

Task Agnostic Restoration of Natural Video Dynamics

Muhammad Kashif Ali Dongjin Kim Tae Hyun Kim*
Department of Computer Science, Hanyang University, Seoul, Korea
{kashifali, dongjinkim, taehyunkim}@hanyang.ac.kr

Abstract

In many video restoration/translation tasks, image processing operations are naïvely extended to the video domain by processing each frame independently, disregarding the temporal connection of the video frames. This disregard for the temporal connection often leads to severe temporal inconsistencies. State-Of-The-Art (SOTA) techniques that address these inconsistencies rely on the availability of unprocessed videos to implicitly siphon and utilize consistent video dynamics to restore the temporal consistency of frame-wise processed videos which often jeopardizes the translation effect. We propose a general framework for this task that learns to infer and utilize consistent motion dynamics from inconsistent videos to mitigate the temporal flicker while preserving the perceptual quality for both the temporally neighboring and relatively distant frames without requiring the raw videos at test time. The proposed framework produces SOTA results on two benchmark datasets, DAVIS and videvo.net, processed by numerous image processing applications. The code and the trained models are available at <https://github.com/MKashifAli/TARONVD>.

1. Introduction

Video sharing social media platforms like Snapchat and TikTok have introduced the common populace to a plethora of computer vision applications such as Style Transfer [9], Colorization [43], Denoising [42], and Dehazing [38]. With this wide-scale integration of these classical computer vision applications in such platforms, various image processing operations are naïvely extended to videos due to scarcity of annotated video datasets and their computational complexity of video processing methodologies. This naïve extension of static image processing methodologies to videos disregards the temporal connection of the video progression and introduces severe temporal flickering in the videos. This temporal flicker can appear for various reasons; for in-

stance, these image processing methods can produce drastically different results for temporally neighboring frames due to slight changes in their global or local content statistics, or it could also happen due to the multi-modality of the application as there could exist a number of valid solutions for images with similar content as highlighted in [3, 41]. Therefore, the extension of these image-to-image translation tasks to the video domain is an active area of research in computer vision, and methods that can help extend various image processing applications with little to no knowledge of the operations used to create videos are quite useful.

Currently, there are several task-dependent temporal consistency correction approaches available such as [4, 23, 34], but due to the complex nature of these applications, only a handful of approaches have been proposed to tackle the problem of blind temporal consistency correction. Bonneel et al. [3] initiated the investigation of blind temporal consistency correction with gradient-domain minimization of per-frame processed video with the unprocessed video to minimize the warping error between the frames. Lai et al. [16] extended the formulation of [3] with the help of recurrent Convolutional Neural Networks (CNN) and introduced a perceptual penalty in their formulation to restrict the deviation of perceptual content of the restored video from the frame-wise processed video. Deep Video Prior (DVP) [19] extended Deep Image Prior (DIP) [35] to the temporal dimension and proposed to formulate enforcing temporal consistency by training a CNN to generate processed video from the unprocessed video without utilizing optical flow. All of these approaches rely on the availability of unprocessed videos for implicit extraction of consistent motion dynamics to use as a restoration guide.

Please note that all the previously proposed approaches for this task define it with the help of unprocessed videos (both task-specific and task-agnostic). Having access to unprocessed videos at test time helps the models in siphoning and transferring consistent motion dynamics to the inconsistent videos to decrease temporal flicker. Despite its efficacy, this implicit definition limits the applicability of the previously proposed approaches to only the videos for which their unprocessed counterparts are available and also

*: Corresponding author.

Key features	Bonneel et al. [3]	Lai et al. [16]	DVP [19]	Ours
Does not require raw videos	×	×	×	✓
Explicit video dynamics design	✓	×	×	✓
Faithful perceptual restoration	×	×	×	✓
Sequential processing	✓	✓	×	✓
Long term consistency	×	×	✓	✓
Sharpness in processed videos	×	✓	×	✓
High-frequency flicker removal	✓	×	×	✓
Test time training	×	×	✓	×
Runtime (per frame avg.)	2.4635*	0.2146	3.6365	0.2236

Table 1: **Key differences in available approaches.** “*” in front of Bonneel et al. [3]’s runtime highlights environment difference and is only presented for reference only. Please note that the runtime is averaged over the same video sequence for all the methods.

introduce an inherent bias towards the unprocessed videos which can compromise the translation effect and quality of the processed videos (as presented in Fig. 8 ~ 10).

In order to overcome these limitation, we aim to learn and infer consistent motion representation solely from temporally inconsistent videos. Doing so does not only alleviate the need for raw videos at test time, but also mitigates the inherent bias towards the raw videos that is commonly present in the currently available approaches for this task. For the task of learning these consistent motion representation, we fine-tune conventional optical flow estimations networks like [11, 31, 33] along with the restorative part of the network in an end-to-end manner. The proposed model pipeline is illustrated in Fig. 1. The proposed model achieves state-of-the-art qualitative and quantitative results. The proposed formulation can also handle the resolution mismatch problem in processed and raw videos and makes it possible to extend the Single Image Super-Resolution [17] methods to Video Super-Resolution methods without any modification. The detailed description of our formulation is presented in 3. We summarize our contributions as follows:

- Identify and propose tailored solutions for various challenges of naïve extensions of image processing applications to videos.
- Learn to infer consistent motion representations from temporally inconsistent videos.
- Utilize the learned consistent motion representations to propose a general framework for task agnostic temporal consistency correction that produce SOTA results.

2. Related Work

The literature is divided into two main streams to generate visually appealing videos through image transformation models. The first stream tackles the task of mitigating temporal inconsistencies with the reformulation of the task at hand. Whereas, The second stream focuses on developing post-processing models that refine the frame-wise processed videos by penalizing the temporal deviation of the processed video from its unprocessed counterpart. The

second stream is further divided into two sub-streams: task-specific and task-agnostic. The following passages summarize the details of both of the above-mentioned streams, respectively.

