

AUTOMATIC PORTRAIT VIDEO MATTING VIA CONTEXT MOTION NETWORK

Anonymous authors

Paper under double-blind review




Figure 1: Our automatic portrait video matting method does not require extra inputs. Most state-of-the-art matting methods rely on semantic segmentation methods to automatically generate the trimap. Their performance is compromised due to the lack of temporal information. Our method exploits semantic information as well as temporal information from optical flow and produces high-quality results.

ABSTRACT

Automatic portrait video matting is an under-constrained problem. Most state-of-the-art methods only exploit the semantic information and process each frame individually. Their performance is compromised due to the lack of temporal information between the frames. To solve this problem, we explore the optical flow between video frames for the automatic portrait video matting. Specifically, we propose the context motion network to leverage semantic information and motion information. To capture the motion information, we estimate the optical flow and design a context-motion updating operator to integrate features between frames recurrently. Our experiments show that our network outperforms state-of-the-art matting methods significantly on the Video240K SD dataset.

1 INTRODUCTION

Portrait video matting aims to recover the foreground portraits from a portrait video. It has attracted extensive interest due to its wide applications in visual effects, video editing, and video meeting. It is an under-constrained problem.

To solve the matting problem, most existing methods rely on extra inputs to provide external information about the foreground or background areas. Trimap-based methods Xu et al. (2017); Hou & Liu (2019); Lu et al. (2019); Li & Lu (2020); Forte & Pitié (2020) require a trimap image as the extra input, where users specify known foreground and background areas, as well as an undefined area. These methods recover alpha matte values in the undefined area by propagating the alpha values from known area. These methods achieve promising performance. However, obtaining accurate trimaps often requires a vast amount of user effort. Alternatively, background-based methods Sengupta et al. (2020); Lin et al. (2020) assume that the background is relatively static, and users can capture the background. The matting problem is reduced to estimate the matte values with a known background. These methods adopt an end-to-end network to estimate the alpha values directly and produce high-quality results. However, it is challenging to capture background for videos with moving backgrounds, which impedes their applications to more general scenarios.

Our research is inspired by recent portrait image matting without any extra inputs Shen et al. (2016); Zhu et al. (2017a). These methods incorporate the semantic information into the matting network to automatically estimate the alpha maps for portraits. For instance, Shen *et al.* Shen et al. (2016) employs a FCN Long et al. (2015) as a part of the whole network to estimate the trimap. They estimate the matte values based on the trimap and achieve promising results. However, these methods are designed for images. Their performance on videos might be compromised due to the lack of temporal information. For instance, as shown in Figure 1, the hand in the semantic segmentation is inaccurate, while the optical flow can provide useful temporal information for the hand. Further-

more, adopting these methods to each frame individually sometimes leads to undesirable temporal incoherence, such as flickering.

Our idea for portrait video matting is to make use of not only semantic information but also motion information from optical flow. There are several advantages of introducing motion information to the portrait video matting. First, the motions of foreground and background are typically different. It will be easier to find the edges between foreground and background. Second, if the foreground moves across the background, some portions of the background will be covered or revealed. We can construct a more complete background to matte against.

We accordingly developed a context motion network dedicated to portrait video matting. Our network takes a sequence of frames as inputs and estimates the alpha maps and foregrounds for each frame. Specifically, our network contains four parts: the feature extraction network, the optical flow estimator, the context motion updating operator, and the upsampler. We first extract the context features for each frame and estimate the optical flow between consecutive frames. The optical flow and context features are fed into the context motion updating operator. To get the motion features, the context motion updating operator backwarps features from neighboring frames using optical flow and calculate the correlation between the current features and the backwarped features. Our operator also encodes the optical flow directly. We employ ConvGRUs to recurrently integrate features from consecutive frames. We then feed the resulting features to upsamplers to predict alpha maps and foreground. Our network is trained in an end-to-end manner. Our experiments show that our method outperforms the state-of-the-art trimap-based and background-based methods significantly.

This paper contributes to the research on automatic portrait video matting by leveraging semantic information and temporal information. First, we propose a novel context motion network. It is the first deep learning based method to utilize optical flow to capture the motion information in the automatic portrait video matting. Second, we design a context motion updating operator to fuse the context and temporal information and recurrently update features across consecutive frames. Third, the proposed method achieves state-of-the-art performances compared to the trimap-based and background-based methods.

2 RELATED WORK

Matting is a classic computer vision problem. For an image \mathbf{I} , it assumes that \mathbf{I} is a linear composition of a foreground \mathbf{F} and a background \mathbf{B} based on an alpha map $\alpha \in [0, 1]$,

$$\mathbf{I} = \alpha\mathbf{F} + (1 - \alpha)\mathbf{B}, \quad (1)$$

where \mathbf{F} , \mathbf{B} , and α are unknown. Matting aims to recover the foreground object, which is an under-constrained problem. In this part, we discuss the most relevant works on image matting and video matting. Due to the space limit, we refer readers to a comprehensive survey article Wang et.al. Wang et al. (2008).

