

# One-Trimap Video Matting

Hongje Seong<sup>1,\*</sup>, Seoung Wug Oh<sup>2</sup>, Brian Price<sup>2</sup>,  
Euntai Kim<sup>1</sup>, and Joon-Young Lee<sup>2</sup>

<sup>1</sup> Yonsei University, Seoul, Korea. {hjseong, etkim}@yonsei.ac.kr

<sup>2</sup> Adobe Research, San Jose, CA, USA. {seoh, bprice, jolee}@adobe.com

**Abstract.** Recent studies made great progress in video matting by extending the success of trimap-based image matting to the video domain. In this paper, we push this task toward a more practical setting and propose One-Trimap Video Matting network (OTVM) that performs video matting robustly using only one user-annotated trimap. A key of OTVM is the joint modeling of trimap propagation and alpha prediction. Starting from baseline trimap propagation and alpha prediction networks, our OTVM combines the two networks with an alpha-trimap refinement module to facilitate information flow. We also present an end-to-end training strategy to take full advantage of the joint model. Our joint modeling greatly improves the temporal stability of trimap propagation compared to the previous decoupled methods. We evaluate our model on two latest video matting benchmarks, Deep Video Matting and Video-Matting108, and outperform state-of-the-art by significant margins (MSE improvements of 56.4% and 56.7%, respectively). The source code and model are available online: <https://github.com/Hongje/OTVM>.

**Keywords:** Video matting, Trimap propagation, Alpha prediction

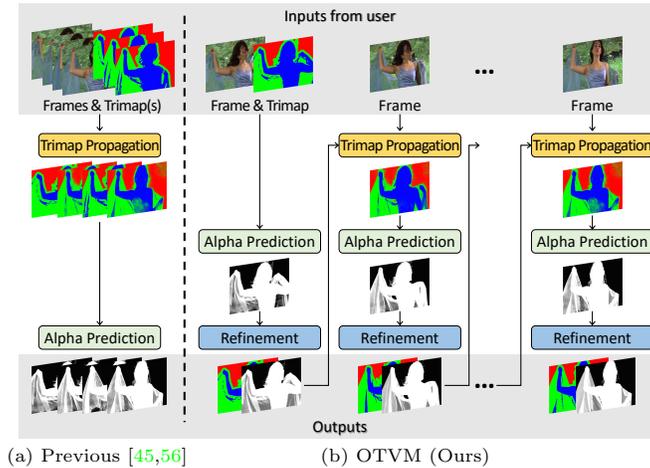
## 1 Introduction

Video matting is the task of predicting accurate alpha mattes from a video. This is an essential step in video editing applications requiring an accurate separation of the foreground and the background layers such as video composition. For each video frame  $I$ , it aims to divide the input color into three components: the foreground color, the background color, and the alpha matte. Formally, for a given pixel, it can be written as,  $I = \alpha F + (1 - \alpha)B$ , where  $F$  and  $B$  are the foreground and background color, and  $\alpha \in [0, 1]$  represents the alpha value. Here, only 3 values ( $I$ ) are known, and the remaining 7 values ( $F$ ,  $B$ , and  $\alpha$ ) are unknown. Given the ill-posed nature of the problem, traditional methods utilize trimaps as additional inputs that indicate pixels that are either solid foreground, solid background, or uncertain. The trimap provides a clue for the target object and effectively improves the stability of the alpha prediction.

Leveraging the latest progress in trimap-based image matting [52,28,14] and mask propagation [54,34,39], recent studies [45,56] successfully developed

---

\* This work was done during an internship at Adobe Research.