Reformulation stream (video-to-video translation):

These approaches are generally termed video-to-video translation applications. Generally, these approaches either consider multiple frames as input and produce multiple output frames or generate a single frame from multiple input frames in the form of frame recurrence such as [5, 23]. There are also cases like video style transfer where content information is propagated to the next time step with the help of optical flow to initiate the optimization of the next frame. This frame-recurrent methodology has also been proven effective in applications like video super resolution [29]. Designing these approaches for each task and training them from scratch is a taxing task and data scarcity can make these formulations unfeasible. Generally, such models often do not adapt well to different tasks; therefore, methods that can restore the temporal consistency of multiple frame-wise processed videos are being actively investigated.

Post-processing methods: Bonneel et al. [3] initiated the investigation of task agnostic temporal consistency correction using a gradient-based optimization strategy in which temporally consistent (unprocessed) videos were used as restoration guides for correcting the temporal inconsistency of frame-wise processed videos. Their formulation motivated various task-dependent and task-agnostic, and temporal consistency correction approaches with slight variations of their original objective formulation. Despite the simplicity of their formulation, there only exist a handful of approaches that address task-agnostic temporal consistency correction due to the difficult nature of the task. Yao et al. [40] proposed a key-frame strategy that accounted for the motion of different objects in those key-frames to handle occlusion as well as temporal consistency. Lai et al.[16] proposed the first deep learning approach for this task by employing ConvLSTM [30] and a perceptual loss [13]. Lei et al. [19] proposed an extension of Deep Image Prior [35] and demonstrated its capability to mitigate the temporal flicker of per-frame processed videos.

Apart from these general frameworks, various task-specific approaches have also been proposed, such as [5, 23, 34]. Defining these task-specific approaches is relatively straightforward. These approaches define a temporal extension around a backbone method and penalize the deviation of generated content with the help of motion information from the unprocessed videos.

In this work, we propose a general formulation for this task that alleviates the requirement of unprocessed videos at test time by learning to infer a consistent motion representation from the temporally inconsistent videos. This inferred

motion representation is then processed with the content of temporally inconsistent videos in a two branched network that restores the natural dynamics of per-frame processed videos. Our proposed network also contains a recurrent bottleneck module that helps consistent propagation of content throughout the video sequence and produces SOTA qualitative and quantitative results by only taking into account the temporally inconsistent videos at test time.

In order to aid the readers, we provide the following table to highlight the key differences between the currently available approaches for this task.

3. Proposed Method

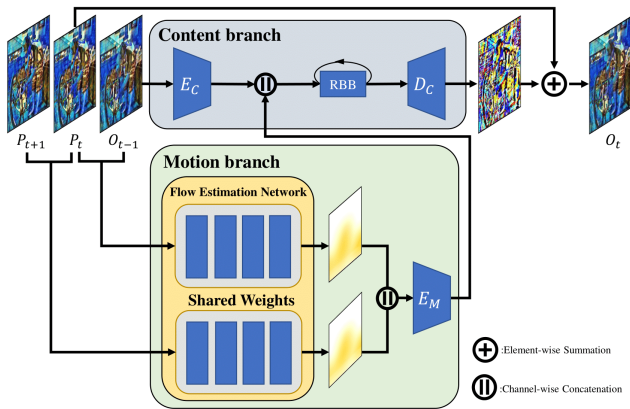


Figure 1: **Model Architecture.** An illustration of the proposed model architecture. The proposed model consists of a two-branch architecture that separately processes motion and content. The combined motion and content features are concatenated along the channel dimension and are then passed through a Recurrent Bottleneck Block (RBB) and a decoder to generate the restored frame.

In this section, we describe the details of the proposed method. Consider a raw (temporally consistent) video $\{I_0, I_1, I_2, \dots, I_n\}$ with n frames and its frame-wise processed (temporally inconsistent) video $\{P_0, P_1, P_2, \dots, P_n\}$ acquired by an image processing function h as $\{h(I_0), h(I_1), h(I_2), \dots, h(I_n)\}$. The goal of this task is to restore the temporal consistency of the temporally inconsistent video and to produce a temporally consistent version $\{O_0, O_1, O_2, \dots, O_n\}$. For the sake of simplicity, these notations are chosen to be in line with the notations used in past literature [3, 16, 19].

The proposed formulation draws inspiration from the disentanglement of a video into its base components: content and motion representations as presented in [22]. In their work [22], they investigated the controlled generation of short length videos by mixing motion and con-




Figure 2: **Temporal flicker in optical flow.** An illustration of the decomposition of optical flow generated from temporally inconsistent frames to consistent optical flow and inconsistent flicker.

tent components from various videos. In this work, we explore a similar concept of controlled disentanglement and re-entanglement strategy to restore natural video dynamics of the frame-wise processed videos. More specifically, we argue that the previously proposed approaches for this tasks implicitly employ this disentanglement and re-entanglement in their formulation with the help of unprocessed video frames. Conventionally, optical flow is used as an inter-frame motion representations. The estimated optical flow can also be used to approximate the content of future frames with the help of past frames provided that there does not exist occlusion between the frames (as presented in Eq. 1).

$$\hat{I}_{t+1} \approx w(I_t, of_{t+1 \Rightarrow t}). \quad (1)$$

This interdependence of optical flow and content can lead to inconsistent motion representation as optical flow estimation networks are not trained explicitly to distinguish temporally consistent content from temporal flicker. Due to this limitation, the evaluated optical flow estimated from temporally inconsistent frames is often compromised with temporal flicker (as presented in Fig. 2). This compromise of the motion information makes it challenging to separate the consistent component of optical flow from the additive flicker. Therefore, the currently available techniques for this task often rely on siphoning implicit motion representations from the raw videos and often suffer in cases where the content of both processed and raw videos deviate significantly. Another problem with employing these implicit representations is the bias introduced due to the lack of control in mixing these implicit representations to restore videos which often leads to the mitigation of translation effect. Therefore, if a mechanism can be developed that produces similar motion representations from both unprocessed and frame-wise processed videos, better and unbiased restorative effects than those produced by the previously proposed approaches can be achieved without relying on the availability of unprocessed videos.