Image matting. To solve the matting problem, most of the current methods rely on the extra inputs for external information. Trimap-based methods require a user-provided trimap or stokes which specify the foreground, background, and unknown areas. To recover the matte values in the unknown area, there are three categories of trimap-based matting methods: color sampling methods, propagation methods, and deep learning methods. Color sampling methods Chuang et al. (2001); Wang & Cohen (2005); He et al. (2011) are based on the observation that neighboring pixels’ alpha values are often similar if their colors are similar. However, the performance of color sampling methods declines sharply when the color distributions of foreground and background have a large overlap. Propagation methods Aksoy et al. (2017); Chen et al. (2013a;b); He et al. (2010); Lee & Wu (2011); Levin et al. (2008a;b); Sun et al. (2004); Levin et al. (2008a); Aksoy et al. (2017); Chen et al. (2013a) propagate alpha values from known to unknown pixels through various affinities between neighboring pixels. However, it is challenging to obtain accurate affinities because of color ambiguity. Deep learning methods use a CNN and estimate the alpha map from the source images Shen et al. (2016); Xu et al. (2017); Cho et al. (2016; 2019); Hou & Liu (2019); Lu et al. (2019); Tang et al. (2019); Li & Lu (2020) or finetune the results from traditional methods Cho et al. (2016; 2019). Instead of using hand-crafted features, these methods learned the features which are more robust to the challenging scenarios. These methods achieved promising results. However, trimap-based methods require an amount of user effort in creating the trimap. The extensive requirements of accurate trimaps pose a challenge, especially when running these methods on videos.

Alternatively, background-based methods Sengupta et al. (2020); Lin et al. (2020) assume that the background image is relatively static, and users can capture the background image. The matting problem is reduced to estimate the foreground objects with a known background. These methods take the image as well as the background image as inputs and estimate the alpha mattes with an end-to-end trained network, which have shown particularly effective and high efficiency when the background is relatively static. However, their performance drops sharply when it has a large mismatch between the user-captured background and the real background. Furthermore, for the videos with a moving background, it is difficult to capture the accurate background.

The semantic-based portrait matting methods Shen et al. (2016); Zhu et al. (2017a) leverage the semantic information to estimate the matte values for the portrait matting without any user interaction. For instance, Shen *et al.* accordingly proposed an end-to-end trained network Shen et al. (2016). It first estimated the trimap. Then they designed a matting layer to estimate the matte values based on the trimap. Their method achieved promising results for images. Compared to these methods, our method improves the accuracy of automatic matting with temporal information.

Recently, some methods Li et al. (2020); Qiao et al. (2020) estimate the matte values for any objects without any extra inputs. However, the generalization capability of these methods is compromised.

Video matting. Most traditional video matting methods solve the video matting problem in two steps Chuang et al. (2002); Apostoloff & Fitzgibbon (2004); Lee et al. (2010); Bai et al. (2011). In the first step, these methods employ semantic segmentation methods to estimate a binary segmentation image of each frame independently. Based on the binary segmentation, these methods generate trimaps for each frame by performing spatio-temporal optimizations. In the second step, these methods apply image matting algorithms to estimate the matte values for each frame using trimaps generated in the first step. While these methods have achieved promising results for some scenes, the quality and stability of predictions from these methods decline when they are applied to challenging scenarios because these methods heavily rely on hand-crafted features. Our method addresses this challenging problem by learning the features with an end-to-end trained network.

Recently, there are deep-learning based methods for video matting, including Ke et al. (2020); Zhang et al. (2021); Sun et al. (2021). For instance, Ke et al. (2020) first estimates the alpha maps for each frame and employs a self-supervised strategy to preserve the temporal coherence. Zhang et al. (2021) employs an attention-based temporal aggregation module to recurrently predicts the alpha maps for videos. Sun et al. (2021) fuses the features from multiple frames to estimate the alpha maps for the target frame. There is also a concurrent work by Lin et al. (2021) that also applies temporal information for video matting. But it targets efficient video matting, while our method focuses on the high quality video matting. Furthermore, most of these works leverage recurrent network to capture the temporal information. Compared to these methods, our method use optical flow to capture the motion information.

Optical flow. Optical flow measures the per-pixel motions between video frames. It has been widely used in high-level video tasks, including video in-painting Xu et al. (2019), moving-object detection Yang et al. (2019), video recognition Zhu et al. (2017b), video segmentation Cheng et al. (2017), and video frame interpolation Niklaus & Liu (2020). We method employ the optical flow to measure the motion in the video matting task.

3 METHOD

Our method takes a sequence of portrait frames $\{\mathbf{I}_t\}_{t=1}^N$ as inputs and aims to estimate their corresponding alpha maps $\{\alpha_t\}_{t=1}^N$ as well as their foregrounds $\{\mathbf{F}_t\}_{t=1}^N$, where N indicates the number of frames. Our network leverages not only semantic information but also **temporal** information.

As shown in Figure 2, we develop a new deep network architecture for automatic portrait video matting. It recurrently estimates the alpha maps and foreground images. Our network can be distilled into four parts: (1) feature extraction, (2) optical flow estimation, (3) context-motion updating operator, and (4) upsampler. All parts are differentiable, and we train the network in an end-to-end manner. Given an input frame \mathbf{I}_t , our network first extracts the context features \mathbf{Fea}_t from an encoder pretrained on the image classification task. We capture the motion information by estimating the backward optical flow $\mathbf{f}_{t \rightarrow t-1}$ from the current frame \mathbf{I}_t to the previous frame \mathbf{I}_{t-1} by the PWC-Net Sun et al. (2018). To fuse the context information and the motion information, we develop a