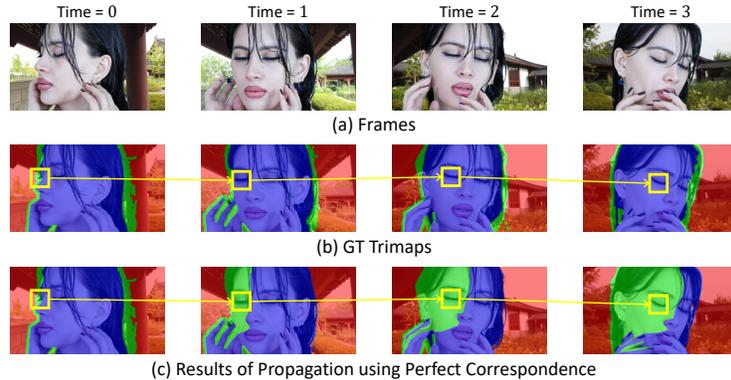


**Fig. 1. The previous decoupled approach vs. Our joint modeling.** (a) Previous video matting methods [45,56] have two decoupled stages: first generate all missing trimaps and then predict the alpha mattes given the trimaps. This approach does not have any interaction between trimap propagation and alpha matting. Since the trimap propagation with no understanding of alpha mattes is often not stable, they require *multiple user-annotated trimaps* to prevent drifting at trimap propagation. (b) In our OTVM, trimap propagation and alpha prediction modules interact each other and the refinement step updates the predictions from the two modules. Our trimap propagation module memorizes and utilizes all information (RGB, trimap, alpha, latent features) to propagate trimaps. It results in accurate and robust predictions, enabling us to perform *one-trimap video matting*.

learning-based video matting techniques. They decouple video matting into two stages, trimap propagation and alpha prediction. They re-purpose the latest mask propagation network [34] to propagate the given trimaps throughout the video, then design alpha prediction networks that take multiple trimaps as input and predict the alpha matte (Fig. 1(a)).

While the decoupled approach effectively simplifies the task, it has a critical limitation as illustrated in Fig. 2. By the nature of the trimap, the unknown region of one frame may be changed into foreground or background and vice versa at a different frame. Therefore, if we propagate a trimap based on visual correspondences without the knowledge for alpha matte [34,5], it may produce inaccurate trimaps and the error can be easily accumulated as shown in Fig. 2(c), leading to the failure of alpha prediction. With this challenge, the existing decoupled methods require multiple user-annotated trimaps to prevent drifting at trimap propagation.

In this paper, we aim to tackle video matting with a single trimap input. To cope with the challenging scenario, we propose One-Trimap Video Matting network (OTVM) that performs trimap propagation and alpha prediction as a joint task, as illustrated in Fig. 1(b). Starting from baseline trimap propagation and alpha prediction modules, we cascade the two modules to alternate trimap propagation and alpha prediction auto-regressively at each frame. We employ the



**Fig. 2. The intrinsic challenge of trimap propagation.** As shown in (b), the same part of an instance can have varying trimap labels at different times, *e.g.*, the label of the left eyelid has changed from the unknown to the foreground. If trimap propagation is conducted using visual correspondences without the knowledge for alpha matte, then it may produce inaccurate trimaps and the error can be easily accumulated as in (c).

space-time memory (STM) network [34] and the FBA matting network [14] as the baseline modules, respectively. To facilitate information flow within the pipeline, we add a refinement module and re-engineer STM accordingly. In addition, we present an end-to-end training pipeline to make OTVM learn the joint task successfully. The major advantage of OTVM is robust trimap propagation that is critical for the practical video matting scenario. Since an alpha matte contains richer information than a trimap, we are able to update the trimap after alpha prediction and this update step prevents error accumulation in the trimap, resulting in robust trimap propagation and accurate alpha prediction.

OTVM produces accurate video mattes even in the challenging one-trimap scenario. We demonstrate that OTVM outperforms previous matting methods with large margins on two latest video matting benchmarks: 56.4% improvement on Deep Video Matting (DVM) [45] and 56.7% improvement on VideoMatting108 [56] in terms of MSE. We also conduct extensive analysis experiments and show that the proposed joint modeling and learning scheme are crucial for achieving robust and accurate video matting results.

## 2 Related Work

**Image Matting.** The image matting task was introduced in [36]. Unlike the image segmentation task that predicts a binary alpha value, the image matting problem aims to predict a high-precision alpha value in a continuous range. Therefore, the matting problem is harder to solve, and most existing works are addressed under some conditions. The most common condition is assuming a human-annotated trimap is given. The trimap is annotated into three different regions: definitely foreground region, definitely background region, and unknown region. The trimap serves to not only reduce the difficulty of the matting problem but also allows some user control over the results.

