

Bringing Old Films Back to Life

Ziyu Wan¹ Bo Zhang² Dongdong Chen³ Jing Liao^{1*}

¹City University of Hong Kong ²Microsoft Research ³Microsoft Cloud + AI

{raywzy, cddlyf}@gmail.com zhanbo@microsoft.com jingliao@cityu.edu.hk

http://raywzy.com/Old_Film



Figure 1. Real-world old films restored by our method. First row: original video frames. Second row: the restored frames. Third row: the restored frames after further colorization. The quality of old films is considerably enhanced after restoration.

Abstract

We present a learning-based framework, recurrent transformer network (RTN), to restore heavily degraded old films. Instead of performing frame-wise restoration, our method is based on the hidden knowledge learned from adjacent frames that contain abundant information about the occlusion, which is beneficial to restore challenging artifacts of each frame while ensuring temporal coherency. Moreover, contrasting the representation of the current frame and the hidden knowledge makes it possible to infer the scratch position in an unsupervised manner, and such defect localization generalizes well to real-world degradations. To better resolve mixed degradation and compensate for the flow estimation error during frame alignment, we propose to leverage more expressive transformer blocks for spatial restoration. Experiments on both synthetic dataset and real-world old films demonstrate the significant superiority of the proposed RTN over existing solutions. In addition, the same framework can effectively propagate the color from keyframes to the whole video, ultimately yielding compelling restored films. The implementation and model will be released at <https://github.com/raywzy/Bringing-Old-Films-Back-to-Life>.

*Corresponding author.

1. Introduction

Old film classics have the lasting power to strike the resonance and fantasies of audiences today. Unfortunately, many of them have become less popular because people are no longer used to the low resolution and disturbing artifacts caused by the photographic film aging. Film restoration techniques have been developed to bring these old films back to life, which nonetheless takes painstaking efforts. The restoration nowadays is conducted digitally, where the artists meticulously examine each frame, manually retouch the blemishes, fix up the flickering and finally perform the colorization frame by frame, so repairing the entire old film brings insurmountable expenses. Hence, people desire an algorithm that performs all of these tedious tasks automatically such that old films can be revived in modern looking.

Old films typically suffer from mixed degradations, which, to the best of our knowledge, only a few works aim to solve. While one can sequentially apply dedicated models for restoration, the models designed for specific tasks cannot generalize well to real-world degradations. Recently, higher-order degradation models [41, 48] have been proposed to characterize the real-world degradations, yet these works mainly consider the photometric degradations, such as blurriness and noises, rather than the structured defects (e.g., scratches, cracks, etc.) that obstruct the most in old

films. Closely related to our work, [37] attempts to address complex degradations in vintage photos, yet its frame-wise processing on old films does not yield temporally consistent results. The DeepRemaster [14], in comparison, targets video restoration as well as colorization, yet this work cannot sufficiently leverage the temporal information with explicit frame alignment and the spatial information with long-range correlations, thus unable to fix up large cracked areas.

In this work, we seek to unify the entire film restoration tasks with a single framework in which we conduct spatio-temporal restoration. The key insight is that most degradations in old films, especially structured defects, are temporally variant, *i.e.*, the structured defects occluded in one frame may reveal its content in successive frames. Therefore, we propose to repair the degradations by leveraging the spatio-temporal context rather than relying on the hallucination. Specifically, we propose a bi-directional recurrent network (Figure. 2) which aggregates the knowledge of the scene across adjacent frames, effectively reducing the film flickering. The hidden state of the recurrent module embeds the representation of the scene content. After the alignment, the restoration for a specific frame fuses such hidden representation as it offers useful knowledge of the film content underlying the defects. Such a recurrent scheme brings three-fold benefits. First, the film degradations, no matter how severe they are, can be fully restored as long as the information is well-preserved in other frames. Second, the explicit maintenance of the hidden knowledge ensures that the restoration for frames is temporally consistent in a long period. More importantly, the structured defects can be localized in an unsupervised manner because these areas show a larger discrepancy between the representation of the current frame and the hidden state. As opposed to [37] that requires a defect segmentation network, such defect localization is more generalizable to real-world old film degradations.

Spatially, we need a module to account for the slight mismatch during the frame alignment. As such, we propose to leverage the Swin Transformer [25] — even if the hidden representation is not accurately aligned, the interaction of the corresponding pixels can still be modeled through self-attention. Indeed, we observe more stabilized training of the recurrent module due to the use of attention. Besides, thanks to the superior expressivity, the transformer blocks offer improved restoration ability for mixed degradations which would be hard to resolve using a specialized ConvNet [34, 46]. Thus, the proposed network makes the best of the recurrent module and transformers: the memorization nature of recurrent modules benefits the temporal coherency whereas the long-range modeling capability of transformers helps the spatial restoration, which significantly outperforms strong baselines on synthetic datasets, and yields unprecedented quality when restoring real old films.