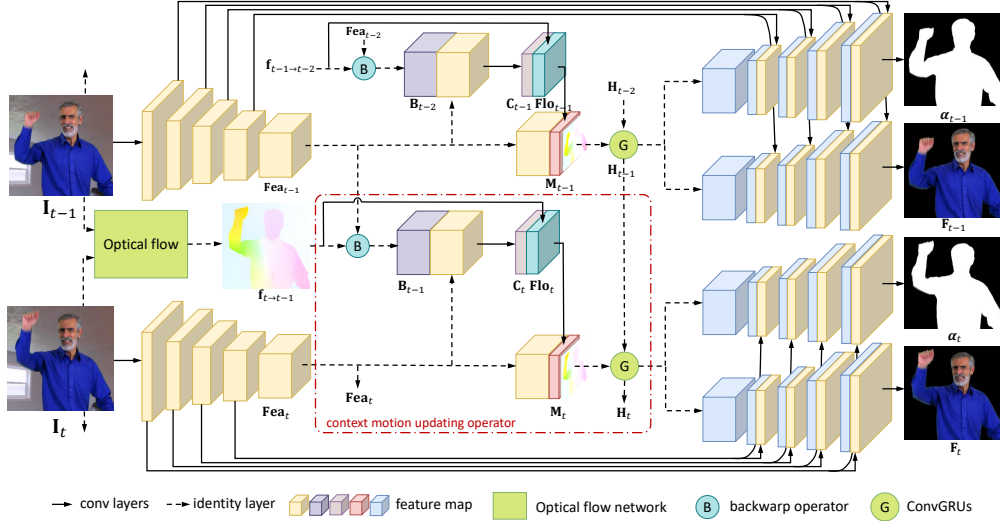


Figure 2: The architecture of our network. Our network takes a sequence of frames as inputs and estimate the alpha maps and foreground image for each frame.

context-motion updating operator, which integrates context features, optical flow as well as the features from the previous frame. Finally, we upscale the resulting features H_t and predict the alpha maps α_t and the foreground F_t . Below we describe the network in detail.

Feature extraction. We extract the features from the input images using a convolutional network. The feature encoder network is applied to all the input frames. It maps the input images to feature maps at a lower resolution as follows

$$\mathbf{Fea}_t = g_{encoder}(\mathbf{I}_t), \quad (2)$$

where $g_{encoder}$ indicates the encoder network. We adopt the modified ResNet50 He et al. (2016) pretrained on the ImageNet dataset. Our encoder outputs features at 1/16 resolution of the source frame as we empirically find that the large receptive field is critical for the semantic information. Aside from the context features, we also extract the multi-scale intermediate features from the encoder to capture the fine structures. We use 3×3 convolutional layers to map these intermediate feature maps to 32 channels. The weights of the feature extraction network are shared for all the input frames.

Optical flow estimation. We calculate the optical flow between the current frame I_t and the previous frame I_{t-1} , which indicates the per-pixel motion between these two frames. Inspired by the success of state-of-the-art optical flow networks Sun et al. (2018); Zhao et al. (2020); Teed & Deng (2020), we leverage the PWC-Net Sun et al. (2018) as a sub-network in our network to estimate the optical flow. It takes two consecutive video frames as input and output the backward optical flow as follows,

$$\mathbf{f}_{t \rightarrow t-1} = g_{flow}(\mathbf{I}_t, \mathbf{I}_{t-1}), \quad (3)$$

where g_{flow} indicates the optical flow network. $\mathbf{f}_{t \rightarrow t-1} \in \mathbb{R}^{H \times W \times 2}$ indicates the backward optical flow. We initialize the the optical flow subnetwork using weights pretrained on the Sintel dataset Butler et al. (2012). During our training, we train the whole network in an end-to-end manner.

Context-motion updating operator. Context-motion updating operator is designed to fuse semantic information and temporal information. The context-motion updating operator is recurrent. The features from the previous frames are utilized in the context-motion updating operator to provide temporal information.

Given the frame feature \mathbf{Fea}_t , \mathbf{Fea}_{t-1} as well as the backward optical flow $\mathbf{f}_{t \rightarrow t-1}$, our goal is to capture the motion information. Inspired by the design of PWCNet Sun et al. (2018), we backwarp the features \mathbf{Fea}_{t-1} from previous frame to the current frame by

$$\mathbf{B}_{t-1} = g_{backwarp}(\mathbf{Fea}_{t-1}, \mathbf{f}_{t \rightarrow t-1}), \quad (4)$$

where $g_{backwarp}(\cdot)$ indicates the operator of backwarping. \mathbf{B}_{t-1} indicates the backwarped features. We calculate the correlation between the current feature \mathbf{Fea}_t and the backwarped feature \mathbf{B}_{t-1}

$$\mathbf{C}_t = g_{corr}([\mathbf{Fea}_t, \mathbf{B}_{t-1}]), \quad (5)$$

where $[\cdot]$ indicates the concatenating operator. g_{corr} indicates the operator to calculate the correlation. Our network adopts two 3×3 convolutional layers with 32 channels.

Our network also encodes the optical flow to capture the motion information from optical flow directly. We first bilinearly downsample the optical flow size to the size of \mathbf{Fea}_t . The optical flow is encoded as

$$\mathbf{Flo}_t = g_{flow.encode}(\mathbf{f}_{t \rightarrow t-1}), \quad (6)$$

where $g_{flow.encode}$ indicates the operator for optical flow encoding. Our network adopts two 7×7 convolutional layers with 32 channels.

