

Temporally Coherent Video Harmonization Using Adversarial Networks

Hao-Zhi Huang¹, Member, IEEE, Sen-Zhe Xu, Jun-Xiong Cai², Wei Liu³, Senior Member, IEEE, and Shi-Min Hu⁴, Senior Member, IEEE

Abstract—Compositing is one of the most important editing operations for images and videos. The process of improving the realism of composite results is often called harmonization. Previous approaches for harmonization mainly focus on images. In this paper, we take one step further to attack the problem of video harmonization. Specifically, we train a convolutional neural network in an adversarial way, exploiting a pixel-wise disharmony discriminator to achieve more realistic harmonized results and introducing a temporal loss to increase temporal consistency between consecutive harmonized frames. Thanks to the pixel-wise disharmony discriminator, we are also able to relieve the need of input foreground masks. Since existing video datasets which have ground-truth foreground masks and optical flows are not sufficiently large, we propose a simple yet efficient method to build up a synthetic dataset supporting supervised training of the proposed adversarial network. The experiments show that training on our synthetic dataset generalizes well to the real-world composite dataset. In addition, our method successfully incorporates temporal consistency during training and achieves more harmonious visual results than previous methods.

Index Terms—Harmonization, video editing, temporal consistency.

I. INTRODUCTION

GENERATING realistic composite videos is a fundamental requirement in video editing tasks [3]–[6]. Given two videos, one of them contains a desired foreground, while the other contains a desired background. To generate a realistic composite video from these two source videos, three steps are needed to be performed correctly. First, extract a foreground object from one of the source videos by computing an alpha matte or a binary mask indicating the pixels which belong to the foreground. Second, paste the foreground to a proper location in the background. Third, adjust the foreground

appearance to make it look natural in the new background. The last step is often called **harmonization**. In this paper, we focus on the video harmonization task which intends to improve the realism of a composite video by performing appearance adjustments on the foreground (see Figure 1).

Traditional methods for handling image harmonization are based on learning statistical relationships between hand-crafted appearance features [7]–[10]. These methods neglect whether the foreground and the background are compatible considering the context of one whole image. Recently, some image harmonization methods are proposed to leverage powerful convolutional neural networks (CNNs) to automatically learn features that capture context and semantic information of the composite images, which generate more appealing harmonization results. Zhu *et al.* [1] trained a CNN to distinguish natural images from generated ones. Then they used the predicted realism score to guide a simple color adjustment of the foreground to obtain more realistic composite images. Unlike the method in [1] which takes realism evaluation and improvement as two separated steps, Tsai *et al.* [2] proposed an end-to-end network which takes a composite image and a foreground mask as input to generate a harmonized image directly.

All the methods mentioned above aim at harmonization for images. For videos, although applying image harmonization frame by frame can generate a video harmonization result, this will introduce obvious flicker artifacts in the absence of considering the temporal consistency between consecutive harmonized frames. In this paper, we propose an end-to-end harmonization network for videos, which is able to simultaneously harmonize the composite frames and maintain the temporal consistency between them. In order to make the harmonization results more realistic, we propose to train the network with a pixel-wise discriminator. Different from the most common global discriminators used in the literature which predict whether an image is real or fake as a whole, our discriminator is trained to precisely distinguish the harmonious pixels from the disharmonious ones. In addition, with the well trained pixel-wise discriminator, we are able to predict foreground masks automatically to relieve the need of input foreground masks. On the other hand, for the purpose of constraining the network to generate temporally consistent results, we train it with a temporal loss term to incorporate temporal information in the training phase and avoid the trouble of computing optical flows in the inference phase.

Manuscript received August 29, 2018; revised March 18, 2019 and April 26, 2019; accepted June 10, 2019. Date of publication July 17, 2019; date of current version September 12, 2019. This work was supported in part by the Natural Science Foundation of China under Project 61521002 and in part by the Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Keigo Hirakawa. (*Corresponding author: Hao-Zhi Huang.*)

H.-Z. Huang and W. Liu are with the Tencent AI Lab, Shenzhen 519000, China (e-mail: huanghz08@gmail.com; wl2223@columbia.edu).

S.-Z. Xu, J.-X. Cai, and S.-M. Hu are with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China, and also with the Beijing Engineering Research Center for Intelligent Processing of Visual Media and Content Security, Beijing 100084, China (e-mail: xusz1993@qq.com; caijunxiong000@163.com; shimin@tsinghua.edu.cn).

This article has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/TIP.2019.2925550

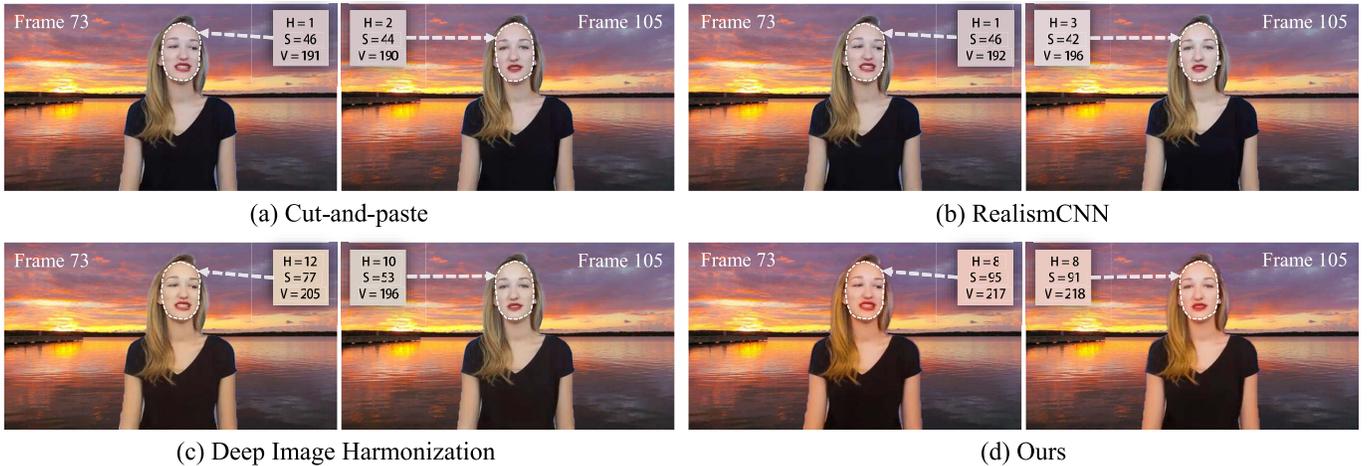


Fig. 1. Given a composite video generated by a direct cut-and-paste operation (a), our method learns to harmonize it to improve its realism. Compared with previous methods RealismCNN [1] (b) and Deep Image Harmonization [2] (c), the harmonized results of our method look more natural and temporally consistent (d). The warm color during sunset is correctly cast to the foreground in our result. The average HSV color values of the facial region are shown for clear comparison.

