is still under-explored. Different from image restoration that focuses on the intrinsic properties in single image [44], video restoration relies more on extracting and utilizing temporal features for better quality.

Recent video restoration methods mainly focus on network structure design for better extracting temporal features. For example, RViDeNet [43] and EDVR [36] use deformable convolution to align the features of neighboring frames. BasicVSR [7] designs a bi-directional feature propagation network. BasicVSR++ [8] introduces the second-order grid propagation network structure and flow-guided

*This work was done when Cong Huang was an intern at Microsoft Research Asia.

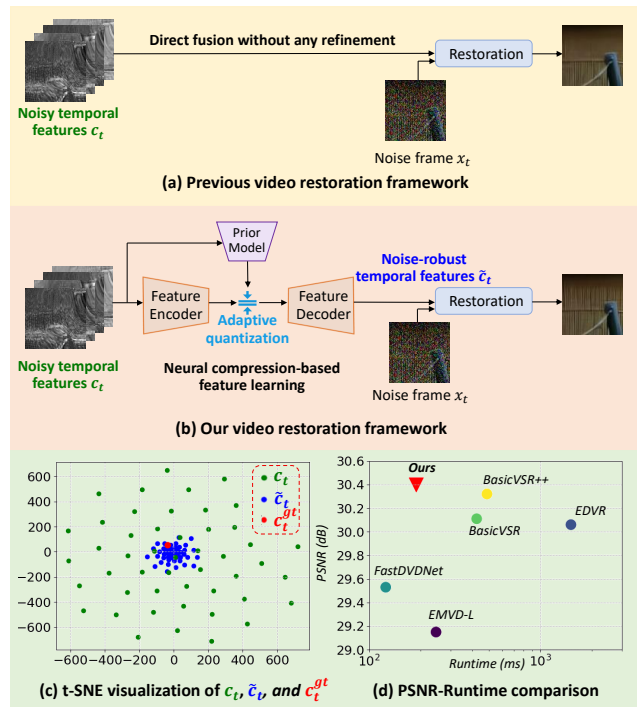


Figure 1. (a) Previous framework without temporal feature refinement. (b) Our framework via neural compression-based feature learning. (c) t-SNE [35] visualization. c_t^{gt} are sampled from clean video (hypersmooth, Set8 [34]). For c_t and \tilde{c}_t , we add different additive white Gaussian noises (same noise σ but different noise random seeds) to the same input video to sample these feature points. It shows that \tilde{c}_t are more robust to noise and get closer to c_t^{gt} . (d) Performance comparison on video denoising (Set8, noise $\sigma = 50$).

deformable alignment network. However, these methods directly use the extracted temporal features without any refinement. The temporal features usually contain lots of noisy and irrelevant information, which interferes with the restoration of the current frame. In this paper, we take video denoising as a case study and explore how to utilize the extracted temporal features efficiently.

We propose a novel neural compression-based solution to refine the features and learn noise-robust feature repre-

sentations. From the perspective of neural codec, the noisy data usually contains lots of high-frequency and is hard to predict. To save the bitrate, codec prefers to discard these noisy and uncorrelated contents. This motivates us to design a neural compression module to purify the temporal features and filter the noisy information therein for video restoration. To achieve robustness to noise, namely let the representations of the noise-perturbed data be mapped to the same quantized representation with the clean data with high probability, the quantization step needs to be properly set. However, most existing neural compression frameworks only support fixed quantization step size. This cannot meet our purpose and even harms the inherent textures. To solve this problem, we design an adaptive quantization mechanism at spatial-channel-wise for our compression module, where the quantization step is learned by our prior model. Our quantization mechanism can adaptively purify the features with different content characteristics. During the training, the cross-entropy loss is used to guide the learning of the compression module and helps preserve the most useful information.

Fig. 1 shows the framework comparison. From the t-SNE [35] visualization shown in Fig. 1 (c), we find, via our neural compression-based feature learning, the features are more robust to noise and get closer to the features generated from the clean video. Fig. 1 (d) is the performance comparison. We observe that, empowered by the noise-robust feature representations, our framework significantly improves the restoration quality, when compared with prior state-of-the-art (SOTA) methods. The major contributions of this paper are summarized as follows:

- We propose a novel neural compression-based feature learning for video restoration. After processed by our compression module, the features are more robust to noise and then improve the restoration quality.
- To achieve robustness to noise and adaptively purify the features with different content characteristics, we design a learnable quantization mechanism at spatial-channel-wise.
- To further boost the performance, we also design an attention module to help the feature learning, and a motion vector refinement module to improve the discontinuous motion vector estimated from noisy video.
- We propose a lightweight framework. Compared with previous SOTA methods, our method achieves a better quality-complexity trade-off on video denoising, deraining, and dehazing.

2. Related Work

2.1. Video Restoration

Existing video restoration methods leveraging temporal correlation can be divided into two categories: the sliding

window-based methods and the recurrent methods.

The sliding window-based methods take several adjacent frames as input for each frame. Some methods [10, 34] do not depend on explicit motion alignments. VNLNet [10] uses the non-local module to search similar patches across frames. FastDVDNet [34] uses stacked U-Net [32] to progressively fuse the unaligned neighboring frames. By contrast, ToFlow [39] and DVDNet [33] use a motion estimation component to explicitly align neighboring frames. To explore more temporal correlation, RViDeNet [43] and EDVR [36] propose feature domain alignment. They align the features of neighboring frames rather than raw pixels and this mechanism is adopted by most recent methods.