We fuse the correlation \mathbf{C}_t and optical flow feature \mathbf{Flo}_t to get the motion features

$$\mathbf{M}_t = g_{motion}([\mathbf{C}_t, \mathbf{Flo}_t]), \quad (7)$$

where $g_{motion}(\cdot)$ indicates the operator for motion encoding. Our network adopts a 3×3 convolutional layer with 62 channels.

By concatenating the semantic feature \mathbf{Fea}_t , motion feature \mathbf{M}_t , as well as the optical flow $\mathbf{f}_{t \rightarrow t-1}$, our method gets the fused feature \mathbf{Fus}_t ,

$$\mathbf{Fus}_t = [\mathbf{Fea}_t, \mathbf{M}_t, \mathbf{f}_{t \rightarrow t-1}]. \quad (8)$$

Inspired by RAFT Teed & Deng (2020), our network adopts separable ConvGRUs to leverage the features from the previous frame. It replaces the fully connected layers in the GRU with separable convolutions. It can be represented as

$$\mathbf{z}_t = g_z([\mathbf{H}_{t-1}, \mathbf{Fus}_t]), \quad (9)$$

$$\mathbf{r}_t = g_r([\mathbf{H}_{t-1}, \mathbf{Fus}_t]) \quad (10)$$

$$\widetilde{\mathbf{H}}_t = g_h([\mathbf{r}_t \cdot \mathbf{H}_{t-1}, \mathbf{Fus}_t]) \quad (11)$$

$$\mathbf{H}_t = (1 - \mathbf{z}_t) \cdot \mathbf{H}_{t-1} + \mathbf{z}_t \cdot \widetilde{\mathbf{H}}_t, \quad (12)$$

where $g_z(\cdot)$, $g_r(\cdot)$, and $g_h(\cdot)$ indicate the gated activation units. Specifically, $g_z(\cdot)$ and $g_r(\cdot)$ employ $sigmoid(\cdot)$ as their activation function. $g_h(\cdot)$ uses $tanh(\cdot)$. Our network adopts two ConvGRUs. In the first ConvGRU, $g_z(\cdot)$, $g_r(\cdot)$ and $g_h(\cdot)$ are one 1×5 convolutional layer. And in the second ConvGRU, they are 5×1 convolutional layers. It can increase the receptive field while keeps the model size small.

Upsampler. As shown in Figure 2, our network contains two decoders: one for the alpha maps and the other for the foreground. They have the identity network architecture. We first bilinearly upsample the features by a factor of 2. Our network concatenates the upsampled features with the intermediate features from the encoder. Our network fuses them by two 3×3 convolutional layers with 32 channels. We repeat the process until the resolution of upsampled features is the same as the input frame. The last layer of the alpha decoder has two convolutional layers. We use the activation function $sigmoid()$ for the second convolutional layer. The last layer of the foreground decoder also has two convolutional layers. The second convolutional layer direct outputs the foreground image without activation function.

Loss functions. For each frame of the input video, we compute the loss over the alpha maps and the foreground images. For the alpha maps, we use the standard ℓ_1 to measure the difference between the predicted alpha maps α_t and the ground truth $\hat{\alpha}_t$ as follows

$$\mathcal{L}_{1,t}^\alpha = \|\alpha_t - \hat{\alpha}_t\|_1. \quad (13)$$

Following Context Matting Hou & Liu (2019), we also employ the Laplacian loss for the alpha maps. Instead of measuring the difference from the image directly, it measures the differences of two Laplacian pyramid as follows,

$$\mathcal{L}_{lap,t}^\alpha = \sum_{i=1}^5 2^{i-1} \|L^i(\alpha_t) - L^i(\hat{\alpha}_t)\|_1, \quad (14)$$




Figure 3: Visual example of backwarping. Backwarping can help construct a more complete background. Please note that our method backwarps the feature maps instead of images.



Figure 4: Examples from our background video dataset. It contains various indoor scenarios.

where $L^i(\cdot)$ indicates the i^{th} level of the Laplacian pyramid.

For the foreground image, we also employ the standard ℓ_1 to measure the difference between the predicted foreground image \mathbf{F}_t and $\hat{\mathbf{F}}_t$. We only calculate the loss where the foreground is visible

$$\mathcal{L}_{1,t}^{fg} = \|\mathbb{I}(\hat{\alpha}_t > 0) \cdot (\mathbf{F}_t - \hat{\mathbf{F}}_t)\|_1, \quad (15)$$

where $\mathbb{I}(\cdot)$ indicates a binary function whose value is 1 if the condition is true and 0 otherwise.

We get the total loss as

$$\mathcal{L}_{total} = \sum_{t=1}^N \mathcal{L}_{1,t}^{\alpha} + \mathcal{L}_{lap,t}^{\alpha} + 0.1\mathcal{L}_{1,t}^{fg}, \quad (16)$$

where N indicates the length of the input video.

Training. We use PyTorch to train our neural network. The feature extraction network is initialized with the official weights pretrained on the ImageNet dataset. And the optical flow network is initialized with the weights pretrained on the Sintel dataset. The other modules are initialized with random weights. We adopt AdamW Loshchilov & Hutter (2017) as our optimizer. Our initial learning rate is set to 10^{-4} . We use the ‘‘OneCycleLR’’ learning rate policy Smith & Topin (2019) in training, which we empirically find it converges faster than other policies. Furthermore, we clip gradients in the range $[-1, 1]$.