The traditional sampling-based methods [10,49,15,19,41] determine alpha values in unknown regions by combining sampled foreground and background pixels. Another traditional approach is the propagation-based [22,17,25,26,24,3,4], which propagates foreground and background pixels into unknown regions based on affinity scores. Recently, deep learning-based methods achieved great success in image matting [7,52,33,21,28,14] by training networks with a trimap input.

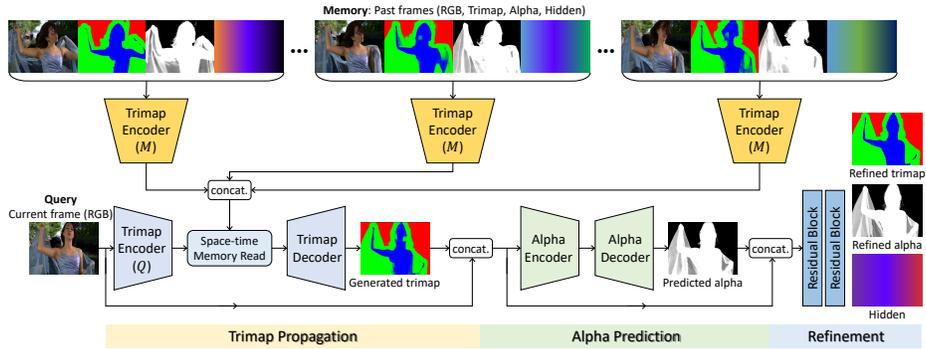
Some approaches try to find good alternatives to human-annotated trimap. Portrait matting [42,57,23] can extract an alpha matte without any external input, but they are only applicable for human subjects. Background matting [38,29] proposes to take complete background information instead of trimap input and predicts high-resolution alpha matte. However, this method is hard to extend to general video matting because it can work only with a near-static background. Mask guided matting [55] proposes to replace the trimap with a coarse binary mask that is more accessible. All image matting methods mentioned above can be extended to video matting by applying frame-by-frame, but the constraint must be met for each frame (*e.g.*, trimap for every frame).

**Video Matting.** Early video matting methods largely extended traditional image matting methods by either extending the propagation temporally [1,8,27,40] or by sampling in other frames [46,40]. While there was some work that would generate trimaps automatically by deriving them from segmentations [16,47] or using interpolation [9] or propagation [46], these were computed independently from the video matting method. Bai *et al.* [2] propagates trimaps based on predicted alpha mattes. Tang *et al.* [48] would use the alpha matte of one frame to help predict the trimap of the next frame, but did so by using the alpha matte to compute a binary segmentation from which a new trimap would be computed.

Recently, large-scale video matting datasets have been proposed [45,56,29] and they fueled video matting researches based on deep learning. Along with a large benchmark dataset, Sun *et al.* [45] proposed a two-stage deep learning-based method by decoupling video matting into trimap generation and alpha matting. They mainly focused on the alpha matting stage to achieve temporally coherent prediction by learning temporal feature alignment [50] and fusion [51]. Concurrently, Zhang *et al.* [56] also approach the problem with a similar motivation. To be specific, [56] proposed a temporal aggregation module based on the guided contextual attention block [28] to maintain temporal consistency during the alpha prediction. They also released a large-scale benchmark dataset, VideoMatting108 [56].

However, both latest learning-based methods [45,56] overlooked the challenge in trimap propagation. Their temporally consistent results can be achieved only if the trimap is accurately generated for every frame. To prevent this issue, additional trimaps need to be manually annotated for several frames which makes the solutions less practical. Although there are trimap-free methods [30,44], those work only on human video and cannot generate user-desired results.

Toward robust and practical trimap-based video matting, we revisit the traditional video matting studies [2,48], where propagating a trimap based on the predicted alpha matte greatly improves the temporal coherency. Inspired by those



**Fig. 3. Overall architecture of OTVM.** Our trimap propagation network is inspired by space-time memory networks [34]. The network predicts a trimap based on the information from the previous frames and predictions that are embedded by the trimap memory encoder. From the given or generated trimap, we initially predict the alpha matte using our alpha prediction network inspired by [14]. Then we refine the generated trimap and predicted alpha matte via two light-weighted residual blocks. The refined trimap, alpha matte, and hidden features (dimension = 16) are fed to the trimap memory encoder so that they can be used for the next frames as a new memory. The framework is trained in an end-to-end manner. We further illustrate the details of each module in the supplementary material.

studies while continuing the success of learning-based video matting, we propose OTVM that jointly learns trimap-based alpha prediction and alpha-based trimap propagation. Our method shows that high-quality and temporally consistent alpha prediction in video is achievable using only a single trimap input.