Moreover, we show that the same framework can be easily

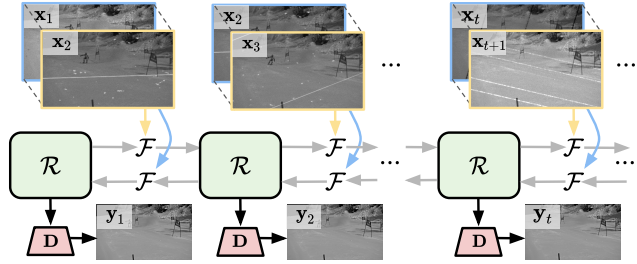


Figure 2. Pipeline overview. Our method follows a bi-directional RNN architecture. \mathcal{R} : Spatial restoration. \mathcal{F} : Feature aggregation. \mathcal{D} : Pixel reconstruction decoder.

adapted for film colorization as well. We follow the colorization pipeline favored by artists that only a few keyframes undergo manual colorization, whose color is then propagated to the rest frames. Our method performs favorably over the leading colorization methods, effectively propagating the color from keyframes to the whole video. As shown in Figure. 1, our method unifies the restoration and colorization, and demonstrates the capability of reviving the old films as if they were captured by yesterday.

2. Related Works

Image Restoration Traditional methods [2, 3, 7, 43] mostly use model-based optimization to restore degraded images, where various image priors are imposed to find suitable solutions. Recently, CNN-based methods have shown remarkable capability for restoration through learning mappings between low-quality and high-quality paired images. However, these works only focus on a single type of degradation like denoising [27, 49, 50], super-resolution [17, 21, 42, 52], inpainting [24, 39, 45], *etc.*, whereas real-world images are plagued with a compound of degradations. Although some methods [23, 41, 48] start to improve the results of real settings via constructing degradation models, they mainly focus on unstructured defects and leave the scratches or dirt on the old medium unfixed. Considering this issue, Wan *et al.* [37, 38] propose a comprehensive solution for old photo restoration by restoring the global and local artifacts in the latent space. Although impressive results can be achieved, they cannot handle the artifacts for dynamic scenes well.

Video Restoration Most existing methods explicitly estimate the dense correspondence among input frames, then reconstruct the clean target frame via CNNs for video denoising [6, 35], video deblurring [33, 53] or video super-resolution [4, 11, 40]. However, they barely consider the real-world mixed degradations and only have limited generation capability. Another research line, video inpainting [8, 16, 44], tries to synthesize plausible contents for the missing regions of each frame conditioned on global information. Yet, these works assume a fixed inpainting mask prescribed by users, whereas the defect areas in old films are typically unknown and temporally varying.

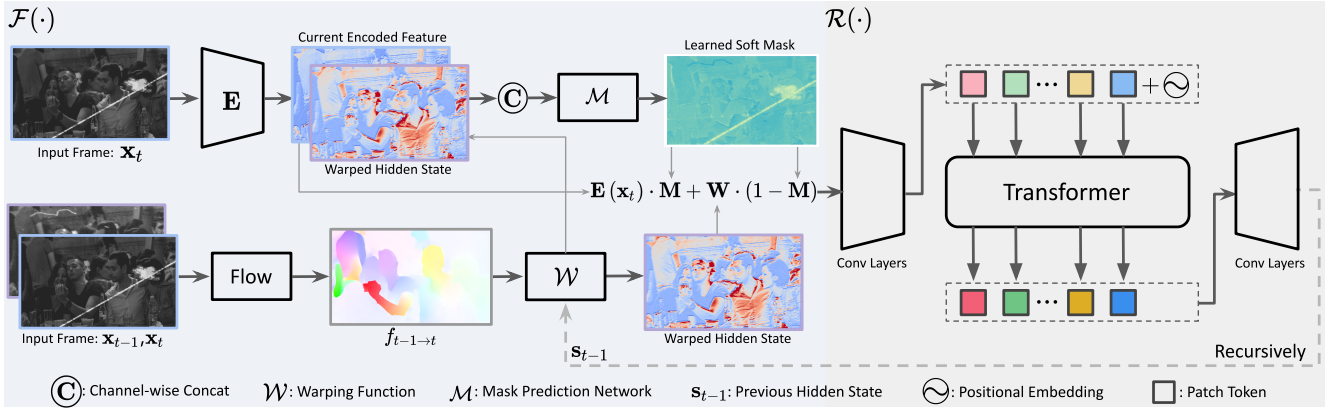


Figure 3. The framework of temporal aggregation module \mathcal{F} and spatial restoration transformer \mathcal{R} in once recurrent forward propagation. Backward propagation follows the same paradigm.