We train our network for 300k iterations on a single NVIDIA Titan RTX GPU. In the first 150k iterations, we freeze the parameters of the optical flow network. We empirically find that freezing the parameters of the optical flow network in the early stages can make the training process more stable. Our batchsize is set to 4. In the following 150k iterations, we unfreeze the parameters of the optical flow network and train the network in an end-to-end manner. We set the batchsize to 2 due to the GPU memory limitation. We set the length of frames as 3 for training. Furthermore, our network freezes the parameters of batch norm layers during the whole training process.

Dataset. We train our network using the video portrait matting dataset, Video240K SD, shared by Lin et al. (2020). This dataset contains 484 training videos with the corresponding alpha maps and the foreground images, including 479 videos for training and 5 videos for testing. The resolution of videos is 768×432 . The Video240K SD dataset only releases the background images, while our method is designed for the videos. We collected 82 indoor videos as the background videos. These videos are under Common Creative Licenses from YouTube. The resolution of these videos is higher than or equal to 1080p. Most videos last more than 10 minutes, and the frame rate is higher than or equal to 24fps. We randomly extracted frame sequences from these videos. Each frame sequence contains 50 consecutive frames. We also downsampled the frames to 720p. We manually removed the sequences with cut shots, visual effects, and extremely fast motion. We also removed the sequences which contain humans. This culling process reduced the total number of sequences to 5636, consisting of 281,800 individual video frames with a typical resolution of 1080×720 . We split them into two parts: 5586 sequences from 68 videos as the training background and 50 sequences from 14 videos as the testing background. There is no overlap among them. As shown in Figure 4, our dataset covers various indoor scenarios. For training, we created the training samples by compositing the foreground video onto a randomly selected background sequences. For the testing, we composite each testing foreground video onto 10 randomly selected background videos. There are 50 testing videos in total.

Data augmentation. As discussed in Hou & Liu (2019), the foreground videos and background videos often contain different artifacts. Directly training the network on the composited videos without augmentation might compromise the generalization capability of the trained network. Following Hou & Liu (2019), we leveraged the JPEG noise to the composited videos. We keep the

Method	Extra input	Alpha				Foreground
		SAD	MSE	Grad	Conn	MSE
Context Matting Hou & Liu (2019)	\mathbf{T}_{seg}	5.39	13.38	18.09	5.34	5.17
	\mathbf{T}_{gt}	1.41	1.78	4.24	1.31	4.26
Index Matting Lu et al. (2019)	\mathbf{T}_{seg}	4.32	10.32	13.03	4.22	-
	\mathbf{T}_{gt}	1.47	1.57	2.74	1.28	-
GCA Matting Li & Lu (2020)	\mathbf{T}_{seg}	4.33	10.91	14.14	4.24	-
	\mathbf{T}_{gt}	1.49	2.32	3.29	1.37	-
FBA Matting Forte & Pitić (2020)	\mathbf{T}_{seg}	4.17	10.18	12.36	4.10	2.50
	\mathbf{T}_{gt}	1.02	0.87	1.84	0.91	2.25
BGMV2 Lin et al. (2020)	\mathbf{B}_{first}	83.56	247.03	16.79	-	3.79
	\mathbf{B}_{gt}	1.07	1.41	2.47	0.94	4.01
MODNet Ke et al. (2020)	-	3.62	8.47	4.11	3.42	9.63
RVM Lin et al. (2021)	-	1.06	0.89	1.85	0.96	3.15
Ours	-	0.44	0.27	0.38	0.26	0.88

Table 1: Comparison on the Video240K SD dataset. \mathbf{T}_{seg} indicates the trimap generated from the semantic segmentation method Chen et al. (2018). \mathbf{T}_{gt} indicates the trimap generated from the ground truth alpha maps. \mathbf{B}_{first} indicates the background image from the first frame. \mathbf{B}_{gt} indicates the ground truth background image for each frame in the video.

quality of the composited videos in the range of [70%, 100%]. Aside from the augmentations of images, we also introduced the H264 noise to our composited videos, which is commonly used in video compression. Our method also uses standard augmentations. Specifically, we randomly crop the composited videos to patches of size 384×384 , as the large cropping size is critical for semantic information. We randomly flip the images horizontally. We also employ the color jitter for the foreground video to increase the color diversity, where we set the brightness in a range of [0.85, 1.15], the contrast in the range of [0.85, 1.15], the hue in the range of [0.9, 1.1], and the saturation in the range of [0.7, 1.3]. We also employ the gamma transformation for the alpha maps, in which we set the gamma value in the range of [0.2, 2].

4 EXPERIMENT

We conduct our experiments on the Video240K SD dataset. To compare with the trimap-based methods, we generate two versions of trimaps. The first version \mathbf{T}_{gt} follows the Adobe Matting Dataset Xu et al. (2017), which is generated by dilating the ground truth alpha mattes. For each video, we randomly select the size of dilation in the range of [5, 15]. Since manually creating the trimaps will take a lot of user effort, some users automatically create the trimaps by dilating the results from the semantic segmentation networks. As a consequence, we generate the second version \mathbf{T}_{seg} by dilating the semantic segmentation using the Deeplab V3+ Chen et al. (2018). We set the size of dilation as 10 pixels.