### 3 Method

The overall architecture of OTVM is illustrated in Fig. 3. Our method aims to perform video matting robustly with only a single trimap. From the given user-provided trimap at the first frame, we sequentially predict a trimap and an alpha matte for every frame in a video sequence. Starting from the user-provided trimap, we first predict the initial alpha matte by feeding the RGB frame and trimap to our alpha prediction network. Then, a lightweight refinement module is followed to correct errors in alpha matte and trimap, resulting in a refined alpha matte and trimap. The refinement module also produces a hidden latent feature map, and all the outputs from the refinement module are encoded as memory by the trimap memory encoder for trimap propagation. On the next frame, our trimap propagation module predicts the trimap by reading relevant information from the memory. This procedure – alpha prediction, refinement, and trimap propagation – is repeated until the end of the video sequence.

Note that recent works [56,45] also generate trimaps and predict alpha mattes, but our approach is completely different from those. In the existing works, the trimap propagation and alpha matting are totally decoupled. They tried to propagate trimap naively (*i.e.* without consideration of the challenge introduced

in Fig. 2), resulting in inaccurate trimap propagation. Therefore, multiple ground-truth (GT) trimaps should be provided to achieve good results. In contrast, our OTVM can extract accurate alpha matte results even with a single human-annotated trimap, thanks to our joint modeling.

### 3.1 Alpha Prediction with Trimap

Given an RGB image and (either a propagated or user-provided) trimap, we predict an alpha matte. Here, we opt for the state-of-the-art image-based matting network, FBA [14], to simplify the problem. The alpha matting network is an encoder-decoder architecture. The alpha encoder first takes a concatenation of the RGB and trimap along the channel dimension as input. Then, the resulting pyramidal features of the alpha encoder are fed into the alpha decoder that produces an alpha matte. To exploit the advantage of the coupled network trained end-to-end, we directly use the soft trimap from the trimap propagation module without binarization when a propagated trimap is given.

We can use any alpha matting network, however, we empirically observe that advanced alpha networks (*e.g.* video alpha networks) make our framework complex and hinder end-to-end training given limited training data for video matting, while the latest image-based alpha networks work surprisingly well as long as a reliable trimap and end-to-end training are provided. Therefore, we take the simple image-based alpha prediction model and focus on developing the joint framework that can reliably propagate the trimaps.

### 3.2 Alpha-Trimap Refinement

In our video matting setting, which takes only one GT trimap, naively propagating the trimap may result in severe drifting as depicted in Fig. 2. To take the advantage of our coupled framework, we have an additional refinement module following the alpha network to provide refined information to the trimap propagation module afterward. The refinement module is light-weighted as it is composed of two residual blocks. The refinement module takes all available information for the current frame: an input RGB frame, the generated trimap, the predicted alpha matte, and the alpha decoder’s latent features. Then, the module produces an updated trimap and a refined alpha matte along with unconstrained hidden features. The hidden features are intended for information that cannot be expressed in the form of trimap and alpha. These features are learnable through end-to-end training. All the outputs of the refinement module will be used for trimap propagation.

### 3.3 Trimap Propagation with Alpha

To propagate trimaps, we repurpose a state-of-the-art video segmentation network, space-time memory network (STM) [34], with important modifications. In the original STM [34], input images and corresponding masks in the past frames are set to the memory, while the image at the current frame is set to the query. Then, the memory and query are embedded through two independent ResNet50 [20] encoders. The embedded memory and query features are fed to the space-time

memory read module. In the module, dense matching is performed and then a value of memory is retrieved based on the matching similarity. The decoder takes the retrieved memory value and query feature and then outputs an object mask. This approach can effectively exploit rich features of the intermediate frames and achieve state-of-the-art performance in video binary segmentation.