Old Film Restoration Traditional approaches [9, 13, 18, 32] focus on removing the structured artifacts with a detection network followed by an inpainting pipeline, but these works rely on hand-crafted features without a semantic understanding of the video content, which limits their inpainting effects. Moreover, focusing on structured defects while ignoring photometric degradations (*e.g.*, blurriness and noises) makes their overall restoration results less appealing. Subsequently, Iizuka *et al.* propose DeepRemaster [14], a fully 3D convolutional method, to restore the old films. This method works well on restoring the artifacts of synthetic videos, but it fails to generalize well on restoring real-old films, especially the ones with large cracked areas because it lacks special designs for complex mixed degradation and fails to sufficiently leverage the long-term temporal information from adjacent frames with 3D convolution.

3. Method

Let $\mathbf{x}_1^T \equiv \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ be a sequence of old film frames, where T is the video length. Our target is to train a deep neural network to automatically restore spatial-temporal deterioration. To comprehensively mitigate the issues of old films, like illuminance flicker, physical structural pollution, or quality degeneration, we propose recurrent transformer networks (RTN), whose details are elaborated in Sec. 3.1. Since the training of RTN requires supervision, we subsequently introduce a data simulation method in Sec. 3.2, which could convert consecutive video frames acquired by modern cameras to corresponding degraded versions. In Sec. 3.3, we further extend the proposed framework to reference-based video colorization. Lastly, we introduce how to optimize RTN in Sec. 3.4.

3.1. Proposed Framework

Temporal Recurrent Network Most of the time, the frames of old films are plagued with severe quality degradation due to the improper storage environment of the film

material and the abrasion caused by mechanical protruding parts of old-fashioned projectors. More seriously, sometimes several frames are totally damaged where only fuzzy structures remain. Hence, it is necessary to leverage the long-term temporal clues to remove the occurred degradations and then render reasonable contents.

Towards this goal, we propose a recurrent architecture for the temporal modeling of old films because of the huge benefits of learned hidden representation and its powerful long-term propagation ability. Specifically, for timestamp t , we estimate the optical flows $f_{t-1 \rightarrow t}$ between two input frames \mathbf{x}_{t-1} and \mathbf{x}_t . Then the previous hidden state \mathbf{s}_{t-1} could be aligned to time t by warping function \mathcal{W} . Meanwhile, we leverage a convolutional encoder \mathbf{E} to project input frame \mathbf{x}_t into feature map $\mathbf{E}(\mathbf{x}_t)$ which shares the same spatial dimension with propagated state, then the newly restored state \mathbf{s}_t could be obtained as follows,

$$\mathbf{s}_t = \mathcal{R}_\uparrow \circ \mathcal{F}_\uparrow(\mathbf{E}(\mathbf{x}_t), \mathcal{W}(\mathbf{s}_{t-1}, f_{t-1 \rightarrow t})), \quad (1)$$

where \mathcal{F} aggregates the states between history and current feature, and \mathcal{R} further recovers and boosts the hidden representation, which will be described in the next section.

Due to the irregular exposure time of early movie cameras, the global brightness or colors may vary from one frame to the next, which is called the *flicker* phenomenon. The flicker is more challenging in our case because the content of old film itself is flickering whereas the flicker in typical video restoration tasks is mainly caused by frame-wise processing. Therefore, we need to aggregate the information from frames in a long period. To ensure that bidirectional information contributes equally, we also learn the backward hidden state $\tilde{\mathbf{s}}_t$ by considering the knowledge of future frames,

$$\tilde{\mathbf{s}}_t = \mathcal{R}_\downarrow \circ \mathcal{F}_\downarrow(\mathbf{E}(\mathbf{x}_t), \mathcal{W}(\tilde{\mathbf{s}}_{t+1}, f_{t+1 \rightarrow t})). \quad (2)$$

Combining the bidirectional states \mathbf{s}_t and $\tilde{\mathbf{s}}_t$, we reconstruct the final pixel-level output via a convolutional decoder.

$$\mathbf{y}_t = \mathbf{D}(\mathbf{s}_t \wedge \tilde{\mathbf{s}}_t) \in \mathbb{R}^{H \times W}, \quad (3)$$

where $\hat{\cdot}$ denotes the concatenation operation along the channel dimension. To this end, we have known how to effectively propagate the temporal information by an RNN, and the overall system pipeline is shown in Figure. 2.

Spatial Transformer Restoration After blending the hidden state and current feature, we need another spatial network to perform the frame-wise restoration. CNN has been a de-facto standard to most video processing tasks due to its efficient locality and excellent pattern restoration ability. However, the same locality property does not apply to old film restoration considering the following several aspects: 1) Old films contain many non-spatial homogeneous degradations like dust, dirt, or scratches, which require the network to explore long-range context information for plausible inpainting. 2) Since old films always deteriorate excessively, the estimated position-wise correspondence between frames may contain errors. If we still use CNN for spatial restoration, the consequent results are extremely unstable training, which has been verified in our experiments. In such cases, we need a mechanism to compensate for the inevitable estimation errors of optical flow.