Please note that another observation that can be made from Fig. 2 (right most figure) is that a significant amount of flicker in frame-wise processed videos, generally occurs near the motion boundaries. This observation helps us in defining objective functions capable of localizing and targeting temporal deviations as presented in Sec. 4.1.

To overcome these problems and to learn and utilize consistent motion representation, we integrate an optical flow estimation network into the restoration network as illustrated in Fig. 1. A similar integration of a small-scale optical flow estimation network (SpyNet [26]) for task-oriented flow estimation is presented in [37] for the task video denoising for its extremely small size. We utilize a relatively larger network (PWC-Net [31]) in our method and provide an ablation study on different sized optical flow models in the accompanying supplementary text. We further extend their idea to a two pass mechanism with an added encoder to encourage generation of consistent dynamics. The optical flow network in the proposed methodology is finetuned in an end-to-end manner without any special supervision.

The proposed model takes in three frames O_{t-1} (P_{t-1} for the first time step), P_t and P_{t+1} as input. These three frames are propagated through both the content and motion branches. The proposed model consists of a UNet [27] like structure with multiple encoder streams and a single decoder stream. The decoder contains skip connections from the content stream to encourage better reconstruction. The bottleneck part of the model contains a recurrent bottleneck block to transmit and manipulate the information of the generated frames to temporally distant frames. Please note that the multi-branch encoder(s)-bottleneck-decoder architecture used for the proposed model is employed due to its proven efficacy in video generation, translation and restoration tasks as described in [4, 8, 14, 15, 16, 23, 32, 34]. The main difference in the proposed architecture comes from the design of our motion branch which allows us to present a standalone solution for this task. The motion branch comprises of two passes through a PWCNet [31] with shared weights followed by a conventional encoder-like architecture. An illustration of the proposed model is presented in Fig. 1. This two-pass strategy in the motion branch along with the attached encoder allows the proposed model in regressing consistent motion dynamics from temporally inconsistent frames. Further architectural details of the proposed network are presented in the accompanying supplementary. Figure 3 presents conventional optical flow along with the learned motion representation generated by the integrated finetuned PWC-Net [31] from temporally consistent (left) and inconsistent frames (right). Please note that the "modified" flow can differ starkly from the conventional flow and can only be interpreted effectively by the network it is trained with as described in [37]. Contrasting to the conventional optical flow modified flow estimated from both consistent and inconsistent videos is quite similar, which suggests that the network can effectively utilize the inferred consistent motion representation to restore natural temporal dynamics in the per-frame videos by only taking into account the temporally inconsistent videos at test time. Another benefit of this disentangled representation




Figure 3: **Conventional and Modified flow like representation.** Comparison of original and Modified flow evaluated from temporally consistent (left) and inconsistent (right) frames. The modified flow evaluated from both the temporally consistent and inconsistent frames is quite similar.

is that it allows the users to further improve the level of restored temporal consistency with the help of an iterative arrangement; please refer to the supplemental for the details of the iterative arrangement and the results generated with such arrangement. The objective functions used for the end-to-end training of the proposed network are described in detail in the next section.

4. Objective Functions

This section provides the details of the learning objectives used in the training phase of the proposed network. These objectives can be classified into two categories, local neighborhood and long-term objectives. Both of these optimization categories are presented below.

4.1. Local Neighborhood Losses

In this work, we aim to find a general solution for correcting the temporal consistency of frame-wise processed videos. Due to the diversity of the applications addressed, a simplistic spatial content matching reconstruction loss cannot justify the generation of temporally consistent frames, as the content can vary significantly in applications like style transfer [9] and image adjustment [38]. This discrepancy in the image space of processed and raw videos hinders a straightforward definition of a feasible loss function for this task. Due to the challenging nature of this task, we propose a flow gradient loss that provides a supervision signal for the reconstruction of motion boundaries (where most of the inconsistencies/flicker appear; as highlighted in Fig. 2) with the help of optical flow acquired through the unprocessed videos. We further extend this loss to only compare the gradients of the optical flow acquired through the synthesized and raw frames. Doing so reduces the redundancy present in the high dimensional optical flow which contains a redundant global component for this task which is only useful in applications like video stabilization [1] as described in [39]. The proposed flow gradient loss encapsulates the local spatio-temporal information necessary

to mitigate temporal flicker near the motion boundaries by removing the unnecessary high-dimensional components from the optical flow. Similar information bottle-necking techniques have been discussed thoroughly in vision based understanding and manipulation literature [7, 24]. The proposed spatio-temporal loss is defined as follows:

$$\mathcal{L}_{fg} = \sum_{t=2}^T \|\nabla(of(O_t, O_{t-1})), \nabla(of(I_t, I_{t-1}))\|_1. \quad (2)$$

Here, \mathcal{L}_{fg} and of represents flow gradient loss and the optical flow estimation network (FlowNet2.0 [11]) respectively. Please note that any optical flow estimation network (SpyNet [26], PWC-Net [31], RAFT [33], etc.) can be used in both the motion branch and the definition of the loss functions, we choose FlowNet2.0 [11] for its relatively lower mean end-point-error, as the estimated optical flow can vary significantly for small scale networks (comparison provided in the supplementary text). In equation 2, the ∇ operator denotes the spatial gradient of the estimated optical flow. A similar loss in optical flow domain has been proposed in [21] where the raw optical flow of the generated video was compared with that of the ground truth video. Although effective for the task of video inpainting, the raw optical flow matching approach fails in applications like style transfer in which the entire content of the generated video and the raw video differ greatly, i.e., contains excessive spatio-temporal variations which contribute to optical flow estimation. In our experiments, the models trained with raw optical flow performed poorly, whereas the model trained with flow gradients converged in a relatively shorter time and produced sharper frames.