Training the proposed end-to-end harmonization network requires a large amount of video data taken at different scenes with ground-truth harmonious and disharmonious pairs, foreground masks, and optical flows, which is still a missing piece in the community. To this end, we propose a simple yet effective way to build up a synthetic dataset which satisfies this demand. We also build up a real-world composite video dataset for evaluating the proposed method in real scenarios, which helps demonstrate that training on our synthetic dataset enables our network to generalize to real-world composite videos. Extensive experiments on the two datasets demonstrate that the proposed method is able to conduct temporally coherent video harmonization while generating more harmonious results than existing methods.

The contributions of our method are two-fold. Firstly, to the best of our knowledge, this is the first end-to-end CNN for video harmonization. The network is trained in an adversarial way with an introduced temporal loss to simultaneously acquire high-quality harmonization results and temporal consistency. Secondly, a synthetic dataset is proposed for supporting the efficient training of the video harmonization network, which contains ground-truth harmonious and disharmonious pairs, foreground masks, and optical flows. This synthetic dataset makes the temporal loss computable and alleviates the issue of lacking training video training data. Thirdly, a pixel-wise adversarial loss is integrated to our system to make the foreground look more harmonious in front of the new background, which is the first time the adversarial loss is leveraged for the video harmonization problem.

II. RELATED WORK

Our work attempts to generate a temporally consistent harmonized video by an adversarial network. This is closely related to the literature on image harmonization, conditional generative adversarial networks, and temporally consistent video editing.

A. Image Harmonization

Traditional methods for image harmonization focus on matching appearance statistics without considering the context of images, such as aligning the statistics of global or local histograms [10], [11], shifting the colors towards predefined harmonious color templates [7], gradient-domain compositing [12]–[14], multi-scale matching of various statistics [9], and maximizing the co-occurrence probability of color distributions [8]. Recently, Zhu *et al.* [1] trained a discriminative model based on a CNN to predict a realism score for a composite image, and then used the score to determine a simple brightness and contrast adjustment for improving the realism. Instead of separating realism evaluation and improvement into two steps, Tsai *et al.* [2] proposed an end-to-end CNN to learn how to harmonize a composite image directly. To improve the ability of capturing semantic information, Tsai *et al.* pretrained the harmonization network with semantic segmentation and used the segmentation branch to provide features to help harmonization. Different from these existing algorithms which are engaged in image harmonization, we propose an end-to-end CNN trained in an adversarial manner with a temporal loss to solve the problem of video harmonization, which generates realistic and temporally consistent harmonized results. With the help of the temporal loss, our method can overcome the flicker artifacts produced by directly applying image harmonization methods to videos.

B. Conditional Image Generation Based on Adversarial Training

Generative adversarial network (GAN) was first proposed by Goodfellow *et al.* [15] to address the problem of realistic image generation from input noise variables. The key idea of GAN is to train a generator and a discriminator in an adversarial fashion. While the discriminator is trained to distinguish fake images from real ones, the generator is trained

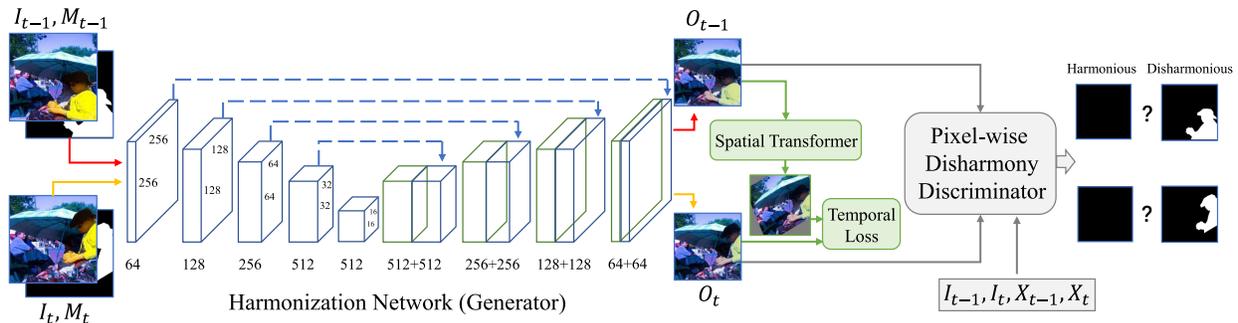


Fig. 2. The training phase of the proposed video harmonization model. The harmonization network is trained in an adversarial manner with the pixel-wise disharmony discriminator. A two-frame coordinated training strategy is adopted to incorporate a regional temporal loss to constrain the consecutive harmonized foregrounds to have visually similar appearances.

to deceive the discriminator and generate images as realistic as possible. Recently, GANs have also been widely used in the task of conditional image/video generation [16]–[26]. Although the Markovian discriminator proposed in [17], [27] performs a patch-level discrimination instead of a global one, it is designed for speeding up the discriminator with a very small receptive field, which is still insufficient for the harmonization task. In this paper, we propose to use an encoder-decoder structure to learn a pixel-wise discriminator which labels each pixel as harmony or disharmony. The encoder-decoder structure provides a large receptive field for calculating a more precise adversarial loss. With the proposed discriminator, we can also complete the harmonization task without input foreground masks.

C. Temporally Consistent Video Editing

Directly applying image harmonization frame by frame for videos inevitably results in flicker artifacts. This is because the corresponding regions in different frames are harmonized in different ways. Plenty of approaches have been proposed to enforce temporal consistency in video editing, such as spatio-temporal smoothing [28]–[30], optimization with a temporal loss [31], [32], frame propagation [33], *etc.* The above methods either require extra post-processing operations or rely on a time-consuming optimization procedure. Recently, some video style transfer methods [34], [35] show that temporal consistency and style transfer can be simultaneously learned by a CNN, which acquires considerable temporal consistency at a very little time cost. The temporal loss is designed for constraining the color consistency under the guidance of optical flow. Optical flow has a long history in extracting video temporal information [36]–[40]. Inspired by the above ideas, our proposed method also chooses to incorporate the temporal consistency during the training phase. However, instead of calculating a global temporal loss, we introduce a regional temporal loss that forces our model to pay more attention to the disharmonious regions, therefore leading to more coherent harmonized results.

III. VIDEO HARMONIZATION NETWORK

In this section, we describe the details of our proposed end-to-end CNN for video harmonization. Figure 2 shows an overview of our network. The harmonization network takes

one frame of a composite video and a foreground mask as input, and performs appearance adjustments on the foreground while keeping the background unchanged. To incorporate temporal consistency between consecutive harmonized frames, a two-frame coordinated training strategy with a regional temporal loss is adopted. Note that in the training phase the two frames are fed to the harmonization network in a coordinate but separate way, while in the testing phase the harmonization network processes a video in a frame by frame way. This kind of setting has been proven to be effective for training a network which conducts smoother transformation with less flicker artifacts in video style transfer [35]. To further enhance the realism of the harmonized results, the harmonization network is trained in an adversarial way with a pixel-wise disharmony discriminator, which distinguishes the disharmonious pixels from the harmonious ones. Moreover, the well-trained discriminator can also be employed to predict the disharmony area in the input, which holds as a replacement of the input foreground mask.