Evaluation metrics. For the alpha maps, we adopt various metrics, including SAD, MSE, Gradient (Grad), and Connectivity (Conn) Xu et al. (2017). For the foreground videos, we evaluate the MSE of the foreground video. In our experiments, we employ the official evaluation codes from the authors of the Video240K SD dataset Lin et al. (2020). We scale MSE, Grad, Conn by 10^3 , 10^{-3} and 10^{-3} , respectively. Please note that the evaluation codes only measure unknown areas on the trimaps, while our method doesn’t require trimaps. For a fair comparison, we evaluate the whole image for all comparing methods. Besides, Conn metric fails on some results, and some methods do not estimate the foreground image. We use “-” to denote them.

4.1 COMPARISON WITH STATE-OF-THE-ART METHODS

We compare our methods with the state-of-the-art trimap-based methods, including Index Matting Lu et al. (2019), Context Matting Hou & Liu (2019), GCA Matting Li & Lu (2020) and FBA Matting Forte & Pitić (2020), and the state-of-the-art background-based method BGMV2 Lin et al. (2020). We also compare our method with the state-of-the-art video matting methods, including MODNet Ke et al. (2020) and RVM Lin et al. (2021). We used the code / model shared by their authors.

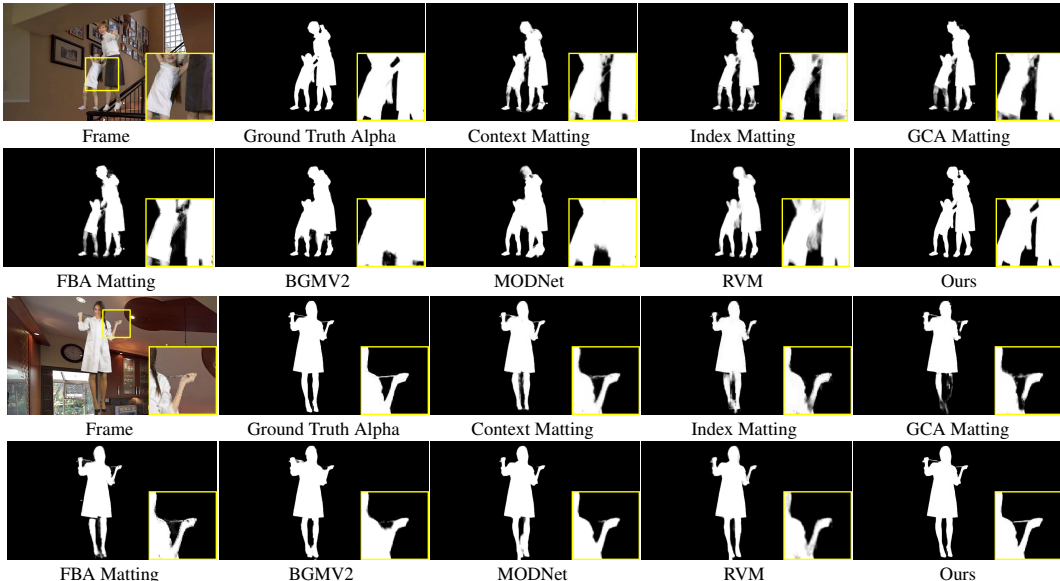


Figure 5: Visual comparison on the Video240K SD dataset. Our method generates more plausible results. We adopt the trimaps \mathbf{T}_{gt} generated from the ground truth alpha maps for the trimap-based methods, including Index Matting Lu et al. (2019), Context Matting Hou & Liu (2019), GCA Matting Li & Lu (2020) and FBA Matting Forte & Pitié (2020). For the BGMV2 Lin et al. (2020), we adopt the ground truth background \mathbf{B}_{gt} for it.

Table 1 reports the quantitative results on the Video240K SD dataset. For each trimap-based method, we report their results on the trimap \mathbf{T}_{seg} which is generated from the semantic segmentation, and \mathbf{T}_{gt} which is generated from the ground truth alpha maps. For the background-based method, we assume that users can capture the background image for the first frame and adopt it as the background for the whole video, whose result is denoted as \mathbf{B}_{first} . We also report the results with the ground truth background for each frame, which is denoted as \mathbf{B}_{gt} . As shown in Table 1, our method outperforms both trimap-based and background-based approaches by a large margin. Specifically, our method wins 0.58 on SAD, 0.60 on MSE, 1.46 on Grad, and 0.65 on Conn for the alpha maps. Our method wins 1.37 on MSE for the foreground. Figure 5 shows several visual examples on the Video240k SD dataset. Our results contain more fine structures. In the first example, the girl’s hair is our result is more accurate than others. In the second example, the stethoscope and arms in our result preserve more details.

4.2 ABLATION STUDY

In this section, we examine several key components of our method.

Feature extraction network. We examine how the feature extraction network affects the final results. In this experiment, we change the feature extraction network while fixing the other parts of the network. Table 2 reports the results of MobileNetV2 Sandler et al. (2018), ResNet18 He et al. (2016) and ResNet50 He et al. (2016). Compared with other backbones, ResNet50 is deeper and has higher computational complexity. It achieves the best performance. We believe that these improvements come from the more representative information of ResNet50, which is consistent with the finding of the previous semantic segmentation methods Chen et al. (2017; 2018) and image matting methods Forte & Pitié (2020).