Please note that this optical flow-based loss by itself is not sufficient for faithfully correcting the temporal consistency, as there could exist multiple solutions for optical flow equation, e.g., it does not take into account the appearance of the synthesized and raw frames. In order to address this issue and to guide the network for the generation of perceptually credible frames, a non-local optical flow based loss in the image space is employed as presented below:

$$\mathcal{L}_{recon} = \sum_{t=2}^T C_t \|O_t - \hat{O}_t\|_1, \quad (3)$$

This warping based reconstruction loss function has been adopted from the [4], as many variants of this objective function with slight modifications have been proven effective for enforcing temporal smoothness in video translation tasks in [3, 5, 16, 23, 34]. Here T represents the total number of frames in a sequence and \hat{O}_t represents the warped frame O_{t-1} (as described in Eq. 1). The working of this loss function can be summarized as the propagation of content from preceding frames to the subsequent frames in image domain to ensure that the generated frames contain similar

content to the previous frames and the inter-frame motion dynamics of the synthesized frames are similar to that of the temporally consistent counterpart. This non-local optical flow based loss function also takes into account the occlusion problem of the Eq. 1 by masking the occluded and de-occluded content as defined below:

$$C_t = \exp\left(-\alpha \|I_t - \hat{I}_t\|_2^2\right). \quad (4)$$

The value for $\alpha = 50$ is chosen according to the previous works where a similar strategy is used for evaluating occlusion [34, 28].

An additional conventional short-term perceptual similarity loss [13] was also introduced in the training phase to minimize the deviation of the synthesized frames from the original processed frames. This loss is defined as follows:

$$\mathcal{L}_p = \sum_t \sum_l \|\phi_l(O_t) - \phi_l(P_t)\|_1, \quad (5)$$

Here $\phi_l(\cdot)$ represents layers of a VGG-16 network till the layer *relu_4_3* (trained on the ImageNet dataset [6]).

4.2. Temporal Constancy Loss

The loss defined in Eq. 3 adequately defines the content propagation from the previous frame to the current frame but lacks the ability to enforce it in temporally distant instances. For instance, consider the task of colorization where a car appears in a portion of the frames and is assigned some random color, and it disappears and re-appears in a later instance. The frame-wise colorization method can assign it a different color in each sequence, and the loss function defined in Eq. 3 will be sufficient to enforce the color constancy of the car in each interval where the car is visible. This, however, does not ensure the constancy of the color assigned to the car in both intervals, i.e., the car could be assigned a solid blue color in the first interval and a solid red color in the next interval which is quite prominent in methods like [16, 34]. To address this problem of temporally distant yet similar instances, we introduce a recurrent module (ConvLSTM [30]) in our proposed model and add explicit constraints that force the model to generate temporally consistent frames. The previously proposed approaches for this task simply try to minimize the deviation of current frames from the first frame of the sequence, which is inadequate for videos containing largely varying frames. This variation of the content can render this loss ineffective, and the inconsistencies can only be weakly penalized with the help of motion dynamics siphoned from the raw videos. Therefore, we introduce a bi-directional penalty that enforces the temporal smoothness in both backward (through the local neighborhood losses) and forward (through the proposed long-term loss) direction. This bi-directional optimization strategy is inspired from the ping-

pong loss dynamic as presented in [5]. The proposed long-term loss is termed as constancy loss and is presented below:

$$\mathcal{L}_{constancy} = \sum_{p=1}^{T-2} \sum_{t=p+2}^T C_{t \Rightarrow p} \|O_t - w(O_p, of(I_{t-1}, I_p))\|_1. \quad (6)$$

Here subscript p highlights that the preceding distant images are compared with all the subsequent frames. Unlike the simplistic long-term temporal constraints used in [16, 34] which only penalized the deviation of content from the first frame, this loss term takes into account the inter-frame flicker of all the frames in a sequence and provides a strong supervision signal capable of penalizing temporal discrepancies throughout the sequence. Due to this stronger temporal penalty accompanied with the proposed flow gradient loss, the proposed network learns to produce smoother videos despite having access to only the frame-wise processed videos at test time. The final loss for each training sequence is given by,

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_{fg} + \lambda_2 \mathcal{L}_{recon} + \lambda_3 \mathcal{L}_p + \lambda_4 \mathcal{L}_{constancy}. \quad (7)$$

The $\lambda(s)$ in the equation above define the contribution of each loss in the optimization phase. The details of these hyperparameters are provided in the supplementary text.

5. Experiments and Results

During our experimentation phase, we tested various optical flow estimation networks such as [11, 26, 31]. Flownet2.0 [11] has approximately 162M trainable parameters, and it produced relatively better results in lesser training iterations as compared to the other two variants. Having a vast number of parameters compromises both the training and testing time. Therefore, we opted for a medium-sized optical flow estimation network (PWC-Net [31]) and a higher number of optimization iterations. The results produced by both variants are discussed in the supplementary text. We also experimented with an iterative arrangement of the trained models where restored videos are again subjected to temporal consistency correction models and evaluated that the proposed model architecture consistently reduces warp error with every iteration (as presented in Fig. 4). Generally, the frames subjected to a high number of consistency correction iterations lose fine details. The generated results by various restoration iterations are provided in the supplemental. Please note that the comparative results provided in this paper are generated through a single restoration iteration.