A. A Synthetic Training Dataset

Before delving deeply into the network architecture, it is essential to describe the way we collect data. For supervised training, our harmonization network needs a composite video and a corresponding harmonized video as a sample pair. Given an arbitrary composite video, it is hard to acquire a high-quality harmonized result even for a human expert. For the image harmonization task, Tsai *et al.* [2] collected images from the MSCOCO dataset [41] which have ground-truth foreground masks, and then applied color transfer between random foreground pairs with the same semantic labels. While the image after a foreground adjustment is used as the input, the original image is used as the ground-truth. The difficulty of extending this idea to the video harmonization task lies in the fact that there are a limited number of videos that have ground-truth foreground masks. Even in the very recent video object segmentation dataset DAVIS [42], there are only 90 annotated videos. This number is far from being enough, because the harmonization network requires training data covering tremendous scenes to learn the natural appearances of foregrounds in various cases. Meanwhile, we also need the ground-truth optical flows between consecutive frames for evaluating the temporal consistency.

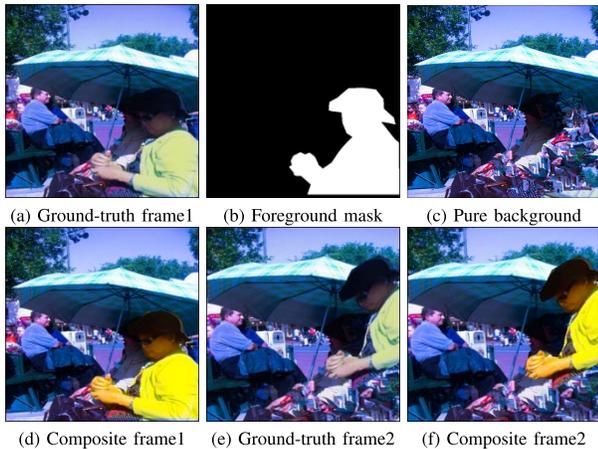


Fig. 3. Building up the synthetic dataset. Given an image (a), we take it as the first ground-truth frame. Then we cut out the foreground and apply inpainting to obtain the pure background (c). By performing color adjustment on the foreground of (a), we obtain the first composite frame (d). By applying a random affine transform to the foregrounds of (a) and (d), we obtain the second ground-truth frame (e) and the second composite frame (f).

To address this data issue, we construct a synthetic dataset named Dancing MSCOCO which contains ground-truth foreground masks and optical flows. Based on the MSCOCO dataset, we apply small-scale random affine transforms to the foregrounds and acquire a series of images containing the same “dancing” foreground, which simulate the consecutive frames in a real video. A similar strategy has been adopted by Dosovitskiy *et al.* [43] to collect data for training a FlowNet to predict optical flows, which shows competitive performance compared to state-of-the-art methods like DeepFlow [44] and EpicFlow [45]. Similarly, Khoreva *et al.* [46] used synthetic frames to help the training of an object tracking system. The success of [43], [46], [47] proves that proper synthetic data is sufficient to some extent for training a deep neural network targeting at dealing with real video data.

Key outputs during the process of building up our Dancing MSCOCO dataset are shown in Figure 3. First, we select images containing ‘people’ as the foregrounds from the MSCOCO dataset, and wipe off the images whose foreground area is smaller than 10% of the whole image. Since the image numbers of different classes are imbalanced in the MSCOCO dataset, training all the classes together inevitably introduces biases. In this paper, we simply focus on images containing people to avoid a class bias. It is easy to transfer the proposed method to other kinds of foregrounds. Further solutions for avoiding a class bias are beyond the scope of this paper. Second, we cut out the foreground and apply inpainting [48] to fill the holes to obtain pure background images. For each background image, we perform random cropping to obtain a distorted copy of the background image and resize it back to the original size, which simulates a background movement in a video. Third, we apply color adjustments to the foregrounds to simulate the composite images. Besides performing color transfer [11] between random foreground pairs as in [2], we also perform random adjustments of the basic color properties including exposure, hue, saturation, temperature,

contrast, and tone curve. This makes our dataset cover more kinds of compositing situations than the dataset created in [2]. Fourth, we apply the same random affine transform to the original foreground and the color adjusted foreground, which simulates a foreground movement in a video, and paste it back to the corresponding randomly cropped background. Since we know the exact affine transform between the foregrounds, it is easy to acquire the corresponding ground-truth optical flow. The affine transform parameters are sampled from a suitable range in order not to conduct large-scale foreground movements. Finally, we achieve a total number of 33,338 pairs of “consecutive original frames” and their corresponding “consecutive composite frames”. 29,818 pairs of them are used as the training data, 1,000 pairs are used as the validation data, and 2,520 pairs are used as the testing data. We have also tried generating more than one distorted copy for each image to simulate the situation of multiple consecutive frames, but found a marginal improvement in the training. This is because only generating two consecutive frames is already enough for encouraging the temporal consistency between neighbors. In the case of video harmonization, long-term temporal consistency is not a problem since the foreground does not undergo a vast color adjustment. Thus, generating more frames for constraining the long-term temporal consistency does not bring apparent improvements.

B. Adversarial Training With a Temporal Loss

Our network contains two parts, a harmonization network G behaving as a generator and a pixel-wise disharmony discriminator D . The harmonization network G processes a composite video frame by frame to generate a harmonized video. At each time step t , G takes a composite frame I_t and a foreground mask M_t as input and generates a harmonized output frame $O_t = G(I_t, M_t)$, which adjusts the appearance of the foreground to make it look more natural in the background. To incorporate temporal consistency, the network G is trained in a two-frame coordinated manner to constrain the consecutive outputs O_{t-1} and O_t to be temporally coherent. Thanks to this training strategy, G is able to generate temporally consistent harmonized frames, though processing a video in a frame-by-frame way. The insight behind this strategy is that by constraining the outputs to be temporally consistent, we train a more stable network that can avoid amplifying small differences in the inputs. This strategy makes the harmonization network learn smoother transformations. To acquire more realistic harmonized results, we propose a pixel-wise disharmony discriminator D to play against G by telling disharmonious pixels from harmonious ones. To deceive the discriminator D , G learns to generate more and more realistic results. Here, we implement the discriminator D with a large receptive field to take the context into account when learning to distinguish whether a pixel is harmonious or not. Instead of giving a global true or false label to indicate whether the overall picture is realistic or not, our pixel-wise disharmony discriminator localizes the disharmonious regions more accurately, which encourages the generator to focus on adjusting the challenging disharmonious regions towards achieving the most realistic results.