STM [34] simply can be extended from binary mask to trimap by increasing the input channel dimension of the memory encoder and the output channel dimension of the decoder. However, there is a fundamental limitation to apply STM directly for trimap propagation. In the binary mask, the foreground region and background region can be estimated by propagating from past binary masks. The trimap, however, cannot be estimated only with propagation because the unknown regions are frequently changed by the view of the foreground object (see Fig. 2) and trimap-only supervision does not provide a consistent clue for the changes. To address this problem, we additionally impose rich cues for generating the trimap into the memory encoder. Since the trimap has been determined by the alpha matte, it effectively helps to learn for trimap generation. We additionally impose a hidden feature extracted from the refinement module. With the hidden features, any errors can be easily propagated backward at training time, resulting in stable training. By imposing those into the memory encoder, we significantly reduce errors that occurred by drifting of the unknown regions.

### 3.4 End-to-End Training

To make OTVM work, it is critical to train the model end-to-end because each module depends on each other’s outputs. However, video matting data are extremely difficult to annotate and existing video supervisions are not sufficient to train the model directly. As a practical solution, we train each module stage-wise then fine-tune the whole network in an end-to-end manner. First, we initialize the trimap propagation and the alpha matting modules with the pretrained weights of off-the-shelf STM [34] and FBA [14], respectively. Specifically, both pretrained models leverage ImageNet [37]. In addition, the STM is trained using image segmentation datasets [13,31,18,43,6] and video object segmentation datasets [35,53]. The FBA is trained on the Adobe Image Matting (AIM) dataset [52]. After the initialization, we pretrain our network modules in three stages (from Stage 1 to Stage 3) on the AIM dataset [52] and then finetune the whole model end-to-end on either VideoMatting108 [56] or DVM [45], depending on the target evaluation benchmark (Stage 4).

#### **Stage 1: Training the alpha matting module and trimap propagation module separately.**

As two modules depend on each other, if we train the alpha matting module and the trimap propagation module simultaneously from scratch, this can lead the model to either poor convergence or simply memorizing training data (*i.e.*, overfitting). It is because both modules cannot learn meaningful features from almost randomly initialized input data which is as the output of other modules. Therefore, we first separately train two modules without the connections between two. Specifically, we train the alpha matting module with GT trimaps and train the trimap propagation module without taking inputs of an alpha matte and hidden features.

**Stage 2: Training the alpha matting and refinement modules with propagated trimaps.** We train the alpha matting model and refinement modules together while the trimap propagation module is frozen. This stage enables the refinement module to take a soft and noisy trimap as input and learn to predict accurate trimap and alpha matte.

**Stage 3: Training the trimap propagation module.** In the trimap propagation module, we activate all input layers for alpha matte and hidden features. Then, we train the trimap propagation module while the parameters for the remainders – alpha prediction and refinement – are frozen. In this stage, we leverage not only the loss from the predicted trimap but also the losses from the alpha prediction. This enables the trimap propagation module to predict a more reliable trimap for estimating the alpha matte. While we are not updating the alpha network and refinement module in this stage, we can leverage the gradients from their losses for updating the trimap propagation module.

**Stage 4: End-to-end training.** Finally, we train the whole network end-to-end using a video matting dataset. With the stage-wise pretraining, we can effectively leverage both image and video data, and achieve stable performance improvement at the end-to-end training.

### 3.5 Training Details

**Data preparation.** During training, we randomly sampled three temporally ordered foreground and background frames from each video sequence. If an image dataset (*e.g.*, AIM [52]) is used, we simulate three video frames from a pair of foreground and background images by applying three different random affine transforms into both foreground and background images. The random affine transforms include horizontal flipping, rotation, shearing, zooming, and translation. For each foreground and background frame, we randomly crop patches into  $320 \times 320$ ,  $480 \times 480$ , or  $640 \times 640$ , centered on pixels in the unknown regions. And then we resize the cropped patches into  $320 \times 320$ . Additionally, we employ several augmentation strategies on both foreground and background frames: histogram matching between foreground and background colors, motion blur, Gaussian noise, and JPEG compression. Then we composite foreground and background on-the-fly to generate an input frame. The GT trimaps are generated by dilating the GT alpha matte with a random kernel size from  $1 \times 1$  to  $26 \times 26$ .