Optical flow network. We examine how the optical flow network affects the final results. In this experiment, we use the MobileNetV2 Sandler et al. (2018) as our feature extraction network. We compare two models: one with PWC-Net Sun et al. (2018) as the optical flow network and the other one with MaskFlowNet Zhao et al. (2020) as the optical flow network. Although MaskFlowNet achieves better performance in the optical flow estimation task Zhao et al. (2020), PWC-Net gets better matting results. As shown in Table 3, PWC-Net wins 0.10 on SAD, 0.28 on MSE, 0.14 on Grad, and 0.09 on Conn on the alpha maps. It also wins 0.29 on the foreground MSE. This is

Method	Alpha				Foreground
	SAD	MSE	Grad	Conn	MSE
MobileNetV2	0.63	0.54	0.80	0.41	1.32
ResNet18	0.48	0.27	0.48	0.29	1.51
ResNet50	0.44	0.27	0.38	0.26	0.88

Table 2: The effectiveness of the feature extraction network on the Video240K SD dataset.

consistent with many other works that PWC-Net has great generalization capability Niklaus & Liu (2018; 2020); Bao et al. (2019).

Context motion updating operator. We also investigate the contribution of each component in the context motion updating operator. In this experiment, we adopt the MobileNetV2 Sandler et al. (2018) as the feature extraction and PWC-Net Sun et al. (2018) as the optical flow network. In particular, we study the impact of context motion updating operator by removing it from the whole network. For the model with “Baseline”, we remove the context motion updating operator. For the model with “+ConvGRU”, we only keep the ConvGRU while removing the other components from the context motion updating operator. For the model with “+Motion”, we keep the context motion updating operator. The performance of these models can be found in Table 4. We can find that ConvGRU can improve all the evaluation metrics on the alpha maps while slightly increasing the foreground MSE. Particularly, the motion plays a crucial role, which improves the results by 0.20 on SAD, 0.42 on MSE, 0.25 on Grad, and 0.16 on Conn. It also improves the foreground MSE by 0.51. It verifies our motivation to introduce the motion information into the video matting network.

5 LIMITATION AND FUTURE WORK

There are several limitations and drawbacks of the proposed automatic portrait video matting method.

Memory. Our method employs ResNet50 as our feature extraction network to capture semantic information. It will take too much GPU memory on high-resolution videos, e.g., videos of 4k resolution. Besides, the optical flow network also requires large memory. The high memory requirements impede our training on high-resolution videos.

Optical flow. Our method currently relies on PWC-Net Sun et al. (2018) to estimate the optical flow between frames. However, PWC-Net might not produce reliable optical flow for the challenging scenarios, e.g., videos with large motion. In these videos, our method will fail to produce correct matting results. An interesting future direction is to design an optical flow network specifically for video matting.

6 CONCLUSION

This paper presented an automatic portrait video matting method. We designed a context motion network for this task. Our network takes a sequence of frames as inputs and does not require extra inputs. Our network leverages context information and temporal information. Our method first extracts the context features for each frame. It also extracts the optical flow between consecutive frames. This is followed by a context motion updating operator to fuse the context and motion features recurrently. Our experiments showed that our method is able to generate high-quality alpha maps and foreground images. Our experiments also showed that the context motion updating operator is helpful to generate high-quality results.

REFERENCES

Yagiz Aksoy, Tunc Ozan Aydin, and Marc Pollefeys. Designing effective inter-pixel information flow for natural image matting. In *Proceedings of the IEEE Conference on Computer Vision and*

Method	Alpha				Foreground
	SAD	MSE	Grad	Conn	MSE
MaskFlowNet	0.73	0.82	0.94	0.52	1.61
PWC-Net	0.63	0.54	0.80	0.41	1.32

Table 3: The effectiveness of the optical flow network on the Video240K SD dataset.

Method	Alpha				Foreground
	SAD	MSE	Grad	Conn	MSE
Baseline	0.95	1.45	1.22	0.75	1.66
+ ConvGRU	0.83	0.96	1.05	0.57	1.83
+ Motion	0.63	0.54	0.80	0.41	1.32

Table 4: The effectiveness of the context motion updating operator on the Video240K SD dataset. We add the components cumulatively from top to bottom.

- Pattern Recognition*, pp. 29–37, 2017.
- Nicholas Apostoloff and Andrew Fitzgibbon. Bayesian video matting using learnt image priors. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pp. I–I. IEEE, 2004.
- Xue Bai, Jue Wang, and David Simons. Towards temporally-coherent video matting. In *International Conference on Computer Vision/Computer Graphics Collaboration Techniques and Applications*, pp. 63–74. Springer, 2011.
- Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3703–3712, 2019.
- D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.) (ed.), *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pp. 611–625. Springer-Verlag, October 2012.
- Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018.
- Qifeng Chen, Dingzeyu Li, and Chi-Keung Tang. Knn matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2175–2188, 2013a.
- Xiaowu Chen, Dongqing Zou, Steven Zhiying Zhou, Qinpeng Zhao, and Ping Tan. Image matting with local and nonlocal smooth priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1902–1907, 2013b.
- Jingchun Cheng, Yi-Hsuan Tsai, Shengjin Wang, and Ming-Hsuan Yang. Segflow: Joint learning for video object segmentation and optical flow. In *Proceedings of the IEEE international conference on computer vision*, pp. 686–695, 2017.
- Donghyeon Cho, Yu-Wing Tai, and Inso Kweon. Natural image matting using deep convolutional neural networks. In *European Conference on Computer Vision*, pp. 626–643. Springer, 2016.
- Donghyeon Cho, Yu-Wing Tai, and In So Kweon. Deep convolutional neural network for natural image matting using initial alpha mattes. *IEEE Transactions on Image Processing*, 28(3):1054–1067, 2019.
- Yung-Yu Chuang, Brian Curless, David H Salesin, and Richard Szeliski. A bayesian approach to digital matting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 2, pp. II–II. IEEE, 2001.
- Yung-Yu Chuang, Aseem Agarwala, Brian Curless, David H Salesin, and Richard Szeliski. Video matting of complex scenes. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pp. 243–248, 2002.
- Marco Forte and François Pitié. f, b , alpha matting. *arXiv preprint arXiv:2003.07711*, 2020.
- Kaiming He, Jian Sun, and Xiaoou Tang. Fast matting using large kernel matting laplacian matrices. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 2165–2172. IEEE, 2010.
- Kaiming He, Christoph Rhemann, Carsten Rother, Xiaoou Tang, and Jian Sun. A global sampling method for alpha matting. In *CVPR 2011*, pp. 2049–2056. IEEE, 2011.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