1) *Pixel-Wise Disharmony Discriminator*: The objective for training D is to classify harmonious pixels into class 0 and disharmonious ones into class 1, respectively, while keeping the parameters of G fixed. Both the input image and the harmonized result can be taken as the fake samples, while the ground-truth images in the dataset are taken as the real samples. The foreground regions in the fake samples are regarded as the disharmonious pixels. The loss function we aim to minimize is:

$$\mathcal{L}_D = \frac{1}{2N} \|D(O_t) - M_t\|_2^2 + \frac{1}{2N} \|D(I_t) - M_t\|_2^2 + \frac{1}{N} \|D(X_t)\|_2^2. \quad (1)$$

Here, M_t is the input foreground mask, in which the foreground pixels are labeled as 1 and the background pixels are labeled as 0. $D(O_t)$ is the discriminator output for O_t , which should be close to M_t for guiding the discriminator to distinguish the harmonized pixels from the ground-truth realistic pixels. In our experiments, we find that training D merely with O_t cannot generalize well to I_t , which is also a kind of image with a disharmonious foreground. I_t can be regarded as some easy samples for D to warm up, since the foreground color of I_t is distinguishing compared to the background. Thus, we also feed I_t to D as a disharmonious sample, and enforce $D(I_t)$ to be close to M_t . We give the loss terms relevant to O_t and I_t the same weight. $D(X_t)$ is the discriminator output for the ground-truth realistic frame X_t , in which all pixels should be labeled as 0. Note that, unlike the most common global discriminators in the literature which consider an image as a whole to be real or fake, our discriminator learns to classify each pixel separately. This is because the background pixels in O_t and I_t are definitely harmonious, which should be treated differently from those disharmonious pixels for training a more precise discriminator. To correctly label a pixel, the discriminator should have a large receptive field to capture context information. Thus a UNet [49] architecture is used as the harmonization network, which will be described in detail later. Theoretically, our method is a variant of LSGAN [50], which shows a more stable and faster convergence than vanilla GAN [15] and WGAN [51].

2) *Harmonization Network*: The objective for training G is to generate temporally coherent harmonized frames that are indistinguishable from realistic ones. Thus, G is learning to deceive the discriminator to minimize the adversarial loss, while encouraging the temporal consistency between consecutive frames. Since we know the ideal harmonized image (the ground-truth image without any color adjustment), a reconstruction loss is added directly in a supervised manner. The loss function that we minimize for training G is:

$$\mathcal{L}_G = \underbrace{\frac{1}{N} \|O_t - X_t\|_2^2}_{\text{Reconstruction Loss}} + \underbrace{\frac{\lambda_1}{N_F} \|M_t \circ (O_t - S(O_{t-1}))\|_2^2}_{\text{Regional Temporal Loss}} + \underbrace{\frac{\lambda_2}{N} \|D(O_t)\|_2^2}_{\text{Adversarial Loss}}. \quad (2)$$

Here, N denotes the total number of pixels in a frame, and N_F denotes the number of pixels in the foreground. The harmonization network is trained with a combination of a reconstruction loss, a regional temporal loss, and an adversarial loss. The reconstruction loss enforces the harmonized frame O_t to be similar to the ground-truth realistic frame X_t . The temporal loss enforces the harmonized frame O_t to be temporally consistent with the previous harmonized frame O_{t-1} . S denotes a Spatial Transformer Network [52], which warps O_{t-1} according to the ground-truth optical flow provided by our Dancing MSCOCO dataset. The Spatial Transformer Network is fully differentiable, which makes our network end-to-end trainable. Instead of treating each pixel equally as in the global temporal loss [35], our regional temporal loss focuses on the foreground region by taking a Hadamard product between the foreground mask and the global element-wise difference between frames: $M_t \circ (O_t - S(O_{t-1}))$. The insight behind this is that since the network only needs to learn a simple identity mapping for the background pixels, it should pay more attention to learning how to generate temporally consistent foregrounds. In our experiments, we find that the generated backgrounds achieve great temporal consistency even without any applied temporal loss.

On the other hand, the adversarial loss encourages G to generate harmonized results that are indistinguishable for the discriminator D , hence enforcing the output of G to be close to the manifold underlying realistic frames. Although the reconstruction loss has enforced the harmonized frames to be similar to the realistic frames, this is not enough because there may be multiple answers besides the ground-truth for harmonizing a composite frame. Different training samples with similar compositing situations may teach G to learn different harmonization solutions for similar inputs. This finally prompts G to output an average of different solutions, which may fall outside the manifold of realistic frames. Leveraging D to provide an adversarial loss for training G can deal with this issue.

As shown in Figure 2, our harmonization network adopts the architecture like UNet [49], which has skip connections to reserve more content details which may be lost during the progressive downsampling in the convolutional layers. Our discriminator also adopts the same architecture for acquiring a large receptive field. The network proposed for deep image harmonization [2] also used skip connections but in the form of element-wise summation instead of feature concatenation in the UNet. We have also tried the residual network used in style transfer [53]. The superiority of the UNet over other networks is clearly validated in Section IV-C.

An algorithm procedure is provided here to give a clearer illustration of the training phase.

C. Harmonization Without Foreground Masks

Previous state-of-the-art harmonization methods all require a foreground mask as the input besides the composite image itself. In this paper, with the help of a well-trained pixel-wise disharmony discriminator, we can predict the disharmonious foreground area automatically. Thus, we are able to accomplish the harmonization task without an input foreground mask

Algorithm 1 Training Phase of the Proposed Method

Input: input images $\{I\}$, foreground masks $\{M\}$,
and ground-truth harmonized results $\{X\}$.

Output: network parameters θ_G and θ_D for the generator and the discriminator, respectively.

Initialization:

- 1: Initialize the network parameters θ_G and θ_D using Xavier [54];

Loop Process:

- 2: **for** $epoch = 0$ to max_epoch_num **do**
- 3: Optimize the discriminator D for one step by minimizing \mathcal{L}_D ;
- 4: Optimize the generator G for one step by minimizing \mathcal{L}_G ;
- 5: Test on the validation set and record the best model up to now;
- 6: **end for**
- 7: **return** θ_G and θ_D .

using the same harmonization network which is trained with foreground masks: $O_t = G(I_t, D(I_t))$. Another solution for achieving the same goal is to train G without input foreground masks in the first place. In our experiments, we find that this solution achieves similar performance to ours but comes with an extra labor of training another network.

IV. EXPERIMENTS

In this section, we present extensive experiments to evaluate the effectiveness of the proposed method.

Real-World Composite Dataset: Besides the synthetic dataset, we also build up a dataset containing real-world composite videos to evaluate the effectiveness of the proposed method. First, we collect videos with the tag ‘fashion’ from Youtube-8M dataset [55], the content of which usually is a person talking in front of a static camera. We require the videos to be taken by a static camera because the incompatible movements of cameras also influence the realism of a composite video, which is out of the scope of this paper. For each video, we cut out a 5-10 seconds clip which contains the desired content. To generate the ground-truth foreground masks, we utilize the Rotobrush tool in Adobe After Effects CC 2017 to manually label the foreground regions. Additionally, we collect various background videos from Videvo.net [56]. For each background video, we apply random adjustments of basic color properties as in Section III-A to get two more distorted copies to cover various composite situations. In total, we acquire 30 foreground videos and 48 background videos. Then we extract the foregrounds paste them to different backgrounds one by one. In the end, we get 1440 composite videos. For this real-world composite dataset, the ground-truth optical flows between frames are unavailable. Instead, we utilize the state-of-the-art method [47] to estimate the optical flows.