The sliding window-based methods suffer from a narrow temporal scope and cannot leverage the information outside the sliding window. By contrast, the recurrent methods learn the temporal features within a long temporal range, and achieve better performance. EMVD [25] recurrently combines all past frames as auxiliary information. Yan *et al.* [40] proposed a recurrent feature propagation framework without explicit alignment. The feature propagation in BasicVSR [7] uses the explicit alignment. Recently, BasicVSR++ [8] achieves excellent performance by using a second-order grid propagation structure and a flow-guided deformable alignment module.

2.2. Video Compression

Traditional video codecs, e.g. H.264 and H.265, adopt the hybrid framework which consists of prediction, transform, quantization, entropy coding, and loop-filter. Benefiting from the progress of neural image compression [3, 4, 27], neural video compression [2, 17, 19, 20, 24] recently also has a great development. For example, Lu *et al.* [24] designed the DVC model, which follows the framework of traditional video codec but uses neural networks to implement all modules therein. Following DVC, Agustsson *et al.* [2] designed a more advanced optical flow estimation in scale space. Recently Li *et al.* [17] proposed a conditional coding-based framework which achieves better performance.

3. Motivation

Our motivation comes from that video compression can filter the noise. Video compression aims at using the least bitrate cost to represent the video. For traditional codec, the residuals of noisy contents are usually large as they are hard to predict from reference frames. These residuals contains lots of high-frequency and will consume many bits. To achieve the bitrate saving, traditional codec uses the quantization to discard the residuals of noisy contents, especially for the high-frequency therein, which is like a low-pass filter. We use the traditional codec x265 [1] to conduct an analysis experiment, as shown in Fig. 2. From Fig. 2 (c), we find that the traditional codec x265 can filter the noise in a




Figure 2. Comparison between traditional codec x265 [1] and neural codec [17] when compressing the noised video (additive white Gaussian noise with $\sigma = 20$). BPP denotes bits per pixel, which measures bitrate cost.

large degree. Fig. 2 (d) shows, when allocated more bits, x265 will encode the noise but in a much smoother way.

Different from traditional codec using linear DCT (discrete cosine transform), neural codec will learn a neural encoder to transform video from pixel domain to latent feature domain. The latent feature is then quantized, and its distribution is estimated to perform arithmetic coding. The distribution is predicted more accurately, more bitrate saving is achieved. However, the distributions of the noisy and uncorrelated contents are hard to predict well. Thus, to save the bitrate, these contents are more inclined to be discarded guided by the cross-entropy loss. Fig. 2 (e) shows the effectiveness of neural codec [17] (model weights are provided by authors of [17]). In particular, the neural codec can much better remove the noise therein and keep more semantic information when compared with x265.

Inspired by this analysis, we propose utilizing a neural codec to help video restoration. The neural codec is used to filter the noisy information in features via the quantization. If the quantization step and data distribution are properly learned, the representations of the noise-perturbed data will be mapped to the same quantized representation with the clean data with high probability. The noise-robust feature representations will improve the final restoration quality.

Another advantage of using neural codec rather than traditional codec is that the neural codec can be end-to-end trainable and will have better performance when jointly trained with other restoration modules.

4. Proposed Method

4.1. Framework Overview

We design a neural compression-based framework for video restoration. Our framework contains three parts: feature alignment, feature refinement for learning noise-robust feature representations, and feature fusion. The framework is illustrated in Fig. 3.

Feature alignment. Given the noisy frames x_{t-1} and x_t , we first use motion estimation to estimate the motion vector (MV) mv_t . Then we design an MV refinement mod-

ule to improve the discontinuous MV mv_t estimated from noisy video. With the refined MV \widetilde{mv}_t , the coarse features \hat{c}_t are obtained via a bilinear warping function.

Feature refinement. As \hat{c}_t contains some noisy and uncorrelated information, we propose a neural compression-based feature refinement to purify the features. It is noted that, our feature refinement part consists of two modules. One is the attention module and the other is the neural compression module used for noise-robust feature learning.

Feature fusion. With the noise-robust features \tilde{c}_t and the current frame x_t , the final output frame y_t is generated through the restoration module. Besides y_t , the restoration module part will also generate the temporal features c_t used for next step.

4.2. Feature Alignment

To align the temporal features from last step to the current frame, we need to predict the MV. In our paper, we use the pre-trained optical flow estimation network SPyNet [30] as our motion estimation module.

However, estimating accurate MV from the degraded frames is quite difficult. As shown in Fig. 4 (a), the MV without any processing suffers from corruption and discontinuity, which is not accurate compared with the MV estimated from clean frames in Fig. 4 (c). To solve this problem, we propose using an MV refinement module to improve the MV. The MV refinement module adopts a lightweight auto-encoder structure. It encodes the corrupted MV into compact representations and then decodes them to the refined MV. The detailed network structure can be found in the supplementary materials. As Fig. 4 (b) shows, with our MV refinement, the MV is cleaner and more similar to the MV from the clean frames.

4.3. Feature Refinement via Neural Compression

Previous recurrent methods directly fuse the current frame and the aligned temporal features without any refinement. Actually, the temporal features may still contain some noisy and uncorrelated information, which disturbs the restoration of the current frame.




Figure 3. The overall framework of our method. The detailed network structures of each module can be found in supplemental materials.