The dataset for training the proposed model contains videos from the DAVIS dataset for video segmentation [25] and videos gathered from Videovo.net by [16]. The training dataset contains videos processed through a diverse range




Figure 4: **Temporal warp error vs. iterations.** Temporal warp error consistently decreases with a higher number of restoration iterations.

of applications such as, Artistic Style Transfer [9, 13], Colorization [36, 44], Image Enhancement [10], Intrinsic Image Decomposition [2], and Image-to-Image Translation tasks [20, 44, 46]. The qualitative and quantitative results are described below.

5.1. Quantitative Results

Task	DAVIS				VIDEVO			
	Bonneel[3]	Lai[16]	Ours	DVP[19]*	Bonneel[3]	Lai[16]	Ours	DVP[19]*
WCT [20]/antimono	0.0029	0.0031	0.0026	0.0022	0.0015	0.0021	0.0015	0.0017
WCT [20]/asheville	0.0047	0.0059	0.0047	-	0.0032	0.0043	0.0039	-
WCT [20]/candy	0.0034	0.0047	0.0035	-	0.0020	0.0032	0.0025	-
WCT [20]/feathers	0.0040	0.0040	0.0027	-	0.0027	0.0030	0.0021	-
WCT [20]/sketch	0.0036	0.0029	0.0021	-	0.0025	0.0022	0.0017	-
WCT [20]/wave	0.0033	0.0035	0.0027	0.0026	0.0023	0.0026	0.0021	0.0019
Fast-neural-style [13]/princess	0.0043	0.0063	0.0042	0.0038	0.0035	0.0053	0.0042	0.0046
Fast-neural-style [13]/udnie	0.0021	0.0023	0.0022	0.0022	0.0012	0.0014	0.0014	0.0022
DBL [10]/expertA	0.0017	0.0011	0.0010	0.0015	0.0014	0.0007	0.0006	0.0015
DBL [10]/expertB	0.0016	0.0010	0.0008	0.0014	0.0011	0.0006	0.0004	0.0014
IID [2]/reflectance	0.0015	0.0008	0.0007	0.0012	0.0011	0.0006	0.0006	0.0006
IID [2]/shading	0.0014	0.0008	0.0008	0.0013	0.0008	0.0004	0.0004	0.0017
CycleGAN [46]/photo2ukiyoe	0.0024	0.0019	0.0015	0.0013	0.0017	0.0013	0.0010	0.0018
CycleGAN [46]/photo2vangogh	0.0026	0.0026	0.0019	0.0014	0.0020	0.0020	0.0015	0.0020
Colorization [43]	0.0016	0.0011	0.0008	0.0013	0.0010	0.0006	0.0004	0.0012
Colorization [36]	0.0015	0.0009	0.0007	0.0013	0.0010	0.0005	0.0003	0.0012
Average	0.0027	0.0027	0.0021	0.0018*	0.0018	0.0019	0.0015	0.0018*

Table 2: **Quantitative comparison of Temporal Warping Error.** The comparison presented in this table is limited to the results produced by [3, 16] and the proposed method. Lower warp error suggests better temporal consistency. "*"** Indicates partially averaged results.

The quantitative results are evaluated on the basis of temporal warp error. Tab. 2 shows the quantitative result produced by the methods proposed in [3, 16, 19] and our method. Please note that there are over 830 video sequences in the evaluation datasets and over 115,000 frames. Due to the extremely tedious and resource-intensive nature of DVP [19], only a portion of the evaluation datasets could be processed and the evaluated metrics in the DVP [19] column are averaged over only the processed videos for each task. In order to have a fair comparison, the metrics evaluated on the videos processed with DVP [19] are excluded from the large-scale quantitative comparison and are only presented for reference only. Furthermore, it is also noteworthy that in their paper [19], the authors do not evaluate the results on the entirety of the dataset and sample tasks for comparative results and report that their warp error is comparable to that of Bonneel et al. [3].

The lower temporal warp error represents better temporal consistency. The lowest temporal warp error is highlighted

in red, and the second best is highlighted in blue. It is worth mentioning that the temporal warp error does not take into account the perceptual quality of the produced frames and assigns a lower value to blurry videos (as produced by [19]). The quality of the results produced by each of the discussed methods can be verified through the accompanying supplementary video. It is evident from Tab. 2 that the proposed model produces state-of-the-art consistency results. Please note that, unlike the proposed method, all of the compared methods require the availability of raw/unprocessed videos for restoring temporal dynamics at both the training and test time.

We further verify the efficacy of the proposed model through a specialized metric (as presented in [12]) for evaluating motion smoothness in videos, on the portion of evaluation datasets processed through the DVP [19] and present our findings below.

It is worth mentioning that some attempts to judge the perceptual quality through metrics like mean PSNR (performance degradation [19]) have been proposed for this task. We intentionally disregarded this metric and argue that this metric might be helpful to evaluate the performance of successive iterations of their prior-based method where each iteration leads to progressively better results, but it can be a bit misleading in judging the overall perceptual quality of restored videos for this task. A similar phenomenon can be observed with other local perceptual metrics such as LPIPS [45] as it is impossible to quantize a meaningful deviation of perceptual distance or PSNR from per-frame processed videos due to the lack of ground truth data.

In order to be consistent with the prior works, we nonetheless present the LPIPS [45] results in 3. The proposed method performs competitively with the method proposed by [16] while achieving better temporal consistency. It is noteworthy that each of the discussed metrics does not take into account the aspects considered by the other metric and only provides limited insights about the results. Therefore, to properly evaluate the models on both perceptual quality and temporal consistency, thorough user studies are essential and are used as the primary metric for the evaluation of this task. We conducted two separate user studies to properly evaluate all of the discussed methods and present our findings in the next sub-section.