Implementation Details: During training, we resize the input and the ground-truth frame to 512×512 , and scale the range of color values to $[-1, 1]$. We set $\lambda_1 = 2 \times 10^{-2}$ and $\lambda_2 = 10^{-2}$ with a fixed learning rate of 2×10^{-4} . We use a batch size

of 1 to alternatively train our harmonization network and pixel-wise disharmony discriminator for 45 epochs. Here, the batch size 1 means for each iteration we feed only one pair of consecutive frames for training. For optimization we use Adam [57] with $\beta_1 = 0.999$. Hyper parameters are chosen by experiments on the validation dataset.

A. Comparison With Previous Methods

We evaluate the proposed method through comparisons with two state-of-the-art methods [1], [2] and the cut-and-paste baseline. The cut-and-paste method means we directly paste the foreground onto the new background without any modification. Note that the cut-and-paste result is the input of the harmonization network. The direct cut-and-paste results do not have any color adjustment, which makes the foreground look unharmonious in front of the new background. Since there is no adjustment on the foreground, perfect temporal coherence is inherited directly from the input video. Here, we choose to compare with previous methods using their original settings and training data to demonstrate the temporal inconsistency of image harmonization methods and the effectiveness of the proposed synthetic dataset.

Table I shows the quantitative evaluation results. PSNR and MSE are calculated on the synthetic dataset. \mathcal{L}_{T1} represents the regional temporal loss computed according to Equation 2 using the ground-truth optical flow in our synthetic dataset. \mathcal{L}_{T2} represents the regional temporal loss computed using the estimated optical flow in the real-world composite dataset. We show that our method achieves better performance compared to the state-of-the-art methods [1], [2] in terms of both PSNR, MSE and temporal losses. Although the cut-and-paste baseline shows better temporal consistency, its realism is far from being satisfactory.

To demonstrate training on our synthetic dataset is able to generalize well to real-world composite videos. We set up a user study similar as [1], [2] using 15 real-world composite videos randomly picked from the real-world composite dataset, in which each user watches four videos and is asked to rank the harmonized results of the four methods regarding to either single frame realism or temporal consistency between frames. As a result, a total of 32 subjects participate in this study with a total of 480 rankings over the four candidate methods. Then we use the Plackett-Luce (P-L) model [58] to compute the global ranking score for each method. Table I shows that compared with the other harmonization methods, our method achieves the highest ranking score according to both single frame realism and temporal consistency. It is no surprise that the cut-and-paste method achieves the best temporal consistency score, but it comes at the cost of realism because no appearance adjustment is applied.

Figure 1, Figure 4 and Figure 5 illustrate the visual results generated by different methods. Overall, our method generates more harmonious and temporally consistent results than previous methods. RealismCNN [1] may generate unsatisfactory results when the realism prediction process fails. Among all the methods, the foreground colors of our results are the most consistent with the backgrounds. In addition, both [1] and [2]

TABLE I

COMPARISON WITH PREVIOUS METHODS. QUANTITATIVE RESULTS ARE EVALUATED ON THE SYNTHETIC DATASET AND THE REAL-WORLD COMPOSITE DATASET. THE USER STUDY RESULTS SHOW THE PLACKETT-LUCE SCORES ACQUIRED BY DIFFERENT METHODS

Method	Quantitative Results				User Study Results	
	PSNR	MSE	\mathcal{L}_{T1}	\mathcal{L}_{T2}	Realism	Temporal
Cut-and-paste	18.19	0.030	0.0006	0.027	-0.004	2.176
Zhu [1]	18.27	0.032	0.0822	0.098	-0.330	-2.064
Tsai [2]	18.42	0.024	0.2325	0.029	-0.179	-1.291
Ours	22.45	0.009	0.0247	0.026	0.551	1.179

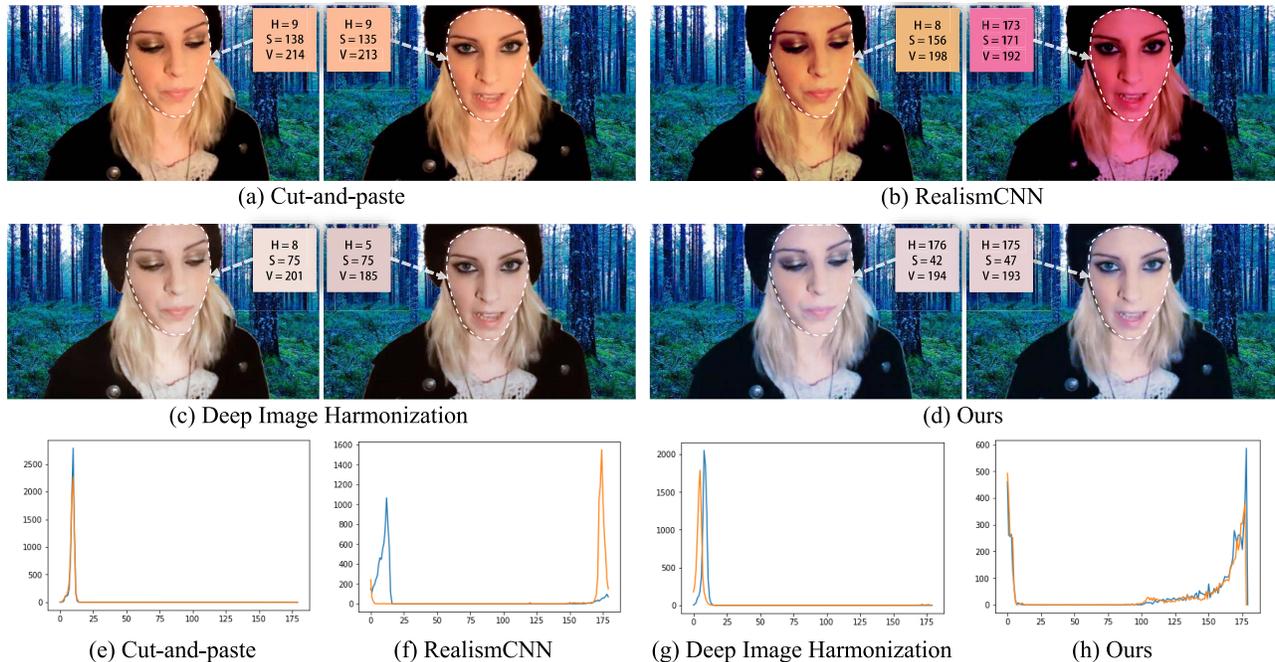


Fig. 4. Visual results on the real-world composite dataset. (a) are the cut-and-paste results, *i.e.*, the inputs to the harmonization methods. (b) are the results of RealismCNN [1]. (c) are the results of Deep Image Harmonization [2]. (d) are our results. Among all methods, our harmonized results appear most harmonious and temporally consistent. The foregrounds of our results are adjusted to cool color tone which looks more harmonious with the background. The average HSV color values and the hue distributions of the facial regions are shown for a clear comparison. The blue curve is for the first frame, and the orange curve is for the second frame.

generate foregrounds in different appearances across frames, which leads to flicker artifacts in the video. More visual comparisons can be found in the supplementary material.