Figure 4. Example of MV comparison. We find MV is more accurate and more similar to the MV from clean frames after ME refinement. Zoom in for a better view.

To solve this problem, we propose a feature refinement process to learn noise-robust feature representations. This process consists of two modules, i.e. attention module and neural compression module. For the attention mechanism, many papers [14, 26, 28, 46] have studied it and proved its effectiveness. Thus, we design an attention module to scale the temporal features to help the feature learning. To achieve a good trade-off between performance and complexity, we design an auto-encoder-based attention network whose detailed network structure can be found in the supplementary materials.

After the attention module, the temporal features \check{c}_t will be purified by the proposed neural compression module. Following the design in neural image/video compression [4, 17, 24], our neural compression module consists of feature encoder-decoder, quantization process, and a prior model.

First, the temporal features \check{c}_t are encoded to be compact latent codes e_t through the feature encoder:

$$e_t = \text{Encoder}(\check{c}_t). \quad (1)$$

To achieve the robustness to noise, the quantization is applied to e_t . The $e_t \in [s_k, s_{k+1})$ is quantized to value $\frac{s_k + s_{k+1}}{2}$, where s_k and s_{k+1} indicate the numerical range. Let $\check{\check{c}}_t = \check{c}_t + \epsilon$ be the noisy feature with noise ϵ . Under the assumption that Encoder is Lipschitz continuous,

$\hat{e}_t = \text{Encoder}(\check{\check{c}}_t)$ and $e_t = \text{Encoder}(\check{c}_t)$ will be located in the same region $[s_k, s_{k+1})$ with high probability if the quantization step $s_{k+1} - s_k$ is relatively large, then they have the same quantized value. That means, the quantized representation is robust to noisy input. However, the robustness is determined by a prerequisite that the data distribution and quantization step are properly learned.

Most existing quantization solutions in neural image/video compression only use fixed quantization step. Actually, the content characteristics spatially vary in a large degree. A fixed quantization step cannot handle various and complex contents well. For example, a fixed small quantization step fails to remove the noisy information. A fixed large quantization step instead causes large information loss (i.e., the intrinsic quantization noise). Thus, we propose an adaptive quantization mechanism, where the quantization step is learned. The illustration is shown in Fig. 5. First, e_t are divided by the learned quantization step q_t after subtracting the learned mean value μ_t . The quotients are then rounded to the closest integers. At last, the quantized latent codes \hat{e}_t are obtained via the opposite operations. The formulation is:

$$\hat{e}_t = \lceil \frac{e_t - \mu_t}{q_t} \rceil * q_t + \mu_t. \quad (2)$$

$\lceil \cdot \rceil$ is the integer rounding operation. With the quantized latent codes \hat{e}_t , the noise-robust temporal features \tilde{c}_t are then decoded through the feature decoder:

$$\tilde{c}_t = \text{Decoder}(\hat{e}_t). \quad (3)$$

As aforementioned, the data distribution and quantization step need to be properly learned to achieve noise robustness. In practice, we do not know the data distribution, thus we use the prior model to estimate it, and then use the

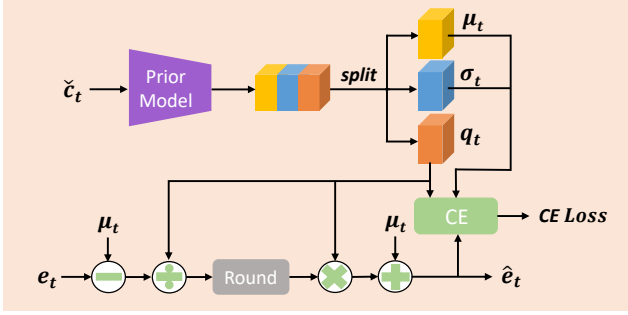


Figure 5. The illustration of adaptive quantization mechanism. CE means cross-entropy.

cross-entropy loss to guide the learning of data distribution and quantization step. The cross-entropy loss is formulated as:

$$\mathcal{L}_{oss_{CE}} = E_{\hat{e}_t}[-\log_2 p_{\hat{e}_t}(\hat{e}_t)], \quad (4)$$

where $p_{\hat{e}_t}(\hat{e}_t)$ is the estimated probability mass function of the latent codes \hat{e}_t . In this paper, we follow [15, 17] and assume that $p_{\hat{e}_t}(\hat{e}_t)$ follows the Laplace distribution. The prior model composed by neural network is used for estimating the distribution parameters. Detailed structures of the prior model can be found in supplemental materials. But different from [15, 17] that only estimate the distribution parameters (μ_t, σ_t) , our prior model also learns the quantization step size q_t . With (μ_t, σ_t, q_t) , the probability estimation of $p_{\hat{e}_t}(\hat{e}_t)$ is calculated as:

$$p_{\hat{e}_t}(\hat{e}_t) = \prod_i (\mathcal{L}(\mu_{t,i}, \sigma_{t,i}^2) * \mathcal{U}(-\frac{q_{t,i}}{2}, \frac{q_{t,i}}{2}))(\hat{e}_{t,i}), \quad (5)$$

where i specifies the spatial position of each element in \hat{e}_t . According the probability mass function in Eq. 5, we can calculate the cross-entropy loss via Eq. 4. The cross-entropy loss guides the compression module to learn proper data distribution and quantization step, and then achieves the robustness to noise.