5.1.1 Smoothness Metric (SM)

Due to the unavailability of ground truth videos, the quantitative metric results presented (both in the main paper and this supplemental) can only give us an approximate insight towards the nature of the restored videos. In order to properly evaluate our work, we present further large-scale quantitative results generated on the portion of the evaluation datasets processed through DVP [19] in this section. Fig. 5

Task	DAVIS				VIDEVO			
	Bonneel[3]	Lai[16]	Ours	DVP[19]*	Bonneel[3]	Lai[16]	Ours	DVP[19]*
WCT [20]/antimono	0.2166	0.0431	0.0485	0.3394	0.2180	0.0451	0.0447	0.2992
WCT [20]/asheville	0.2056	0.0373	0.0449	–	0.2680	0.0514	0.0485	–
WCT [20]/candy	0.2588	0.0435	0.0497	–	0.2175	0.0534	0.0514	–
WCT [20]/feathers	0.2096	0.0467	0.0508	–	0.1519	0.0590	0.0540	–
WCT [20]/sketch	0.1451	0.0499	0.0511	–	0.1871	0.0517	0.0496	–
WCT [20]/wave	0.1777	0.0444	0.0485	0.3423	0.2583	0.0749	0.0793	0.3451
Fast-neural-style [13]/princess	0.2238	0.0646	0.0744	0.4782	0.1610	0.0555	0.0625	0.2566
Fast-neural-style [13]/udnie	0.1473	0.0503	0.0609	0.3660	0.0800	0.0514	0.0604	0.1487
DBL [10]/expertA	0.0795	0.0379	0.0495	0.1453	0.0742	0.0504	0.0584	0.1349
DBL [10]/expertB	0.0530	0.0359	0.0481	0.1361	0.1359	0.0639	0.0778	0.1947
IID [2]/reflectance	0.1011	0.0493	0.0536	0.1389	0.0783	0.0659	0.0731	0.2635
IID [2]/shading	0.0646	0.0527	0.0643	0.1558	0.1191	0.0506	0.0592	0.2302
CycleGAN [46]/photo2ukiyoe	0.0908	0.0394	0.0502	0.2351	0.1462	0.0505	0.0575	0.2571
CycleGAN [46]/photo2vangogh	0.1230	0.0390	0.0494	0.2565	0.1224	0.0519	0.0551	0.1701
Colorization [43]	0.0616	0.0402	0.0468	0.1555	0.0867	0.0495	0.0566	0.1441
Colorization [36]	0.0329	0.0373	0.0475	0.1474	0.1317	0.0081	0.0220	0.1558
Average	0.1369	0.0445	0.0524	0.2414*	0.1523	0.0521	0.0556	0.2167*

Table 3: **Quantitative comparison of LPIPS [45].** The comparison presented in this table is limited to the results produced by [3, 16] and the proposed method. Lower LPIPS suggests better perceptual similarity. "*" Indicates partially averaged results.




Figure 5: **Motion Smoothness Results.** This figure presents the quantitative results acquired through the motion smoothness metric as described above. **The lower motion smoothness represents smoother transitions in the videos.**

presents the temporal smoothness results generated through the motion smoothness metric as presented in [12]. This metric tries to quantify the smoothness of videos over various window sizes through differential optical flow as follows:

$$D = (F_{I_1 \rightarrow I_2} - F_{I_0 \rightarrow I_1}) - (F_{O_1 \rightarrow O_2} - F_{O_0 \rightarrow O_1}),$$

$$M(s) = \log \left(\sum_{d \in D} \mathbf{1}_{[s, s+1]} (\|d\|_2) \right) - \log(|D|), \quad (8)$$

Here, $F_{a \rightarrow b}$ denotes optical flow estimated from frame a to frame b . The preceding I and O refers to the raw/unprocessed and restored frames respectively. In Eq. 8, d denotes each of the $[\nabla u, \nabla v]$ vectors in the differential optical flow and s represents the various window sizes for evaluation and $|D|$ denotes the dimensionality of the differential optical flow. $\mathbf{1}_S(d)$ is the masking factor which is equal to 1 if d belongs to the set S . The resultant $M(s)$ gives a measure of temporal smoothness over 3 consecutive frames. Fig. 5 presents the average results on the portion of evaluation datasets processed through [3, 16, 19] and the proposed model. A lower score suggests better smoothness. The proposed model again achieves SOTA results in terms of motion smoothness even over bigger window sizes as well.

Please note that all the presented metrics are subjective

in nature and do not account for the overall quality and consistency of the results due to the lack of ground truth; therefore, user studies should be considered as the main criteria for judging the results of all the proposed approaches for this task. We discuss the findings of the conducted user studies below.

5.1.2 User Study




Figure 6: **User Study.** This donut chart highlights the user preferences collected through the user study. A majority of the participants preferred the results produced by the proposed method.

We conducted two separate user studies; a comprehensive and a factorized user study, to properly evaluate the performance of the proposed model. The first user study consisted of 36 participants with 150 different scenes processed through each of the proposed models. Each participant was asked to judge 5 videos processed through each of the proposed methods for this task. The users were shown a randomly sampled scene from the video pool processed through [3, 16, 19] and this work, simultaneously. The participants were instructed to judge the videos on the basis of naturalness, quality, consistency, content, and style preservation. Figure 6 presents the findings of the first (comprehensive) user study. On average, 41% of the users preferred the videos restored by the proposed method. Furthermore, the results generated by the method proposed in Lai et al. [16] also received a significant preference in the user study but it is worth noting that all the previously proposed approaches require the unprocessed counterparts of the frame-wise processed videos at test time to restore the temporal consistency, whereas the proposed model restores the temporal consistency by only taking into account the inconsistent video at the test time.