B. Effectiveness of the Synthetic Dataset

To evaluate the effectiveness of the synthetic dataset, we train our harmonization network using different datasets, and test on the synthetic dataset and real-world composite dataset. The datasets we have tested include the MSCOCO dataset [41], the DAVIS dataset [42], and our Dancing MSCOCO dataset. For the MSCOCO dataset, similar as the preprocessing for our synthetic dataset, we only keep those images containing people as the foregrounds and those with foreground occupying over 10% of the whole image, and apply color transfers between foregrounds to acquire composite images as in [2]. For the DAVIS dataset, we apply random adjustments to the basic color properties to create 62 recolored copies for each of the original 70 videos, resulting

in 4410 videos in the end. When training on the DAVIS dataset, we use the same temporal loss setting as in our proposed model. We also train on our dataset without the temporal loss for comparison. The last two rows in Table II shows that training on the synthetic dataset with the proposed temporal loss \mathcal{L}_T achieves smaller MSE and temporal losses than training on the other datasets. This means that our Dancing MSCOCO dataset generalizes the model to more composite cases, while the temporal loss incorporated during the training phase leads to temporally consistent harmonized results in the inference phase. Although training on the DAVIS dataset leads to a smaller temporal loss on the real-world composite dataset, it comes at the sacrifice of the single frame realism. This is because DAVIS dataset covers limited number of scenarios, which leads to overfitting and makes generalizing to other dataset very difficult. Training without the temporal loss on the proposed dataset will achieve the best PSNR and MSE scores, because these two scores are consistent with the remaining L_2 reconstruction loss.



Fig. 5. Visual results on the synthetic dataset. The first column are the ground-truth harmonious frames. The second column are the input frames. The third column are the results of RealismCNN [1]. The fourth column are the results of Deep Image Harmonization [2]. The last column are our results. Among all methods, our harmonized results are closest to the ground-truths and most temporally coherent.

TABLE II
PERFORMANCES OF TRAINING ON DIFFERENT DATASETS

Dataset	PSNR	MSE	\mathcal{L}_{T1}	\mathcal{L}_{T2}
MSCOCO	21.38	0.011	0.051	0.027
DAVIS	17.11	0.031	0.060	0.015
Ours (w/o \mathcal{L}_T)	23.29	0.008	0.043	0.033
Ours (w/o adv)	22.13	0.009	0.022	0.023
Ours	22.45	0.009	0.025	0.026

TABLE III
PERFORMANCES OF DIFFERENT NETWORK ARCHITECTURES

Model	PSNR	MSE	\mathcal{L}_{T1}	\mathcal{L}_{T2}
VGG+color adjust [1]	18.29	0.029	0.083	0.097
ResNet [53]	21.38	0.011	0.020	0.026
DIH [2]	19.82	0.014	0.014	0.022
UNet [49]	22.13	0.009	0.022	0.023

However, in this case, the temporal consistency between consecutive frames will be lost, which will introduce obvious flicker artifacts in the result videos. Thus, although introducing the temporal loss will result in a worse PSNR/MSE score, it is still necessary for the video harmonization task. It is also interesting to see that, training without the adversarial loss (w/o adv) achieves a slightly better temporal consistency but at the cost of lower single-frame quality. This is because without the adversarial loss, the training process can give more attention on minimizing the temporal loss. We also conduct a user study to evaluate the effectiveness of the adversarial loss for improving the single-frame quality in Section IV-C. Note that, introducing the adversarial loss improves the single-frame quality and only introduces a slight increase of the temporal loss, which does not harm the temporal consistency too much.

C. Analysis of the Network Architecture

1) *Model Choice*: To justify the choice of the UNet architecture, we train different networks without the discriminator and using the same setting on the Dancing MSCOCO dataset. The network architectures we have tested include the UNet [49], the deep image harmonization network (DIH) proposed in [2], the style transfer residual network proposed in [53]. The number of parameters in different architectures are constrained to the same level. We also finetune the VGG classification network for 20,000 iterations using our dataset following a similar procedure in [1] and use the same color adjustment method to harmonize the images to give a fair comparison. Table III shows that the UNet achieves the best PSNR and MSE scores. Firstly, dividing the harmonization into two stages as in [1] is not a good choice obviously. Under the end-to-end training setting, UNet achieves the best result, which is because the skip connections used in the UNet reserve image details well and keep extra information that is needed by the harmonization task. Although DIH obtains the smallest \mathcal{L}_{T1} , its MSE is the largest and PSNR is the lowest. As ensuring the harmonization quality is the first priority, we choose to train a UNet for accomplishing the video harmonization task in an end-to-end manner.

2) *Regional Temporal Loss v.s. Global Temporal Loss*: To demonstrate the effectiveness of the proposed regional temporal loss, we train our harmonization network without the discriminator using either the regional temporal loss or the global temporal loss. After converging to similar PSNRs and MSEs, we stop the training and compare the temporal losses in the foreground regions. While using regional temporal loss results in $\mathcal{L}_{T1} = 0.022$, $\mathcal{L}_{T2} = 0.023$, using global temporal loss results in $\mathcal{L}_{T1} = 0.038$, $\mathcal{L}_{T2} = 0.025$. It shows that using regional temporal loss can enforce the model to create more temporally consistent foregrounds.

3) *Ablation Study on the Adversarial Loss*: To evaluate the effectiveness of the adversarial loss, besides the proposed



Fig. 6. The first row shows the results generated by the model trained with a pixel-wise disharmony discriminator. The second row shows the results generated by the model trained without a discriminator, in which the foreground boundaries are more obvious.



Fig. 7. Disharmony maps predicted by our pixel-wise discriminator. The first row are the inputs with the ground-truth foreground masks at the top-right corner. The second row are the predicted disharmony maps.

model, we also train a model without the discriminator. Figure 6 shows harmonized results generated by the two models, from which we can see that the model trained with a discriminator produces harmonized foregrounds with more realistic appearances and the foreground boundaries are less obvious. We also set up a user study using 15 videos randomly picked from the real-world composite dataset, in which each user watches a pair of videos generated by the two methods at a time, and is asked to choose the one that looks more realistic. A total of 36 subjects participate in this study with a total of 540 pairwise comparisons over the two candidate methods. While 53% of the choices prefer the model with the discriminator, 38% of the choices prefer the one without the discriminator. The rest 9% think that they are equally good. This demonstrates that the pixel-wise disharmony discriminator contributes to the generation of more realistic harmonized results.