In our framework, the quantization step q_t is learnable at spatial-channel-wise. It can be adaptive to regions with different content characteristics. We visualize one channel example of quantization step map in Fig. 6. The pixel intensity represents the size of quantization step. The larger pixel intensity represents the more noisy information that should be eliminated. As Fig. 6 shows, the quantization step size of the smooth region is usually larger because the noisy information therein is easier to remove. By contrast, the quantization step size of texture region (e.g., the table in the background poster, and there actually exist many details) is usually smaller.

4.4. Feature Fusion

The feature fusion part contains a restoration module. It will fuse the noise-robust temporal features \tilde{c}_t with the current frame x_t , and then generates the final output frame \tilde{y}_t .




Figure 6. Visualization example of the input frame and the corresponding learned quantization step size map. The frame is from Johnny, HEVC Class E dataset [5].

It is noted that, besides final output frame \tilde{y}_t , the restoration module also generates the temporal features \tilde{c}_t used for the next step, like [7]. Our restoration module is based on a lightweight U-Net [32]. The detailed network structures can be found in supplemental materials.

4.5. Loss Function

In our method, the loss function includes two items:

$$\mathcal{L}_{oss} = \sum_{t=1}^n \mathcal{L}_{oss_{L2}}(y_t, \tilde{y}_t) + \lambda \cdot \mathcal{L}_{oss_{CE}}(\hat{e}_t). \quad (6)$$

y_t and \tilde{y}_t are the clean and estimated frames, respectively. $\mathcal{L}_{oss_{L2}}$ is the L2 loss and $\mathcal{L}_{oss_{CE}}$ is the cross-entropy loss. To learn noise-robust feature representation and then let it help the final reconstruction, we adopt a two-stage training scheme and the details are in the supplementary materials.

5. Experiment

We evaluate our method on several video restoration tasks, including denoising, deraining, and dehazing.

5.1. Dataset

Video denoising. Both synthetic dataset and real-world dataset are tested. For the synthetic dataset, we follow the setting in FastDVDNet [34]. DAVIS2017 train-val set containing 90 videos is used for training. Set8 is used for testing. We add the additive white Gaussian noise (AWGN) to the clean video to synthesize the noisy video. Five noise levels, i.e. $\sigma=10, 20, 30, 40, 50$ are tested. For the real-world dataset, we follow the setting in EMVD [25] and use the dataset from RViDeNet [43]. It consists of a captured raw video dataset (CRVD) and a synthetic raw video dataset (SRVD). Following EMVD and RviDeNet, we use CRVD scene 1~6 plus SRVD for training and CRVD scene 7~11 for testing.

Video deraining. Following [42], we test our method on RainSynComplex25 [21] and RainSynAll100 [42] datasets. RainSynComplex25 contains 190 videos for training and 25 videos for testing. RainSynAll100 contains 900 videos for training and 100 videos for testing.

σ	VNLnet [10]	DVDNet [33]	FastDVDNet [34]	EMVD-L [25]	EMVD-S [25]	EDVR [36]	BasicVSR [7]	BasicVSR++ [8]	Ours
10	37.10/0.9637	36.08/0.9592	36.44/0.9624	36.56/0.9624	35.01/0.9442	37.16/0.9658	37.12/0.9674	37.27/0.9682	37.17/0.9684
20	33.88/0.9360	33.49/0.9307	33.43/0.9334	33.27/0.9320	31.65/0.8927	34.09/0.9379	34.13/0.9397	34.25/0.9411	34.22/0.9437
30	31.95/0.9096	31.79/0.9023	31.68/0.9066	31.40/0.9032	29.94/0.8678	32.31/0.9125	32.33/0.9157	32.55/0.9168	32.57/0.9184
40	30.55/0.8814	30.55/0.8745	30.46/0.8812	30.05/0.8761	28.64/0.8328	31.02/0.8887	31.05/0.8929	31.28/0.8936	31.39/0.8970
50	29.47/0.8561	29.56/0.8480	29.53/0.8573	29.15/0.8528	27.83/0.8082	30.06/0.8660	30.11/0.8690	30.32/0.8696	30.45/0.8770
FLOPs (G)	-	-	665	1106	5	3089	2947	3402	771

Table 1. PSNR/SSIM comparison with SOTA video denoising methods on synthetic dataset Set8. The best performance is highlighted in red (1st best) and blue (2nd best). Our method achieves the best SSIM on all noise levels.

	FastDVDNet [34]	EDVR [36]	RViDeNet [43]	EMVD-L [25]	EMVD-S [25]	BasicVSR [7]	BasicVSR++ [8]	Ours	Ours-L
PSNR	44.30	44.71	44.08	44.48	42.63	44.80	44.98	44.72	45.09
SSIM	0.9881	0.9902	0.9881	0.9895	0.9851	0.9903	0.9903	0.9906	0.9909
Runtime (ms)	132	1511	1254	246	59	425	488	188	275

Table 2. Comparison with SOTA video denoising methods on real-world dataset CRVD [43]. Our method with default setting outperforms other fast methods and gets close to the slow method. With a more powerful restoration module (i.e. 'Ours-L'), we can achieve SOTA performance in terms of both PSNR and SSIM. The runtime is the average frame runtime for the whole dataset on a P100 GPU.

Video dehazing. We use REVIDE [45] dataset which captures the pairs of hazy and corresponding haze-free videos in the same scene by an acquisition system. It contains 42 videos for training and 6 videos for testing.