- Qiqi Hou and Feng Liu. Context-aware image matting for simultaneous foreground and alpha estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4130–4139, 2019.
- Zhanghan Ke, Kaican Li, Yurou Zhou, Qiuhua Wu, Xiangyu Mao, Qiong Yan, and Rynson WH Lau. Is a green screen really necessary for real-time portrait matting? *arXiv preprint arXiv:2011.11961*, 2020.
- Philip Lee and Ying Wu. Nonlocal matting. In *CVPR 2011*, pp. 2193–2200. IEEE, 2011.
- Sun-Young Lee, Jong-Chul Yoon, and In-Kwon Lee. Temporally coherent video matting. *Graphical Models*, 72(3):25–33, 2010.
- Anat Levin, Dani Lischinski, and Yair Weiss. A closed-form solution to natural image matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):228–242, 2008a.
- Anat Levin, Alex Rav-Acha, and Dani Lischinski. Spectral matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(10):1699–1712, 2008b.
- Yaoyi Li and Hongtao Lu. Natural image matting via guided contextual attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 11450–11457, 2020.
- Yaoyi Li, Qingyao Xu, and Hongtao Lu. Hierarchical opacity propagation for image matting. *arXiv preprint arXiv:2004.03249*, 2020.
- Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. Real-time high-resolution background matting. *arXiv*, pp. arXiv–2012, 2020.
- Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. Robust high-resolution video matting with temporal guidance, 2021.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Hao Lu, Yutong Dai, Chunhua Shen, and Songcen Xu. Indices matter: Learning to index for deep image matting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3266–3275, 2019.
- Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1701–1710, 2018.
- Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- Yu Qiao, Yuhao Liu, Xin Yang, Dongsheng Zhou, Mingliang Xu, Qiang Zhang, and Xiaopeng Wei. Attention-guided hierarchical structure aggregation for image matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13676–13685, 2020.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018.
- Soumyadip Sengupta, Vivek Jayaram, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. Background matting: The world is your green screen. In *Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Xiaoyong Shen, Xin Tao, Hongyun Gao, Chao Zhou, and Jiaya Jia. Deep automatic portrait matting. In *European Conference on Computer Vision*, pp. 92–107. Springer, 2016.

- Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, pp. 1100612. International Society for Optics and Photonics, 2019.
- Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8934–8943, 2018.
- Jian Sun, Jiaya Jia, Chi-Keung Tang, and Heung-Yeung Shum. Poisson matting. In *ACM SIG-GRAPH 2004 Papers*, pp. 315–321. 2004.
- Yanan Sun, Guanzhi Wang, Qiao Gu, Chi-Keung Tang, and Yu-Wing Tai. Deep video matting via spatio-temporal alignment and aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6975–6984, 2021.
- Jingwei Tang, Yagiz Aksoy, Cengiz Oztireli, Markus Gross, and Tunc Ozan Aydin. Learning-based sampling for natural image matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3055–3063, 2019.
- Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pp. 402–419. Springer, 2020.
- Jue Wang and Michael F Cohen. An iterative optimization approach for unified image segmentation and matting. In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, volume 2, pp. 936–943. IEEE, 2005.
- Jue Wang, Michael F Cohen, et al. Image and video matting: a survey. *Foundations and Trends® in Computer Graphics and Vision*, 3(2):97–175, 2008.
- Ning Xu, Brian Price, Scott Cohen, and Thomas Huang. Deep image matting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2970–2979, 2017.
- Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3723–3732, 2019.
- Yanchao Yang, Antonio Loquercio, Davide Scaramuzza, and Stefano Soatto. Unsupervised moving object detection via contextual information separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 879–888, 2019.
- Yunke Zhang, Chi Wang, Miaomiao Cui, Peiran Ren, Xuansong Xie, Xian-sheng Hua, Hujun Bao, Qixing Huang, and Weiwei Xu. Attention-guided temporal coherent video object matting. *arXiv preprint arXiv:2105.11427*, 2021.
- Shengyu Zhao, Yilun Sheng, Yue Dong, Eric I-Chao Chang, and Yan Xu. Maskflownet: Asymmetric feature matching with learnable occlusion mask. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Bingke Zhu, Yingying Chen, Jinqiao Wang, Si Liu, Bo Zhang, and Ming Tang. Fast deep matting for portrait animation on mobile phone. In *Proceedings of the 25th ACM international conference on Multimedia*, pp. 297–305, 2017a.
- Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2349–2358, 2017b.

A APPENDIX

You may include other additional sections here.