To alleviate the above-mentioned issues, we propose to employ a powerful transformer network for spatial restoration. The problem is the quadratic computational cost of transformer architecture, which makes it difficult to model high-dimensional data. To efficiently process old films for high-resolution images, we adopt two modifications. Let's denote the fused representation derived from $\mathcal{F}(\cdot)$ as $\mathbf{h} \in \mathbb{R}^{H \times W \times C}$. First, we use strided convolutions to down-sample the spatial resolution by a factor of 2. Second, instead of directly modeling all the correlations among $\frac{HW}{4}$ tokens, we follow Swin Transformer [25] and employ the window-based attention and shifted strategy to improve the efficiency further. More specifically, in each layer of transformer, given one feature representation $\mathbf{z} \in \mathbb{R}^{\frac{H}{2} \times \frac{W}{2} \times C}$, we partition it into several non-overlapping $M \times M$ local windows. In each local window, we project the tokens into queries \mathbf{Q} , keys \mathbf{K} and values \mathbf{V} respectively through MLPs. Then the attention could be computed as follows,

$$\text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = f(\mathbf{Q}\mathbf{K}^\top / \sqrt{d} + \mathbf{B})\mathbf{V}, \quad (4)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{M^2 \times d}$, $f(\cdot)$ is the softmax operation and \mathbf{B} is the learnable relative positional embedding. Despite higher efficiency, the local attention would inevitably hurt the global modeling capability. To solve this problem, we also involve the cross-window connections with a cyclic shift. With stacking the consecutive transformer block, the model is able to perceive global context information for better degradation restoration and compensation of flow errors.

Learnable Guided Mask Intuitively, to restore the contaminants lying on frames, we could first localize the exact mask positions of such structured defects and then perform video inpainting. This scheme sounds reasonable, but ac-

tually infeasible here due to the following two aspects: 1) Similar to [37], we also tried to train an extra detection model to distinguish the structured degradations. Nonetheless, the generalization capability of this detection model is extremely deficient, especially for real-world data. The underlying reason is that, unlike single images, the diversities and variations of old film contaminants are larger. Although mixing some real-world paired data in training may alleviate this problem, the labeling cost is too large to afford, especially for video data. 2) Instead of losing all pixel information, by contrast, the dirt in one frame may be transparent, whose preserved contents may guide us for better restoration.

Motivated by that the contaminants mostly have large motion, shape and material variations across different frames, our solution is to learn a soft guided mask by contrasting the hidden state and current frame features in an unsupervised manner. More specifically, as shown in Figure. 3, given warped clean state $\mathbf{W} = \mathcal{W}(s_{t-1}, f_{t-1 \rightarrow t})$, which has aggregated all available information from frame 0 to frame $t - 1$, we employ a shallow network \mathcal{M} to regress the blending mask \mathbf{M} conditioned on current features $\mathbf{E}(\mathbf{x}_t)$,

$$\mathbf{M} = \mathcal{M}(\mathbf{E}(\mathbf{x}_t) \hat{\cdot} \mathbf{W}) \in \mathbb{R}^{H \times W \times 1}, \quad (5)$$

where $\mathcal{M}(\cdot)$ is composed of several convolutions and sigmoid function. With the learned soft mask, we aggregate the temporal priors and current frames by the following method,

$$\mathcal{F}(\mathbf{E}(\mathbf{x}_t), \mathbf{W}) = \mathbf{E}(\mathbf{x}_t) \cdot \mathbf{M} + \mathbf{W} \cdot (1 - \mathbf{M}). \quad (6)$$

3.2. Video Degradation Model

To train the RTN, designing a realistic video degradation model to generate paired data is essential. We achieve this target by involving the following degradation procedures.

Contaminant Blending To model the scratches, dirt or dust of old films, we collect 1k+ texture templates from the Internet, which are further augmented with random rotation, local cropping, contrast change, and morphological operations. Then we use *addition*, *subtract* and *multiply* blending modes with various levels of *opacity* $\in [0.6, 1.0]$ to combine the scratch textures with the natural frames.

Quality Degradation As mentioned before, most old films suffer from blurring, noise and unsharpness due to long-term storage and usage. Hence, we also downgrade the overall video qualities from these aspects. 1) Gaussian noise and speckle noise with $\sigma \in [5, 50]$. 2) Isotropic and anisotropic Gaussian blur kernels with covariance matrix $\Sigma = \mathbf{R} \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} \mathbf{R}^T$, where \mathbf{R} denotes the rotation matrix whose rotation angle $\theta \in [0, \pi]$, and $\sigma_1, \sigma_2 \in (0, 1)$ control the standard deviation of two principal axes respectively. 3) Random downsampling and upsampling with different interpolation methods. 4) JPEG compression whose level is in the range of [40, 100]. 5) Random color jitter through adjusting the *brightness* $\in [0.8, 1.2]$ and *contrast* $\in [0.9, 1.0]$.



Figure 4. Examples of rendered video frames with degradation.