5.2. Result on Video Denoising

We compare our method with these baselines: VNL-Net [10], DVDNet [33], FastDVDNet [34], EMVD [25], EDVR [36], BasicVSR [7], BasicVSR++ [8], and RViDeNet [43]. EMVD has several network structure configurations with different complexities. The large (EDVR-L) and small (EMVD-S) models are tested (more details about configurations are in supplementary materials). The original BasicVSR/BasicVSR++ are bi-directional methods that leverage the temporal features from both the future and the past frames. To compare with other methods more fairly, we modified BasicVSR/BasicVSR++ to uni-directional methods that only use the temporal features from the past frames.

Quantitative Comparison. We use peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) as quantitative evaluation metrics. We present the results on synthetic noise video in Table 1 and the results on real-world noisy video in Table 2. For the synthetic video, as Table 1 shows, our method achieves the best SSIM on all noise levels. For PSNR, our method outperforms the second-best method BasicVSR++ [8] with at least 0.11dB gain when the noise level is high ($\sigma = 40$ or 50). In addition, we find the quality improvement over BasicVSR++ is larger when the noise level is higher. It verifies that our proposed neural compression module can effectively filter the noise. It is also worth noting that the FLOPs of our method is only 0.23 times of BasicVSR++, which shows that our method achieves a much better trade-off between quality and complexity. Compared with the low-complexity methods FastDVDNet [34] and EMVD-L [25], our method achieves significant quality improvement. For the real-world noisy video, it should be admitted that our method with default setting currently cannot

outperform BasicVSR and BasicVSR++ in terms of PSNR but is better than them in terms of SSIM. When compared with low-complexity method FastDVDNet and EMVD-L, our method can achieve the best quality. Besides, if we change our U-Net-like [32] restoration network to a W-Net-like [38] restoration network with more complexity (more details are in supplemental materials), denoted as 'Our-L' in Table 2, we can achieve the best PSNR and SSIM at the same time, but the complexity is still much less than that of BasicVSR and BasicVSR++.

Qualitative Comparison. Fig. 7 shows the visual quality comparison. As Fig. 7 shows, FastDVDNet without feature alignment suffers from serious distortion in the text region. The results of BasicVSR++ are quite blurry caused by the propagated noise in temporal features. By contrast, our neural compression-based method can learn noise-robust features and is able to restore much clearer textures. More visual comparisons are in the supplementary materials.

5.3. Result on Video Deraining

We compare our method with prior SOTA video deraining methods, including MS-CSC [18], SE [37], SpacCNN [9], FastDerain [16], J4RNet-P [21], FCRVD [41], RMFD [42], and BasicVSR++ [8]. Since RainSynAll100 uses the rain accumulation degradation to generate the rainy video, parts of the baseline methods including SE, MS-CSC, SpacCNN, and FastDerain could not handle this degradation thus MRF [6] is used as post-processing. More details could be found in [42]. FCRVD, RMFD, BasicVSR++, and our method could handle this degradation without extra post-processing. As Table 3 shows, BasicVSR++ beats RMFD in terms of PSNR and SSIM on RainSynComplex25 but is slightly worse than RMFD in terms of SSIM on RainSynAll100. By contrast, with the robust temporal features, our method achieves the best PSNR and SSIM on both datasets. Our method brings PSNR gain of 0.44 dB and SSIM gain of 0.0063 on RainSynAll100. We also test RainSynLight25 [21] and NTRain [9]. Their

		MS-CSC [18]	SE [37]	SpacCNN [9]	FastDerain [16]	J4RNet-P [21]	FCRVD [41]	RMFD [42]	BasicVSR++ [8]	Ours
RainSynAll100	PSNR	16.19	15.29	18.39	17.09	19.26	21.06	25.14	27.67	28.11
	SSIM	0.5078	0.5053	0.6469	0.5824	0.6238	0.7405	0.9172	0.9135	0.9235
RainSynComplex25	PSNR	16.96	16.76	21.21	19.25	24.13	27.72	32.70	33.42	34.27
	SSIM	0.5049	0.5273	0.5854	0.5385	0.7163	0.8239	0.9357	0.9365	0.9434

Table 3. Comparison with SOTA video deraining methods on RainSynComplex25 [21] and RainSynAll100 [42]. We train the BasicVSR++ [8] using the same setting as ours. Other baseline results are provided by RMFD [42] paper.

	DCP [12]	GNet [22]	DuRN [23]	KDDN [13]	MSBDN [11]	FFA [29]	VDN [31]	EDVR [36]	CG-IDN [45]	BasicVSR++ [8]	Ours
PSNR	11.03	19.69	18.51	16.32	22.01	16.65	16.64	21.22	23.21	21.68	23.63
SSIM	0.7285	0.8545	0.8272	0.7731	0.8759	0.8133	0.8133	0.8707	0.8836	0.8726	0.8925

Table 4. Comparison with SOTA video dehazing methods on REVIDE [45] testset. We train the BasicVSR++ [8] using the same setting as ours. Other baseline results are provided by CG-IDN [45] paper.

	M_a	M_b	M_c	M_d
MVR		✓	✓	✓
NCFL			✓	✓
FA				✓
PSNR	29.75	29.87	30.29	30.45

Table 5. The ablation study on different modules. Tested on Set8 ($\sigma = 50$). MVR is the MV refinement, NCFL is the neural compression-based feature learning. FA is the feature attention.