D. Evaluation of Estimated Foreground Masks

With a well-trained pixel-wise disharmony discriminator, we are able to use the predicted disharmony map instead of the ground-truth foreground mask as input to generate harmonized

TABLE IV
RESULTS OF DIFFERENT SOLUTIONS WITHOUT
INPUT FOREGROUND MASKS

Solution	PSNR	MSE
All-Zero Mask	17.72	0.031
All-One Mask	14.00	0.048
Training from Scratch	20.04	0.016
Adversarial D	19.87	0.018
Adversarial D (with finetuning)	20.21	0.015
Full Model with Mask	22.45	0.009

results. Figure 7 shows that the disharmony maps predicted by our discriminator are close to the ground-truth foreground mask and are considerable to be used as a replacement of the input foreground masks. For comparison, we test alternative solutions including using an all-zero mask or an all-one mask with our trained harmonization network, and training a new harmonization network from scratch taking no foreground mask as input. We also try to finetune the trained full model with the predicted disharmony maps. Table IV shows that while all-zero mask and all-one mask solutions result in large MSEs and low PSNRs, using a predicted disharmony map (Adversarial D) achieves similar performance as training a new network from scratch. After finetuning our full model with the predicted disharmony maps, we achieve better PSNR and MSE scores when the ground-truth foreground masks are not available. This demonstrates that the discriminator endow the proposed network with the flexibility of being used either with or without input foreground masks.

V. CONCLUSION

In this paper, we proposed an end-to-end CNN for tackling video harmonization. The proposed network contains two

parts: a generator (*i.e.* the harmonization network) and a pixel-wise discriminator. While the harmonization network takes a composite video as input and outputs a harmonized video that looks more realistic, the pixel-wise discriminator plays against it by learning to distinguish disharmonious foreground pixels from those harmonious ones. To maintain temporal consistency between consecutive harmonized frames, a regional temporal loss was adopted to enforce the harmonization network to pay more attention to generating temporally coherent harmonized foregrounds. To address the issue of lacking suitable video training data, a synthetic dataset Dancing MSCOCO was constructed for achieving harmonization and temporal correspondence ground-truths at the same time. With the help of a pretrained pixel-wise disharmony discriminator, we can relieve the necessity of an input foreground mask for gaining a considerable harmonization quality. The extensive experiments demonstrate the superiority of our method over the state-of-the-arts.

REFERENCES

- [1] J.-Y. Zhu, P. Krahenbuhl, E. Shechtman, and A. A. Efros, "Learning a discriminative model for the perception of realism in composite images," in *Proc. ICCV*, Dec. 2015, pp. 3943–3951.
- [2] Y.-H. Tsai, X. Shen, Z. Lin, K. Sunkavalli, X. Lu, and M.-H. Yang, "Deep image harmonization," in *Proc. CVPR*, Jul. 2017, pp. 3789–3797.
- [3] T. Chen, J.-Y. Zhu, A. Shamir, and S.-M. Hu, "Motion-aware gradient domain video composition," *IEEE Trans. Image Process.*, vol. 22, no. 7, pp. 2532–2544, Jul. 2013.
- [4] F.-L. Zhang, X. Wu, H.-T. Zhang, J. Wang, and S.-M. Hu, "Robust background identification for dynamic video editing," *ACM Trans. Graph.*, vol. 35, no. 6, 2016, Art. no. 197.
- [5] H. Huang, X. Fang, Y. Ye, S. Zhang, and P. L. Rosin, "Practical automatic background substitution for live video," *Comput. Vis. Media*, vol. 3, no. 3, pp. 273–284, 2017.
- [6] T. Gupta, D. Schwenk, A. Farhadi, D. Hoiem, and A. Kembhavi, "Imagine this! Scripts to compositions to videos," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 598–613.
- [7] D. Cohen-Or, O. Sorkine, R. Gal, T. Leyvand, and Y.-Q. Xu, "Color harmonization," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 624–630, 2006.
- [8] J.-F. Lalonde and A. A. Efros, "Using color compatibility for assessing image realism," in *Proc. ICCV*, Oct. 2007, pp. 1–8.
- [9] K. Sunkavalli, M. K. Johnson, W. Matusik, and H. Pfister, "Multi-scale image harmonization," *ACM Trans. Graph.*, vol. 29, no. 4, 2010, Art. no. 125.
- [10] S. Xue, A. Agarwala, J. Dorsey, and H. Rushmeier, "Understanding and improving the realism of image composites," *ACM Trans. Graph.*, vol. 31, no. 4, 2012, Art. no. 84.
- [11] E. Reinhard, M. Adhikmin, B. Gooch, and P. Shirley, "Color transfer between images," *IEEE Comput. Graph. Appl.*, vol. 21, no. 5, pp. 34–41, Jul./Aug. 2001.
- [12] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 313–318, 2003.
- [13] J. Jia, J. Sun, C.-K. Tang, and H.-Y. Shum, "Drag-and-drop pasting," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 631–637, 2006.
- [14] M. W. Tao, M. K. Johnson, and S. Paris, "Error-tolerant image compositing," in *Proc. ECCV*, 2010, pp. 31–44.
- [15] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 2672–2680.
- [16] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proc. CVPR*, Jun. 2016, pp. 2536–2544.
- [17] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. CVPR*, Jul. 2017, pp. 1125–1134.
- [18] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. ICCV*, Oct. 2017, pp. 2223–2232.
- [19] M. Mathieu, C. Couprie, and Y. LeCun, "Deep multi-scale video prediction beyond mean square error," in *Proc. ICLR*, 2016, pp. 1–14.
- [20] X. Liang, L. Lee, W. Dai, and E. P. Xing, "Dual motion GAN for future-flow embedded video prediction," in *Proc. ICCV*, Oct. 2017, pp. 1744–1752.
- [21] X. Gong, H. Huang, L. Ma, F. Shen, W. Liu, and T. Zhang, "Neural stereoscopic image style transfer," in *Proc. ECCV*, 2018, pp. 56–71.
- [22] M. Li, H. Huang, L. Ma, W. Liu, T. Zhang, and Y. Jiang, "Unsupervised image-to-image translation with stacked cycle-consistent adversarial networks," in *Proc. ECCV*, 2018, pp. 184–199.
- [23] W. Xiong, W. Luo, L. Ma, W. Liu, and J. Luo, "Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2364–2373.
- [24] H. Zhu, X. Peng, V. Chandrasekhar, L. Li, and J.-H. Lim, "DehazeGAN: When image dehazing meets differential programming," in *Proc. IJCAI*, 2018, pp. 1–7.
- [25] K. Zhang, W. Luo, Y. Zhong, L. Ma, W. Liu, and H. Li, "Adversarial spatio-temporal learning for video deblurring," *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 291–301, Jan. 2019.
- [26] W. Ren *et al.*, "Deep video dehazing with semantic segmentation," *IEEE Trans. Image Process.*, vol. 28, no. 4, pp. 1895–1908, Apr. 2019.
- [27] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb, "Learning from simulated and unsupervised images through adversarial training," in *Proc. CVPR*, Jul. 2017, pp. 2107–2116.
- [28] M. Lang, O. Wang, T. Aydin, A. Smolic, and M. Gross, "Practical temporal consistency for image-based graphics applications," *ACM Trans. Graph.*, vol. 31, no. 4, 2012, Art. no. 34.
- [29] N. Bonneel, K. Sunkavalli, S. Paris, and H. Pfister, "Example-based video color grading," *ACM Trans. Graph.*, vol. 32, no. 4, 2013, Art. no. 39.
- [30] T. O. Aydin, N. Stefanoski, S. Croci, M. Gross, and A. Smolic, "Temporally coherent local tone mapping of HDR video," *ACM Trans. Graph.*, vol. 33, no. 6, 2014, Art. no. 196.
- [31] N. Bonneel, J. Tompkin, K. Sunkavalli, D. Sun, S. Paris, and H. Pfister, "Blind video temporal consistency," *ACM Trans. Graph.*, vol. 34, no. 6, 2015, Art. no. 196.
- [32] M. Ruder, A. Dosovitskiy, and T. Brox, "Artistic style transfer for videos," in *Proc. German Conf. Pattern Recognit.*, 2016, pp. 26–36.
- [33] G. Ye, E. Garces, Y. Liu, Q. Dai, and D. Gutierrez, "Intrinsic video and applications," *ACM Trans. Graph.*, vol. 33, no. 4, 2014, Art. no. 80.
- [34] A. Gupta, J. Johnson, A. Alahi, and L. Fei-Fei, "Characterizing and improving stability in neural style transfer," in *Proc. ICCV*, Oct. 2017, pp. 4067–4076.
- [35] H. Huang *et al.*, "Real-time neural style transfer for videos," in *Proc. CVPR*, Jul. 2017, pp. 783–791.
- [36] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 568–576.
- [37] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 1933–1941.
- [38] H. Zhu *et al.*, "YouTube: Searching action proposal via recurrent and static regression networks," *IEEE Trans. Image Process.*, vol. 27, no. 6, pp. 2609–2622, Jun. 2018.
- [39] L. Bao, B. Wu, and W. Liu, "CNN in MRF: Video object segmentation via inference in a CNN-based higher-order spatio-temporal MRF," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5977–5986.
- [40] S. Xu, D. Liu, L. Bao, W. Liu, and P. Zhou, "MHP-VOS: Multiple hypotheses propagation for video object segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 314–323.
- [41] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.
- [42] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 DAVIS challenge on video object segmentation," 2017, *arXiv:1704.00675*. [Online]. Available: <https://arxiv.org/abs/1704.00675>
- [43] A. Dosovitskiy *et al.*, "FlowNet: Learning optical flow with convolutional networks," in *Proc. ICCV*, Dec. 2015, pp. 2758–2766.
- [44] P. Weinzaepfel, J. Revaud, Z. Harchaoui, and C. Schmid, "DeepFlow: Large displacement optical flow with deep matching," in *Proc. ICCV*, Dec. 2013, pp. 1385–1392.
- [45] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid, "EpicFlow: Edge-preserving interpolation of correspondences for optical flow," in *Proc. CVPR*, Jun. 2015, pp. 1164–1172.