**Loss functions.** We set objective functions for all outputs of the models, except for the hidden features. For both initially predicted and refined trimaps, we use the cross-entropy loss to compare with the GT. For the first frame where the GT trimap is provided as the input, we only apply the loss to the refined trimap. Ideally, there should be no change after refinement. We find penalizing any change after the refinement is helpful to prevent it from corrupting already accurate trimap. For the alpha predictions, we leverage the temporal coherence loss [45] and image matting losses used in FBA [14]. Different from some previous methods that only compute alpha losses on unknown regions (*e.g.* [52]), we compute our losses on every pixel.

In addition to trimap and alpha losses, we also employ losses for the foreground and background color predictions. We estimate foreground and background colors from the alpha decoder and refinement module following [14]. We minimize all foreground and background losses used in [14], and additionally employ temporal coherence loss on both foreground and background. For the foreground color, we compute the losses only where an alpha value is greater than 0 because the exact foreground color is available only in those regions. More detailed explanations of loss functions are given in the supplementary material.

**Other training details.** We opt for RAdam optimizer [32] with a learning rate of  $1e-5$ . We drop the learning rate to  $1e-6$  once at 90% iteration for each training stage. We freeze all the batch normalization layers in the networks. We used a mini-batch size of 4 and trained with four NVIDIA GeForce 1080Ti GPUs. At the first pretraining stage, we trained the alpha matting model about 100,000 iterations and we trained the trimap propagation model about 400,000 iterations. We trained about 50,000 iterations at each of the second and third training stages. Finally, we trained about 80,000 iterations at the last end-to-end training stage.

### 3.6 Inference Details

We used full-resolution inputs to achieve high-quality alpha matte results. For the memory management in the trimap propagation module, we generally follow STM [34] that stores the first and the previous frames to the memory by default, and additionally saves new memory periodically. We add the intermediate frames to the memory for every 10 frames. To avoid GPU memory overflow, we store only the last three intermediate frames and discard old frames.

## 4 Experiments

### 4.1 Evaluation Datasets and Metrics

We present experimental results and analysis on two latest benchmarks, Video-Matting108 [56] and DVM [45].

**VideoMatting108 dataset** [56] includes 28 foreground video sequences paired with background video sequences in the validation set. The evaluation is conducted with three different trimap settings: narrow, medium, and wide. The groundtruth trimaps are generated by discretizing the groundtruth alpha mattes into the trimaps, followed by dilating the unknown regions with different kernel sizes;  $11 \times 11$  for narrow,  $25 \times 25$  for medium, and  $41 \times 41$  for wide. This benchmark contains long video sequences and the average length of the videos is about 850 frames. Therefore, predicting alpha mattes with only a single trimap at the first frame is challenging in this dataset.

**DVM dataset** [45] provides 12 foreground videos and 4 background videos for validation. The validation set is constructed with an additional 50 foreground images from the AIM dataset [52]. In total, 62 foregrounds are composited with every 4 background videos, resulting in 248 test videos. Following the DVM paper [45], we discard 106 foreground images from the AIM dataset [52] and use the remaining 325 foreground images during our training because there are overlaps in the evaluation.

**Evaluation metrics.** For a fair comparison, we follow the evaluation metrics from two large-scale video matting benchmarks [56,45]. To evaluate on VideoMatting108 [56], we compute SSDA (average sum of squared difference), MSE (mean squared error), MAD (mean absolute difference), dtSSD (mean squared difference of direct temporal gradients), and MESSDdt (mean squared difference between the warped temporal gradient) [12]. To evaluate on DVM [45], we compute SAD (sum of absolute difference), MSE, Grad (gradient error), Conn (connectivity error), dtSSD, and MESSDdt. For all computed metrics, lower is better.

In the original works [56,45], the metrics are computed only on the unknown regions of the GT trimaps. We follow this rule for fair comparisons with the existing methods. However, using only the unknown region cannot capture the errors in the foreground and background regions that occurred by inaccurate trimap propagation, which is important for evaluating the performance of end-to-end video matting methods. Therefore, we present *the modified versions of the metrics* that compute the scores on the full-frames, suffixed with “-V”. The modified metrics are used for our analysis and ablation experiments.

## 4.2 Analysis Experiments

To validate our hypotheses experimentally, we conduct a set of analysis experiments. We use the one-trimap setting, where GT trimap is given only at the first frame. All these analysis experiments are conducted on the VideoMatting108 benchmark with the medium trimap setting.