	NCFL-AdapQ	NCFL-FixedQ	NCFL-NoQ
PSNR	30.29	29.86	29.98

Table 6. The ablation study on quantization. Tested on Set8 ($\sigma = 50$). NCFL-AdapQ is our default model with adaptive quantization, i.e. M_c in Table 5. NCFL-FixedQ means we use the fixed quantization step as in existing neural video codecs. NCFL-NoQ removes the quantization and is only a vanilla auto-encoder.

results are provided in the supplementary materials. Fig. 7 also shows the visual quality comparison. We can see that our model could remove rain streak well, and produces clearer and more visual pleasing results.

5.4. Result on Real-World Video Dehazing

Table 4 shows the comparison between our method and prior SOTA real-world video dehazing method: DCP [12], GNet [22], DuRN [23], KDDN [13], MSBDN [11], FFA [29], VDN [31], EDVR [36], CG-IDN [45], and BasicVSR++ [8]. As Table 4 shows, BasicVSR++ outperforms EDVR but is worse than MSBDN and CG-IDN that are specially designed for dehazing task. By contrast, our method has 0.42 dB PSNR and 0.0089 SSIM improvements when compared with the second-best method CG-IDN. In addition, the parameters of our method are 16M, and only 0.70 times of CG-IDN that has 23M parameters. As Fig. 7 shows, the result of our method is more visual pleasing.

5.5. Ablation Study

This paper proposes three key modules: the MV refinement (MVR) for improving MV, the neural compression-based feature learning (NCFL) incorporated with adaptive

quantization, and feature attention (FA). We study the effect of these modules and report the results in Table 5. Without MVR, NCFL, and FA, the baseline model only contains a motion estimation module, bilinear warping, and a restoration module.

MV refinement (MVR). As Table 5 shows, the baseline model M_a only achieves PSNR 29.75dB. It suffers from the discontinuous MV estimated from noisy video. When enabling our MVR, the MV is refined and M_b reaches PSNR 29.87 dB. Our MVR brings 0.12 dB PSNR improvement.

Neural compression-based feature learning (NCFL). If we further combine NCFL and MVR, M_c reaches PSNR 30.29 dB and improves 0.42 dB compared with M_b . The significant improvement verifies that the effectiveness of NCFL. In addition, we also study two variants of NCFL. As Table 6 shows, the PSNR of a vanilla auto-encoder without quantization (i.e. NCFL-NoQ) drops to 29.98dB. This denotes that the improvement brought by NCFL mainly comes from the adaptive quantization mechanism rather than the increase of model parameters. In addition, we also test NCFL-FixedQ where a fixed quantization step is used as many existing neural video codecs do. The PSNR of NCFL-FixedQ drops to 29.89 dB. Its performance is even worse than NCFL-NoQ. This shows that a fixed quantization step instead loses some useful information and fails at learning noise-robust representations. By contrast, a learnable quantization step at spatial-channel-wise can adaptively filter the noise and purify the temporal features with different content characteristics, which is quite important.

Feature attention (FA). In this paper, we also propose a FA module to further help the feature learning. As Table 5 shows, M_d achieves PSNR 30.45 dB. FA boosts the PSNR by 0.16 dB, which shows its effectiveness.

5.6. NCFL on Different Degradations

Table 5 and Table 6 investigate NCFL under AWGN degradation. However, our NCFL is not confined to AWGN. It is also very effective for other complex degradations, such as real-world denoising, deraining, and dehazing. Table 7 shows the comprehensive study on multiple degradations. For example, the comparison between M_1 and M_3

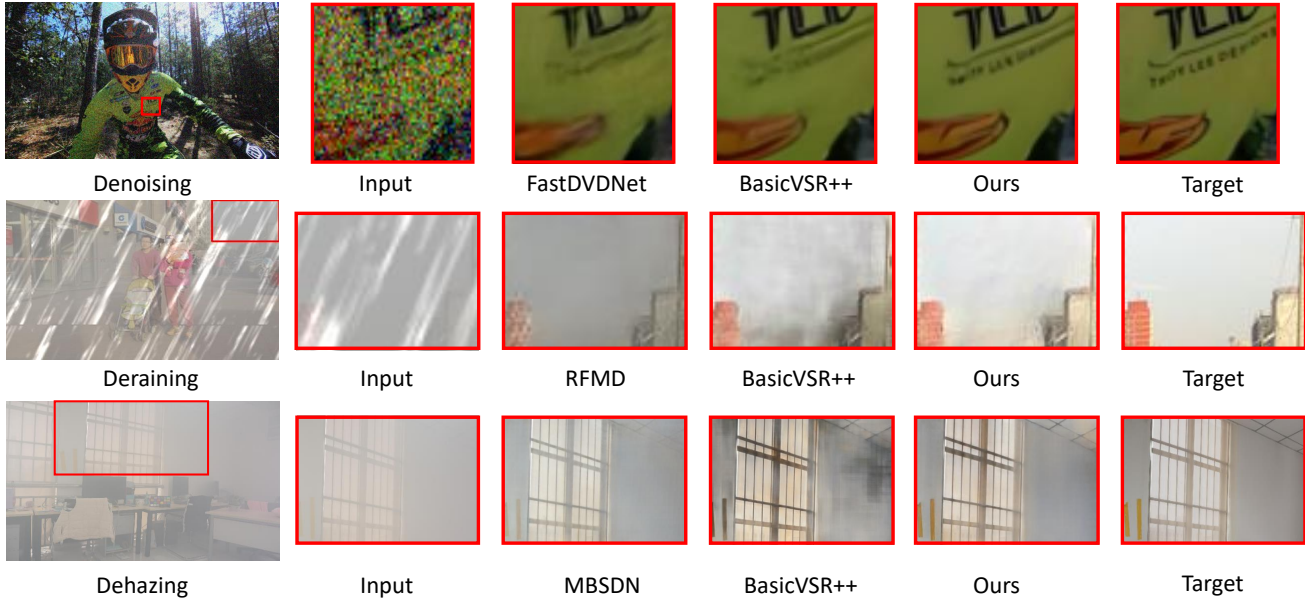


Figure 7. Denoising: *motorbike* video from Set8 testset with noise variance 50. Deraining: *0985* video from RainSynAll100 testset. Dehazing: *L006* video from REVIDE dataset.