5.1.3 Factorized User Study

Apart from the comprehensive user study, we also conducted a secondary in-depth user study to evaluate the reasoning of participants and their opinion on their selections of videos processed with all the methods on challenging scenarios.

In this user study, each user was instructed to judge 15 randomly selected videos (5 videos from each of the applications mentioned above) processed by each of the proposed models for this task. Please note that in this user study users were also shown the per-frame processed video in an extended window to judge the videos on content and style preservation as well. Apart from recording their preference, the participants were also instructed to record the reasoning for their selection. $\sim 86\%$ of the participants remarked that they chose the videos based on lesser flickering and better quality preservation. The results of the second user study for each of the selected tasks are provided in Fig. 7. On average, $\sim 47\%$ of the users preferred the videos processed through the proposed model for these extremely challenging tasks.

5.2. Qualitative Results

Figures 8 ~ 10 present some of the comparative qualitative results produced by the methods proposed in [3, 16, 19] and the proposed model. The results produced by [3] fail to retain the perceptual quality of video sequences where occlusion and deocclusion occur. Whereas the frames processed by DVP [19] are significantly blurry and lose the translation effect. In contrast to the comparison with DVP [19] and Boneel et al. [3], the comparison with the results produced through the method proposed by Lai et al. [16] can be a bit tricky. In order to aid the readers, we provide an epipolar comparison between the frames restored through their method and ours in Fig. 10. Due to the implicit definition of motion dynamics in their method [16], their method fails to faithfully retain the perceptual quality in the restored frames and often flips the colors/appearance of objects in the restored video according to the raw video. This bias towards the raw videos is quite prominent in methods like DVP [19] and Lai et al. [16]. Contrasting to their formulation, the explicit design of motion dynamics in the proposed method mitigates the need for raw videos at test time which consequently eliminates this bias and enables the proposed method to faithfully restore videos. Further qualitative comparisons and higher quality results, along with the results on super resolution [17], inpainting [18], and multiple restoration iterations are presented in the supplemental.

6. Limitations and Future work

The proposed model has difficulties in restoring videos with very low frame rates and in the case of style-transferred videos, some finer inconsistencies like finer brush strokes are lost, which makes the resulting video seem a little dull (which is to be expected from all the proposed approaches for this task). A good direction for future work in this field can be the direct denoising of conventional optical flow with deterministic warp operators. The developed consis-

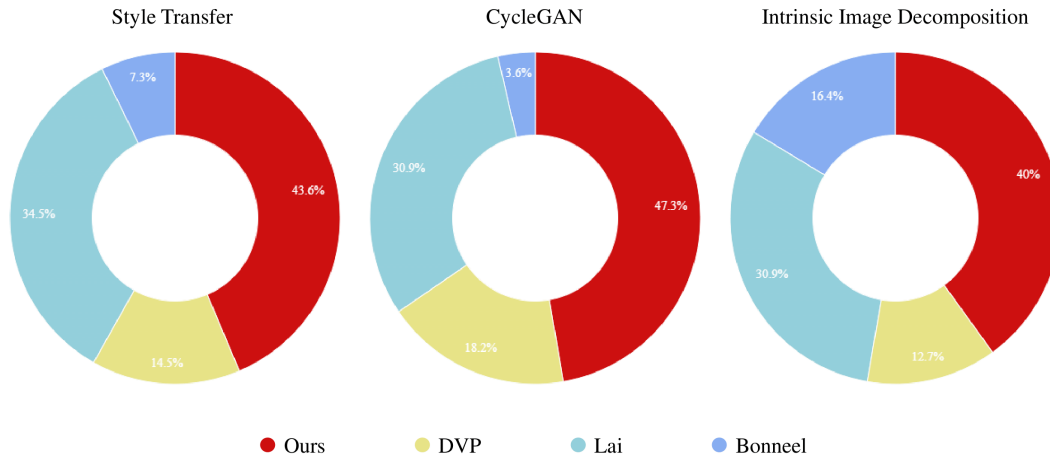


Figure 7: **Factorized user study.** The user preference scores for the tasks of style transfer, CycleGAN, and intrinsic image decomposition (left to right respectively). The majority of the users preferred the videos processed through the proposed method in each category of tasks.




Figure 8: **Comparison with Bonneel et al. [3]** The method proposed by Bonneel et al. [3] has difficulty in handling occlusion and long term consistency and produce temporal artifacts.



Figure 9: **Comparison with DVP [19]** DVP [19] produces blurry frames and compromises the translation effect. This phenomenon is prominent in videos where new content is generated e.g. style transfer [13], CycleGAN [46].

tent flow models can also be used for a number of video processing applications like denoising and dehazing.

7. Conclusion

In this work, we present a task-agnostic temporal consistency correction framework that restores natural video dynamics by inferring and utilizing consistent motion repre-




Figure 10: **Comparison with Lai et al. [16]** This figure produces epipolar temporal comparison of videos restored by Lai et al. [16] and the proposed method. The method proposed in [16], due to its dependence on raw videos and implicit siphoning/mixing of motion dynamics from the raw video introduces bias towards the raw videos.

sentations from temporally inconsistent videos. The proposed strategy mitigates the need for siphoning video dynamics from the unprocessed videos at test time consequently mitigating the compromise in translation effects that is prominent in the currently proposed approaches for this task and expands the scope of its employ-ability to a wider set of applications. Through extensive experimentation and user studies, we demonstrate that despite the limited access to data at test-time the proposed method compares favorably to the pre-existing methods available for this task and achieves SOTA performance.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2022-0-00156, Fundamental research on continual meta-learning for quality enhancement of casual videos and their 3D meta-verse transformation) and Samsung Electronics Co., Ltd, and Samsung Research Funding Center of Samsung Electronics under Project Number SRFCIT1901-06.