**Effectiveness of the joint modeling.** We first validate the importance of the joint modeling of trimap propagation and alpha prediction in video matting. For this purpose, we design a simple baseline model for video matting, STM+FBA, that cascades STM [34] for trimap propagation and FBA [14] for alpha prediction from the propagated trimap. Note that we do not use OTVM because our proposals (*i.e.*, trimap and alpha refinement and using hidden features) are only applicable for the joint modeling. To show the effect of joint modeling, we train the baseline model with two different training strategies. One is obtained by training trimap propagation (STM) and alpha prediction (FBA) separately (*i.e.*, decoupled), and the other is by training both modules jointly (*i.e.*, joint). As shown in Table 1(a), the joint modeling greatly improves the alpha matte quality in the practical one-trimap scenario.

**Efficacy of the stage-wise training.** We evaluate the importance of the proposed stage-wise training in OTVM and summarize the result in Table 1(b). When we activate all input layers for alpha matte and hidden features and end-to-end train OTVM from the pretraining stage (*i.e.*, joint), the model marginally surpasses the STM+FBA (joint). If our stage-wise training is applied, OTVM significantly outperforms STM+FBA (joint). The results demonstrate the superiority of our stage-wise training.

**Efficacy of each training stage.** In this experiment, we use OTVM and validate our training strategy. The result is summarized in Table 1(c). To learn the trimap propagation model with the hidden features of the refinement module, we applied the last training stage (*i.e.*, Stage 4: end-to-end training) for all cases.

**Table 1.** Analysis experiments on VideoMatting108 validation set. For all experiments, we use 1-trimap setting where GT trimap is given only at the first frame. “-V” denotes the error has been computed in all regions of the frames (see Sec. 4.1).

(a) Joint modeling.						
Model	Training method	SSDA-V	MSE-V	MAD-V	dtSSD-V	MESSDdt-V
STM+FBA	decoupled	83.61	10.62	22.12	36.31	3.45
	joint	<b>75.36</b>	<b>9.40</b>	<b>21.01</b>	<b>29.64</b>	<b>2.74</b>

(b) Stage-wise training.						
Model	Training method	SSDA-V	MSE-V	MAD-V	dtSSD-V	MESSDdt-V
OTVM	joint	72.07	9.29	20.81	30.20	2.79
	joint + stage-wise	<b>54.67</b>	<b>2.61</b>	<b>13.02</b>	<b>29.87</b>	<b>1.78</b>

(c) Ablation on training stages. Each training stage is described in Sec. 3.4.

Train stages				SSDA-V	MSE-V	MAD-V	dtSSD-V	MESSDdt-V
Stage 1	Stage 2	Stage 3	Stage 4					
			✓	87.31	11.16	23.35	33.29	3.15
✓			✓	76.55	9.68	23.10	31.64	3.14
✓	✓		✓	75.33	9.54	22.30	31.44	3.06
✓	✓	✓	✓	<b>54.67</b>	<b>2.61</b>	<b>13.02</b>	<b>29.87</b>	<b>1.78</b>

(d) Ablation on modules. The modules are depicted in Fig. 3.

Refinement module (output)		Trimap module (input)		SSDA-V	MSE-V	MAD-V	dtSSD-V	MESSDdt-V	Time (sec/frame)
Alpha	Trimap	Alpha	Hidden						
				83.56	12.17	23.59	30.11	2.80	0.799
✓				83.51	11.26	22.49	31.78	2.89	0.951
	✓			78.70	9.80	21.51	29.97	2.80	0.946
✓	✓			77.58	10.55	22.42	31.04	2.98	0.952
✓	✓	✓		62.47	4.27	15.06	29.98	1.85	0.955
✓	✓	✓	✓	<b>54.67</b>	<b>2.61</b>	<b>13.02</b>	<b>29.87</b>	<b>1.78</b>	0.964

(e) Different image matting backbones.