Temporal Frames Rendering Unlike the degradation of a single photo, where we could inject these defects randomly, one observation of old films is a similar degradation pattern is shared across frames, which makes sense since most of the frames of one video potentially experience the same mechanical abrasion. Hence, we first define a set of template parameters for each video. Then we further apply predefined parameters with slight perturbations on consecutive temporal frames to achieve more realistic rendering. The synthetic data example is shown in Figure. 4.

3.3. Video Colorization

To make the restored old films more vivid, an effective solution is to colorize the videos further. Due to the ill-posed property of colorization, the automatic algorithms always struggle to generate appealing results [22]. A more practical scenario is, given one frame which has been colorized, how to propagate the known information to all the rest frames, meanwhile preserving the spatial-temporal coherence. Thus we further extend our RTN to the video colorization task under this setting, with only several minor modifications. First, we adjust the input color space from RGB to LAB, and only predict the chrominance channels given black-white frames. Second, to further prevent the color fading of long-term propagation, following [47] we search the correspondence between two frames via comparing the semantic similarity to get a coarse colorization result. Then the coarse AB channel will be concatenated with gray input and sent into the RTN. With the temporal propagation and spatial refinement capability of RTN, the reference-based colorization could achieve more stable performance.

3.4. Learning

The ultimate target of our method is to generate temporal-coherent, higher-quality and visual-pleasant videos given consecutive frames from old films. To accomplish these objectives, we apply the following losses.

L1 Loss The pixel-wise reconstruction loss \mathcal{L}_1 penalizes the difference between restored results and GT frames. Specifically, let \hat{y}_t denote target clean video frame and y_t denote the generated frame by our method at time t , we have

$$\mathcal{L}_1 = \frac{1}{T} \sum_{t=1}^T \|y_t - \hat{y}_t\|_1. \quad (7)$$

Perceptual Loss To further improve the visual quality of the restored frames, we also adopt the perceptual loss [15]:

$$\mathcal{L}_{perc} = \frac{1}{T} \sum_{t=1}^T \sum_{p \in P} \omega_p \|\Phi_p^{y_t} - \Phi_p^{\hat{y}_t}\|, \quad (8)$$

where $\Phi_p^{y_t}$ and $\Phi_p^{\hat{y}_t}$ denote the activation from the p^{th} layer of pretrained VGG19 network. P denotes the selected set (*relu2.2* to *relu5.2*) of layers to compute the perceptual loss, and ω_p controls the importance of different layers. The loss is accumulated over all frames in the generated video.

Spatial-Temporal Adversarial Loss Adversarial training, whose main principle behind it is to process a zero-sum game between a generator and a discriminator, has shown promising results in many tasks [10, 37]. For the old film restoration task, we adopt the Temporal-PatchGAN as suggested in [5] to enhance both perceptual quality and spatial-temporal coherence. Specifically, we employ a discriminator D composed of 3D convolutions to distinguish each spatial-temporal feature as real or fake by hinge loss,

$$\mathcal{L}_D = \mathbb{E}_{y \sim Y} [\text{ReLU}(1 - D(y))] + \mathbb{E}_{\hat{y} \sim \hat{Y}} [\text{ReLU}(1 + D(\hat{y}))], \quad (9)$$

where Y and \hat{Y} are the restored frames of RTN and clean video frames respectively. Then we use RTN to fool the discriminator by:

$$\mathcal{L}_G = -\mathbb{E}_{y \sim Y} [D(y)]. \quad (10)$$

Full Objective Combining all the above losses, the the overall objective we aim to optimize is

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_1 + \lambda_p \mathcal{L}_{perc} + \lambda_G \mathcal{L}_G. \quad (11)$$

We empirically set the weights for different losses as: $\lambda_1 = 1.0$, $\lambda_p = 1.0$ and $\lambda_G = 0.01$. For the video colorization task, we calculate the \mathcal{L}_1 loss in LAB color space and use a differentiable operator to convert LAB space to RGB space to compute the perceptual loss and adversarial loss.

4. Experiments

4.1. Implementation

We train the proposed RTN for 20 epochs using the ADAM optimizer [19] with $(\beta_1, \beta_2) = (0.9, 0.99)$. The learning rate is set to $2e-4$ for both generators and discriminators in the first 20 epochs, with linear decay to zero thereafter. We employ the off-the-shelf method RAFT [36] for flow estimation, whose parameters are fixed during the first 5 epochs, and then jointly optimized with other modules together using $lr = 2.5e-5$. During optimization, we randomly crop 256 patches from REDS [30] dataset and apply the proposed video degradation model on the fly. The batch size is set to 4 and the whole training takes ~ 2 days on 4 RTX 2080Ti GPUs.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$E_{warp}\downarrow$
Input	19.982	0.699	0.456	0.0167
Old Photo+TS [20,37]	21.962	0.768	0.315	0.0041
BasicVSR [4]	23.363	0.808	0.328	0.0053
Video Swin [26]	22.758	0.774	0.319	0.0061
DeepRemaster [14]	20.634	0.728	0.427	0.0066
DeOldify [1]	20.051	0.708	0.436	0.0149
Ours	24.465	0.840	0.192	0.0019
Ours w/o bi-direction	24.251	0.831	0.207	0.0036
Ours w/o soft mask	24.297	0.827	0.243	0.0025
Ours w/o transformer	24.342	0.830	0.229	0.0023