	Deraining	Dehazing	RWD	AWGN
M_1 : w/o NCFL	27.30	23.07	44.48	29.99
M_2 : w/ NCFL (w/o CE)	27.64	23.30	44.56	30.20
M_3 : w/ NCFL	28.11	23.63	44.72	30.45

Table 7. Study on NCFL under different degradation types. CE means the cross-entropy loss. RWD means real-world denoising. AWGN indicates additive white Gaussian noise.

shows that NCFL can achieve 0.81 dB gain for deraining. These substantial improvements verify the effectiveness of our NCFL. In addition, the comparison between M_2 and M_3 shows that cross-entropy loss can effectively guides the learning of NCFL under multiple degradations.

5.7. Bi-directional Video Denoising

In previous experiments, we focus on the uni-directional setting where the temporal features only come from the past time. For the bi-directional setting, both the temporal features come from the past time and the future time can be used. One advantage of our method is that our method can be extended to bi-directional setting easily. We test the bi-directional model of BasicVSR [7], BasicVSR++ [8], and our bi-directional model. The PSNR and complexity comparison is shown in Table 8. As Table 8 shows, the bi-directional setting brings BasicVSR 0.59 dB gain, BasicVSR++ 0.78 dB gain, and our method 0.76 dB gain with about 2x complexity. Under bi-directional setting, our method still outperforms BasicVSR++ by 0.11 dB with only 0.21x FLOPs.

Direction	Method	PSNR	FLOPs (G)
Uni-direction	BasicVSR	30.11	2947
Uni-direction	BasicVSR++	30.32	3402
Uni-direction	Ours	30.45	771
Bi-direction	BasicVSR	30.68	5855
Bi-direction	BasicVSR++	31.10	7097
Bi-direction	Ours	31.21	1522

Table 8. Bidirectional video denoising on Set8 with $\sigma = 50$.

6. Conclusion and Limitation

In this paper, we have designed a neural compression-based video restoration framework. We are inspired by the fact that neural video codec can naturally filter the noise, and then propose using neural compression to purify the temporal features and learn noise-robust feature representations. To solve the problem that the fixed quantization step harms the inherent textures, we propose a learnable quantization mechanism at spatial-channel-wise to achieve robustness to noise. At the same time, an attention module and an MV refinement module are proposed to further boost the performance. Experimental results show that the proposed method achieves a much better quality-complexity trade-off than previous SOTA methods.

Although our method is faster than most previous SOTA methods, the inference speed of our method still does not meet the requirements of real-time scenarios. In the future, we will continue to improve the efficiency of our method for real-time video restoration.