- [46] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele, "Lucid data dreaming for video object segmentation," 2017, *arXiv:1703.09554*. [Online]. Available: <https://arxiv.org/abs/1703.09554>
- [47] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *Proc. CVPR*, Jul. 2017, pp. 2462–2470.
- [48] A. Criminisi, P. Pérez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1200–1212, Sep. 2004.
- [49] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.
- [50] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. ICCV*, Oct. 2017, pp. 2794–2802.
- [51] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial network," in *Proc. ICML*, 2017, pp. 214–223.
- [52] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial transformer networks," in *Proc. NIPS*, 2015, pp. 2017–2025.
- [53] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. ECCV*, 2016, pp. 694–711.
- [54] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, 2010, pp. 1–8.
- [55] S. Abu-El-Haija *et al.*, "YouTube-8M: A large-scale video classification benchmark," 2016, *arXiv:1609.08675*. [Online]. Available: <https://arxiv.org/abs/1609.08675>
- [56] (2017). *Video Free Footage*. [Online]. Available: <http://www.videonet.net>
- [57] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015, pp. 1–15.
- [58] L. Maystre and M. Grossglauser, "Fast and accurate inference of Plackett–Luce models," in *Proc. NIPS*, 2015, pp. 172–180.



Hao-Zhi Huang received the bachelor's degree from Tsinghua University, Beijing, China, in 2012, and the Ph.D. degree from the Department of Computer Science and Technology, Tsinghua University, in 2017. He is currently a Senior Researcher with the Tencent AI Lab. His current research interests include image and video editing, generative adversarial networks, and human pose estimation.



Sen-Zhe Xu received the bachelor's degree from the Huazhong University of Science and Technology, Wuhan, China, in 2015. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Technology, Tsinghua University. His research interests include image and video processing and deep learning.



Jun-Xiong Cai received the B.S. degree in software engineering from Jilin University, China, in 2015. He is currently pursuing the Ph.D. degree in computer science with Tsinghua University. His research interests include computer graphics, geometry processing, deep learning, and robot programming.



Wei Liu (SM'19) received the Ph.D. degree in EECS from Columbia University, New York, NY, USA. He was a Research Scientist with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. He is currently a Distinguished Scientist with the Tencent AI Lab and the Director of the Computer Vision Center. He has long been devoted to research and development in the fields of machine learning, computer vision, information retrieval, big data, and so on. He is an Elected Member of the International Statistical Institute (ISI). His research works win a number of awards and honors, such as the 2011 Facebook Fellowship, the 2013 Jury Award for best thesis of Columbia University, the 2016 and 2017 SIGIR Best Paper Award Honorable Mentions, and the 2018 "AI's 10 To Watch" Honor. He currently serves on the Editorial Boards for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE ACCESS, *Pattern Recognition*, *Neurocomputing*, the *Journal of Visual Communication and Image Representation*, and the *International Journal of Applied Intelligence*.



Shi-Min Hu received the Ph.D. degree from Zhejiang University in 1996. He is currently a Professor with the Department of Computer Science and Technology, Tsinghua University. He has published over 100 papers in journals and refereed conferences. His research interests include digital geometry processing, video processing, rendering, computer animation, and computer aided geometric design. He is the Editor-in-Chief of *Computational Visual Media* (Springer), and on the editorial boards of several journals, including *Computer Aided Design* and *Computer and Graphics*.