Table 1. Quantitative restoration comparisons on synthetic dataset. Our method achieves better performance on all metrics.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow
Input	27.100	0.945	0.189	110.559
DeOldify* [1]	26.271	0.937	0.149	59.686
DeepExemplar [47]	30.064	0.952	0.091	37.971
DeepRemaster [14]	29.253	0.950	0.127	40.385
Ours	32.838	0.977	0.065	31.992

Table 2. Quantitative colorization comparisons on REDS [30] dataset. DeOldify*: Non-reference based video colorization.

4.2. Setup

Baseline Methods We conduct comprehensive experiments by comparing RTN with the following approaches:

- *Old Photo Restoration* [37] + *TS* [20]: We adopt a state-of-the-art old photo restoration [37] algorithm for per-frame processing and generate coherent video through blind temporal smoothing [20].
- *BasicVSR* [4]: An advanced RNN-based method, tailored for high-quality video super-resolution.
- *Video Swin* [26]: A fully transformer-based architecture, which performs attention mechanisms in both spatial and temporal dimensions to capture global correlations.
- *DeepRemaster* [14] is a state-of-the-art old film restoration method using 3D convolutions, which also supports video colorization given one referenced frame.
- *DeOldify* [1]: An open-source tool for restoring old films.

For fair comparison, we re-train *BasicVSR* [4] and *Video Swin* [26] from scratch, and fine-tune *DeepRemaster* [14] using the same training data as ours. For *Old Photo Restoration* [37] and professional restoration tool *DeOldify* [1], we directly utilize their pre-trained models for inference.

Evaluation Metrics 1) We employ peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) to measure the low-level discrepancy between the restored output and the ground truth for synthetic data. 2) To better match the judgment of human perception, we also calculate the learned perceptual image patch similarity (LPIPS) [51]. 3) We involve the temporal warping error E_{warp} following [20] to measure the temporal consistency of the restored frames. 4) Since the restoration ground truth for old films

Method	NIQE \downarrow	BRISQUE \downarrow
Input	18.9907	53.6776
Old Photo+TS [20,37]	17.5110	48.1470
BasicVSR [4]	17.6842	62.7381
Video Swin [26]	18.9462	52.4758
DeepRemaster [14]	17.9697	49.9638
DeOldify [1]	17.9062	51.2813
Ours	15.4254	42.1422

Table 3. Quantitative restoration comparisons on real old films.

is unavailable, two non-reference frame quality assessment metrics NIQE [29] and BRISQUE [28] are used to evaluate the algorithm performance on real-world data. 5) For the video colorization task, we also adopt Fréchet Inception Distance (FID) [12] to measure the semantic discrepancy between the colorized output and the natural colorful frames.

4.3. Results

Quantitative Comparisons We report quantitative results on both synthetic dataset and real-world old films. We create the synthetic data by blending random degradations with clean frames of DAVIS [31] dataset, which contains large camera motions and scene diversities. As shown in Table. 1, among these baselines, BasicVSR [4] obtains good PSNR and SSIM performances. However, because video SR mainly considers unstructured degradation, the visual restored results become over-smooth while directly training with various defects, which is also indicated by the descent of LPIPS. Moreover, although per-frame old photo restoration [37] followed by temporal smoothing [20] pipeline has acceptable results on LPIPS and E_{warp} , the PSNR and SSIM inevitably decrease without leveraging the temporal clues well. By contrast, our method outperforms the baselines on all metrics. The same conclusion still holds on real-world experiments, where we collect 63 old films from the internet for evaluation. In Table. 3, the NIQE and BRISQUE results of our method are significantly better than others, which demonstrate the powerful ability of RTN to restore old films.

Qualitative Comparisons We further conduct qualitative comparisons in Figure. 5. As we can see, the old photo restoration method [37] only resolves a small proportion of the structured degradations since this method is designed for image restoration and does not leverage temporal information. Besides, DeOldify [1] could remove some noise but leave many other degradations unsolved. VideoSwin [26] recovers some textures through exploring the spatial-temporal correlations but struggles to remove the existing contaminants well. DeepRemaster [14], as the state-of-the-art old film restoration method, is capable of handling slight dust and scratches. Nonetheless, their method could not render reasonable contents while meeting severe abrasion like the first and the last columns of Figure. 5, meanwhile failing to produce sharp textures given blurry frames. With the elaborately designed modules, the restored videos by our RTN



Figure 5. Qualitative restoration comparisons on real-world old films. Our method could handle complicated degradations of old films.

effectively overcome these issues. Please refer to the supplementary material and video demo for more comparisons.