References

- [1] Ffmpeg. <https://www.ffmpeg.org/>. 2, 3
- [2] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8503–8512, 2020. 2
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*, 2016. 2
- [4] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018. 2, 4
- [5] Frank Bossen et al. Common test conditions and software reference configurations. In *JCTVC-L1100*, volume 12, 2013. 5
- [6] Bolun Cai, Xiangmin Xu, and Dacheng Tao. Real-time video dehazing based on spatio-temporal mrf. In *Pacific Rim conference on multimedia*, pages 315–325. Springer, 2016. 6
- [7] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. 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*, pages 4947–4956, 2021. 1, 2, 5, 6, 8
- [8] Kelvin CK Chan, Shangchen Zhou, Xiangyu Xu, and Chen Change Loy. BasicVSR++: Improving Video Super-Resolution with Enhanced Propagation and Alignment. *arXiv preprint arXiv:2104.13371*, 2021. 1, 2, 6, 7, 8
- [9] Jie Chen, Cheen-Hau Tan, Junhui Hou, Lap-Pui Chau, and He Li. Robust video content alignment and compensation for rain removal in a cnn framework. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6286–6295, 2018. 6, 7
- [10] Axel Davy, Thibaud Ehret, Jean-Michel Morel, Pablo Arias, and Gabriele Facciolo. Non-local video denoising by cnn. *arXiv preprint arXiv:1811.12758*, 2018. 2, 6
- [11] Hang Dong, Jinshan Pan, Lei Xiang, Zhe Hu, Xinyi Zhang, Fei Wang, and Ming-Hsuan Yang. Multi-scale boosted dehazing network with dense feature fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2157–2167, 2020. 7
- [12] Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *IEEE transactions on pattern analysis and machine intelligence*, 33(12):2341–2353, 2010. 7
- [13] Ming Hong, Yuan Xie, Cuihua Li, and Yanyun Qu. Distilling image dehazing with heterogeneous task imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3462–3471, 2020. 7
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 4
- [15] Zhihao Hu. Pytorch video compression, <https://github.com/zhihaohu/pytorchvideocompression>, 2020. 5
- [16] Tai-Xiang Jiang, Ting-Zhu Huang, Xi-Le Zhao, Liang-Jian Deng, and Yao Wang. Fastderain: A novel video rain streak removal method using directional gradient priors. *IEEE Transactions on Image Processing*, 28(4):2089–2102, 2018. 6, 7
- [17] Jiahao Li, Bin Li, and Yan Lu. Deep contextual video compression. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 3, 4, 5
- [18] Minghan Li, Qi Xie, Qian Zhao, Wei Wei, Shuhang Gu, Jing Tao, and Deyu Meng. Video rain streak removal by multi-scale convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6644–6653, 2018. 6, 7
- [19] Jianping Lin, Dong Liu, Houqiang Li, and Feng Wu. M-LVC: multiple frames prediction for learned video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [20] Jerry Liu, Shenlong Wang, Wei-Chiu Ma, Meet Shah, Rui Hu, Pranaab Dhawan, and Raquel Urtasun. Conditional entropy coding for efficient video compression. *arXiv preprint arXiv:2008.09180*, 2020. 2
- [21] Jiaying Liu, Wenhan Yang, Shuai Yang, and Zongming Guo. Erase or fill? deep joint recurrent rain removal and reconstruction in videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3233–3242, 2018. 5, 6, 7
- [22] Xiaohong Liu, Yongrui Ma, Zhihao Shi, and Jun Chen. Grid-dehazenet: Attention-based multi-scale network for image dehazing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7314–7323, 2019. 7
- [23] Xing Liu, Masanori Suganuma, Zhun Sun, and Takayuki Okatani. Dual residual networks leveraging the potential of paired operations for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7007–7016, 2019. 7
- [24] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. DVC: an end-to-end deep video compression framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11006–11015, 2019. 2, 4
- [25] Matteo Maggioni, Yibin Huang, Cheng Li, Shuai Xiao, Zhongqian Fu, and Fenglong Song. Efficient multi-stage video denoising with recurrent spatio-temporal fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3466–3475, June 2021. 2, 5, 6
- [26] Yiqun Mei, Yuchen Fan, Yulun Zhang, Jiahui Yu, Yuqian Zhou, Ding Liu, Yun Fu, Thomas S Huang, and Humphrey Shi. Pyramid attention networks for image restoration. *arXiv preprint arXiv:2004.13824*, 2020. 4
- [27] David Minnen, Johannes Ballé, and George Toderici. Joint autoregressive and hierarchical priors for learned image compression. *arXiv preprint arXiv:1809.02736*, 2018. 2
- [28] Ben Niu, Weilei Wen, Wenqi Ren, Xiangde Zhang, Lianping Yang, Shuzhen Wang, Kaihao Zhang, Xiaochun Cao, and Haifeng Shen. Single image super-resolution via a holistic attention network. In *European Conference on Computer Vision*, pages 191–207. Springer, 2020. 4

- [29] Xu Qin, Zhilin Wang, Yuanchao Bai, Xiaodong Xie, and Huizhu Jia. FFA-Net: Feature fusion attention network for single image dehazing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11908–11915, 2020. 7
- [30] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4161–4170, 2017. 3
- [31] Wenqi Ren, Jingang Zhang, Xiangyu Xu, Lin Ma, Xiaochun Cao, Gaofeng Meng, and Wei Liu. Deep video dehazing with semantic segmentation. *IEEE Transactions on Image Processing*, 28(4):1895–1908, 2018. 7
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2, 5, 6
- [33] Matias Tassano, Julie Delon, and Thomas Veit. DVDnet: A fast network for deep video denoising. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1805–1809. IEEE, 2019. 2, 6
- [34] Matias Tassano, Julie Delon, and Thomas Veit. Fastdvdnet: Towards real-time deep video denoising without flow estimation. In *CVPR*, pages 1354–1363, 2020. 1, 2, 5, 6
- [35] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008. 1, 2
- [36] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. EDVR: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1, 2, 6, 7
- [37] Wei Wei, Lixuan Yi, Qi Xie, Qian Zhao, Deyu Meng, and Zongben Xu. Should we encode rain streaks in video as deterministic or stochastic? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2516–2525, 2017. 6, 7
- [38] Xide Xia and Brian Kulis. W-net: A deep model for fully unsupervised image segmentation. *arXiv preprint arXiv:1711.08506*, 2017. 6
- [39] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019. 2
- [40] Bo Yan, Chuming Lin, and Weimin Tan. Frame and feature-context video super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5597–5604, 2019. 2
- [41] Wenhan Yang, Jiaying Liu, and Jiashi Feng. Frame-consistent recurrent video deraining with dual-level flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1661–1670, 2019. 6, 7
- [42] Wenhan Yang, Robby T Tan, Jiashi Feng, Shiqi Wang, Bin Cheng, and Jiaying Liu. Recurrent multi-frame deraining: Combining physics guidance and adversarial learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 5, 6, 7
- [43] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 2301–2310, 2020. 1, 2, 5, 6
- [44] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14821–14831, 2021. 1
- [45] Xinyi Zhang, Hang Dong, Jinshan Pan, Chao Zhu, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Fei Wang. Learning to restore hazy video: A new real-world dataset and a new method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9239–9248, 2021. 6, 7
- [46] Yulun Zhang, Kunpeng Li, Kai Li, Bineng Zhong, and Yun Fu. Residual non-local attention networks for image restoration. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. 4