References

- [1] Muhammad Kashif Ali, Sangjoon Yu, and Tae Hyun Kim. Deep motion blind video stabilization. *BMVC*, 2021.
- [2] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. *ACM TOG*, 33(4):1–12, 2014.
- [3] Nicolas Bonneel, James Tompkin, Kalyan Sunkavalli, Deqing Sun, Sylvain Paris, and Hanspeter Pfister. Blind video temporal consistency. *ACM TOG*, 34(6):1–9, 2015.
- [4] Dongdong Chen, Jing Liao, Lu Yuan, Nenghai Yu, and Gang Hua. Coherent online video style transfer. In *ICCV*, pages 1105–1114, 2017.
- [5] Mengyu Chu, You Xie, Laura Leal-Taixé, and Nils Thuerey. Temporally coherent gans for video super-resolution (teco-gan). *arXiv preprint arXiv:1811.09393*, 1(2):3, 2018.
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009.
- [7] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Counting out time: Class agnostic video repetition counting in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10387–10396, 2020.
- [8] Wei Gao, Yijun Li, Yihang Yin, and Ming-Hsuan Yang. Fast video multi-style transfer. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3211–3219, 2020.
- [9] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*, 2015.
- [10] Michaël Gharbi, Jiawen Chen, Jonathan T Barron, Samuel W Hasinoff, and Frédo Durand. Deep bilateral learning for real-time image enhancement. *ACM TOG*, 36(4):1–12, 2017.
- [11] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, pages 2462–2470, 2017.
- [12] Meiguang Jin, Zhe Hu, and Paolo Favaro. Learning to extract flawless slow motion from blurry videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8112–8121, 2019.
- [13] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016.
- [14] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Deep blind video decaptioning by temporal aggregation and recurrence. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4258–4267, 2019.
- [15] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Recurrent temporal aggregation framework for deep video inpainting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:1038–1052, 2019.
- [16] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning blind video temporal consistency. In *ECCV*, 2018.
- [17] 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 *CVPR*, pages 4681–4690, 2017.
- [18] Eunhye Lee, Jeongmu Kim, Jisu Kim, and Tae Hyun Kim. Restore from restored: Single-image inpainting. *arXiv preprint arXiv:2102.08078*, 2021.
- [19] Chenyang Lei, Yazhou Xing, and Qifeng Chen. Blind video temporal consistency via deep video prior. *NeurIPS*, 2020.
- [20] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. *NeurIPS*, 30, 2017.
- [21] Zhen Li, Cheng-Ze Lu, Jianhua Qin, Chun-Le Guo, and Ming-Ming Cheng. Towards an end-to-end framework for flow-guided video inpainting. *arXiv preprint arXiv:2204.02663*, 2022.
- [22] Xunyu Lin, Victor Campos, Xavier Giro-i Nieto, Jordi Torres, and Cristian Canton Ferrer. Disentangling motion, foreground and background features in videos. *arXiv preprint arXiv:1707.04092*, 2017.
- [23] Yihao Liu, Hengyuan Zhao, Kelvin CK Chan, Xintao Wang, Chen Change Loy, Yu Qiao, and Chao Dong. Temporally consistent video colorization with deep feature propagation and self-regularization learning. *arXiv preprint arXiv:2110.04562*, 2021.
- [24] Cheng Ma, Yongming Rao, Yean Cheng, Ce Chen, Jiwen Lu, and Jie Zhou. Structure-preserving super resolution with gradient guidance. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7769–7778, 2020.
- [25] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, pages 724–732, 2016.
- [26] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, pages 4161–4170, 2017.
- [27] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241. Springer, 2015.
- [28] Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos. In *GCPR*, pages 26–36. Springer, 2016.
- [29] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *CVPR*, pages 6626–6634, 2018.
- [30] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *NeurIPS*, 28, 2015.
- [31] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, pages 8934–8943, 2018.
- [32] Ximeng Sun, Huijuan Xu, and Kate Saenko. Twostreamvan: Improving motion modeling in video generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2744–2753, 2020.

- [33] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, pages 402–419. Springer, 2020.
- [34] Hugo Thimonier, Julien Despois, Robin Kips, and Matthieu Perrot. Learning long term style preserving blind video temporal consistency. In *ICME*, pages 1–6. IEEE, 2021.
- [35] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *CVPR*, pages 9446–9454, 2018.
- [36] Michael Waechter, Nils Moehrle, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *ECCV*, pages 836–850. Springer, 2014.
- [37] 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.
- [38] Zhicheng Yan, Hao Zhang, Baoyuan Wang, Sylvain Paris, and Yizhou Yu. Automatic photo adjustment using deep neural networks. *ACM TOG*, 35(2):1–15, 2016.
- [39] Yuchen Yang, Ye Xiang, Shuaicheng Liu, Lifang Wu, Boxuan Zhao, and Bing Zeng. Gln-net: Global and local motion estimation via task-oriented encoder-decoder structure. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4211–4219, 2021.
- [40] Chun-Han Yao, Chia-Yang Chang, and Shao-Yi Chien. Occlusion-aware video temporal consistency. In *ACMMM*, pages 777–785, 2017.
- [41] Bo Zhang, Mingming He, Jing Liao, Pedro V Sander, Lu Yuan, Amine Bermak, and Dong Chen. Deep exemplar-based video colorization. In *CVPR*, pages 8052–8061, 2019.
- [42] 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.
- [43] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, pages 649–666. Springer, 2016.
- [44] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, pages 649–666. Springer, 2016.
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.
- [46] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pages 2223–2232, 2017.