Colorization Comparisons In this section, we show the video colorization comparisons with other baselines. The gray version of the subset of REDS [30] is adopted as the test set. For each video, we predict the colors of the first

50 frames by taking the 100th frame as the colorization reference. We choose two state-of-the-art reference-based methods DeepExemplar [47] and DeepRemaster [14], and one old film colorization tool DeOldify [1] as baselines. Quantitatively, although DeepRemaster [14] leverage the reference image, perceptual metrics like LPIPS and FID in




Figure 6. Qualitative video colorization comparisons. [1]*: Colorization without reference frame.

Table 2 are not satisfactory. DeepExemplar demonstrates better performance compared with others, but color bleeding frequently appears like in Figure 6, which hurts the colorization quality a lot. Although our method is not specifically designed for video colorization, the combination of spatial transformer and temporal RNN solves this task well.

4.4. Ablation Study

Temporal Bi-directional RNN We first consider changing the bi-directional RNNs into a unidirectional setting. To ensure the fairness of this setting, we further increase the parameter number of other networks to match the bi-directional setting. Even so, the temporal consistency still vastly downgrades as shown in Table 1, which demonstrates the importance of perceiving both past and future temporal information while restoring the flicker artifact of old films.

Learnable Guided Mask Next, we replace the learnable soft mask with a direct channel-wise concatenation operation for the aggregation of the previous hidden state and the current frame feature to show its significance. Without such crucial spatial clues, the network could not distinguish the original content and contaminant artifacts well, as shown in Figure 8, thereby failing to restore the structured degradations, which is consistent with the performance drop on the SSIM metric in Table 1. Moreover, we also visualize the learned soft mask of real-world old films in Figure 7. Although this structured degradation is never seen in the training and occupies a large portion, the predicted mask is very accurate and effective to advance restoration.

Spatial Transformer Last but not least, we employ a CNN-based architecture instead of transformer layers to perform the spatial restoration. In this setting, the training becomes very unstable due to the underlying error of flow prediction and limited receptive field of convolutions. Therefore, we pick the best checkpoint ahead of the occurrence of gradient explosion for evaluation. Compared with the baseline, which uses a unidirectional network and transformer, although the low-level metrics like PSNR and SSIM maintain similar performance, the perceptual score, *i.e.* LPIPS, drops a lot due to defective spatial restoration. Besides, we also show its qualitative result in Figure 8, where the scratches are not



Figure 7. Visualization of the learned guidance mask. The soft mask could effectively help resolve the structured degradations.




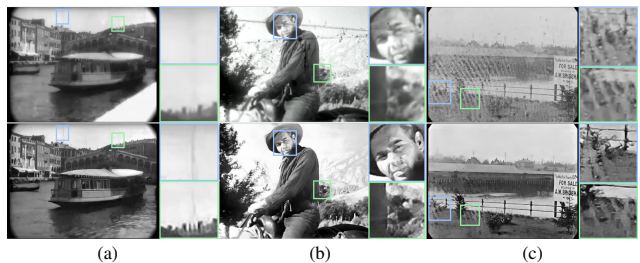
Figure 8. Visual results of each ablation study.

handled well without the help of the spatial transformer even though this baseline is contaminant-aware.

5. Conclusion and Limitation

In this paper, we present a recurrent transformer network to solve the mixed degradations of old films by leveraging the temporal modeling of recurrent neural network and the spatial modeling of transformers. Extensive visual comparisons and quantitative evaluation demonstrate that our approach performs well on both synthetic data and real old films. We also extend the RTN to achieve better reference-based video colorization compared with prior baselines.

However, there are still some limitations waiting to be resolved: a) The model may fail to distinguish the contaminant from the frame content due to their ambiguity like the black lines, which are amplified since they are wrongly recognized as the smoke of the scene; b) GAN may synthesize inadequate high-frequency details and artifacts; c) It is challenging to restore severely degraded frames with barely recognizable content. We will explore these challenges in the future work.



Acknowledgements: We would like to thank anonymous reviewers for their constructive comments. This work was supported by the Hong Kong Research Grants Council (RGC) Early Career Scheme under Grant 9048148 (CityU 21209119).

References

- [1] Deoldify. <https://github.com/jantic/DeOldify>. 6, 7, 8
- [2] S Derin Babacan, Rafael Molina, and Aggelos K Katsaggelos. Total variation super resolution using a variational approach. In *2008 15th IEEE International Conference on Image Processing*, pages 641–644. IEEE, 2008. 2
- [3] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 60–65. IEEE, 2005. 2
- [4] 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. 2, 6
- [5] Ya-Liang Chang, Zhe Yu Liu, Kuan-Ying Lee, and Winston Hsu. Free-form video inpainting with 3d gated convolution and temporal patchgan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9066–9075, 2019. 5
- [6] Michele Claus and Jan van Gemert. Videnn: Deep blind video denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 2
- [7] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006. 2
- [8] Chen Gao, Ayush Saraf, Jia-Bin Huang, and Johannes Kopf. Flow-edge guided video completion. In *European Conference on Computer Vision*, pages 713–729. Springer, 2020. 2
- [9] Ioannis Giakoumis, Nikos Nikolaidis, and Ioannis Pitas. Digital image processing techniques for the detection and removal of cracks in digitized paintings. *IEEE Transactions on Image Processing*, 15(1):178–188, 2005. 3
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 5
- [11] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Recurrent back-projection network for video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3897–3906, 2019. 2
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. 6
- [13] Zhang Hongying, Wu Yadong, and Kuang Zhonglin. An efficient scratches detection and inpainting algorithm for old film restoration. In *2009 International Conference on Information Technology and Computer Science*, volume 1, pages 75–78. IEEE, 2009. 3
- [14] Satoshi Iizuka and Edgar Simo-Serra. Deepremaster: temporal source-reference attention networks for comprehensive video enhancement. *ACM Transactions on Graphics (TOG)*, 38(6):1–13, 2019. 2, 3, 6, 7, 8
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 5
- [16] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Deep video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5792–5801, 2019. 2
- [17] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016. 2
- [18] Seong-Whan Kim and Ki-Hong Ko. Efficient optimization of inpainting scheme and line scratch detection for old film restoration. In *Pacific Rim International Conference on Artificial Intelligence*, pages 623–631. Springer, 2006. 3
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [20] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning blind video temporal consistency. In *Proceedings of the European conference on computer vision (ECCV)*, pages 170–185, 2018. 6, 7
- [21] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 2
- [22] Chenyang Lei and Qifeng Chen. Fully automatic video colorization with self-regularization and diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3753–3761, 2019. 5
- [23] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1833–1844, 2021. 2
- [24] Hongyu Liu, Ziyu Wan, Wei Huang, Yibing Song, Xintong Han, and Jing Liao. Pd-gan: Probabilistic diverse gan for image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9371–9381, 2021. 2
- [25] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10012–10022, October 2021. 2, 4
- [26] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 6, 7
- [27] Xiaojiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder

- networks with symmetric skip connections. *Advances in neural information processing systems*, 29:2802–2810, 2016. [2](#)
- [28] Anish Mittal, Anush K Moorthy, and Alan C Bovik. Blind/referenceless image spatial quality evaluator. In *2011 conference record of the forty fifth asilomar conference on signals, systems and computers (ASILOMAR)*, pages 723–727. IEEE, 2011. [6](#)
- [29] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. [6](#)
- [30] Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *CVPR Workshops*, June 2019. [5](#), [6](#), [7](#)
- [31] 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 preprint arXiv:1704.00675*, 2017. [6](#)
- [32] Takahiro Saito, Takashi Komatsu, Tomohisa Hoshi, and Toshiaki Ohuchi. Image processing for restoration of old film sequences. In *Proceedings 10th International Conference on Image Analysis and Processing*, pages 709–714. IEEE, 1999. [3](#)
- [33] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1279–1288, 2017. [2](#)
- [34] Masanori Suganuma, Xing Liu, and Takayuki Okatani. Attention-based adaptive selection of operations for image restoration in the presence of unknown combined distortions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9039–9048, 2019. [2](#)
- [35] Matias Tassano, Julie Delon, and Thomas Veit. Fastvdnet: Towards real-time deep video denoising without flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1354–1363, 2020. [2](#)
- [36] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020. [5](#)
- [37] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Bringing old photos back to life. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2747–2757, 2020. [2](#), [4](#), [5](#), [6](#), [7](#)
- [38] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Old photo restoration via deep latent space translation. *arXiv preprint arXiv:2009.07047*, 2020. [2](#)
- [39] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4692–4701, October 2021. [2](#)
- [40] 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. [2](#)
- [41] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1905–1914, 2021. [1](#), [2](#)
- [42] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. [2](#)
- [43] Yair Weiss and William T Freeman. What makes a good model of natural images? In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [2](#)
- [44] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2019. [2](#)
- [45] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4471–4480, 2019. [2](#)
- [46] Ke Yu, Chao Dong, Liang Lin, and Chen Change Loy. Crafting a toolchain for image restoration by deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2443–2452, 2018. [2](#)
- [47] Bo Zhang, Mingming He, Jing Liao, Pedro V Sander, Lu Yuan, Amine Bermak, and Dong Chen. Deep exemplar-based video colorization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8052–8061, 2019. [5](#), [6](#), [7](#), [8](#)
- [48] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. *arXiv preprint arXiv:2103.14006*, 2021. [1](#), [2](#)
- [49] 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, 2017. [2](#)
- [50] 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, 2018. [2](#)
- [51] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [6](#)
- [52] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018. [2](#)
- [53] Shangchen Zhou, Jiawei Zhang, Jinshan Pan, Haozhe Xie, Wangmeng Zuo, and Jimmy Ren. Spatio-temporal filter adaptive network for video deblurring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2482–2491, 2019. [2](#)