Backbone	Model	SSDA-V	MSE-V	MAD-V	dtSSD-V	MESSDdt-V
DIM [52]	STM+DIM	102.77	12.58	28.35	47.13	5.14
	OTVM	<b>92.30</b>	<b>11.04</b>	<b>25.98</b>	<b>39.27</b>	<b>4.21</b>
GCA [28]	STM+GCA	100.22	12.25	27.33	41.72	4.36
	OTVM	<b>84.73</b>	<b>10.04</b>	<b>24.23</b>	<b>36.02</b>	<b>3.46</b>

(f) Trimap propagation. “-T” is presented to estimate trimap quality and denotes that the unknown region in GT trimap has been modified (see Sec. 4.2).

Method	Precision-T	Recall-T	Average
Decoupled STM [45,56]	96.98	93.58	95.28
OTVM	<b>98.17</b>	<b>95.92</b>	<b>97.05</b>

As shown in the table, each pretraining stage consistently improves performance. The results demonstrate that our training strategy, stage-wise pretraining then end-to-end finetuning, is effective.

**Effectiveness of the proposed modules.** We thoroughly evaluate our proposed modules in Table 1(d). As shown in the table, the alpha matte and trimap refinements lead to complementary performance improvement. More importantly, feeding the alpha matte and hidden features of the refinement module to the trimap propagation module significantly improves performance with small extra computation. The result demonstrates the importance of reliable trimap propagation in the practical video matting setting. It also shows that our joint learning framework and the proposed modifications effectively address the problem.

**Applicability to different image matting networks.** Table 1(e) compares the results with other image matting backbones. We tested DIM [52] and GCA [28].

**Table 2.** Comparison with state-of-the-art methods on public benchmarks. The trimap setting indicates how many GT trimaps are given as input, *i.e.*, “full-trimap” for all frames, “20/40-frame” for every 20/40th frames, “1-trimap” for only at the first frame.

(a) **Comparison on VideoMatting108 validation set.** Results for other methods are directly copied from [56].  
<sup>†</sup> denotes our reproduced results using our training setup.

Trimap Setting	Methods	SSDA	MSE	MAD	dtSSD	MESSDdt
full-trimap	DIM [52]	61.85	9.99	44.38	34.55	2.82
	IndexNet [33]	58.53	9.37	43.53	33.03	2.33
	GCA [28]	55.82	8.20	40.85	31.64	2.15
	TCVOM (GCA) [56]	50.41	7.07	37.65	27.28	1.48
	TCVOM (FBA) <sup>†</sup> [56]	<b>39.76</b>	<b>4.01</b>	<b>28.68</b>	<b>22.93</b>	<b>1.06</b>
1-trimap	STM + TCVOM (GCA) [56]	77.23	22.15	57.40	32.18	2.97
	STM + TCVOM (FBA) <sup>†</sup> [56]	69.96	19.80	51.21	29.76	2.72
	STM + FBA <sup>†</sup> [14]	70.63	20.18	51.20	31.00	2.86
	OTVM	<b>50.51</b>	<b>8.58</b>	<b>37.16</b>	<b>28.28</b>	<b>1.63</b>

(b) **Comparison on DVM validation set.** Results for other methods are directly copied from [45].

Trimap Setting	Methods	SAD	MSE	Grad	Conn	dtSSD	MESSDdt
full-trimap	DIM [52]	54.55	0.030	35.38	55.16	23.48	0.53
	IndexNet [33]	53.68	0.028	27.52	54.44	19.50	0.49
	Context-Aware [21]	51.78	0.027	28.57	49.46	19.37	0.50
	GCA [28]	47.49	0.022	26.37	45.23	18.36	0.33
	DVM [45]	<b>40.91</b>	<b>0.014</b>	<b>19.02</b>	<b>40.58</b>	<b>15.11</b>	<b>0.25</b>
20-frame	DVM [45]	43.66	0.016	26.39	42.23	16.34	0.28
	OTVM	<b>37.90</b>	<b>0.013</b>	<b>19.13</b>	<b>36.48</b>	<b>14.76</b>	<b>0.22</b>
40-frame	DVM [45]	52.85	0.026	-	-	19.23	-
	OTVM	<b>38.24</b>	<b>0.014</b>	<b>19.29</b>	<b>36.83</b>	<b>14.82</b>	<b>0.22</b>
1-trimap	DVM [45]	65.33	0.039	-	-	35.46	-
	OTVM	<b>41.02</b>	<b>0.017</b>	<b>20.17</b>	<b>39.79</b>	<b>14.85</b>	<b>0.25</b>

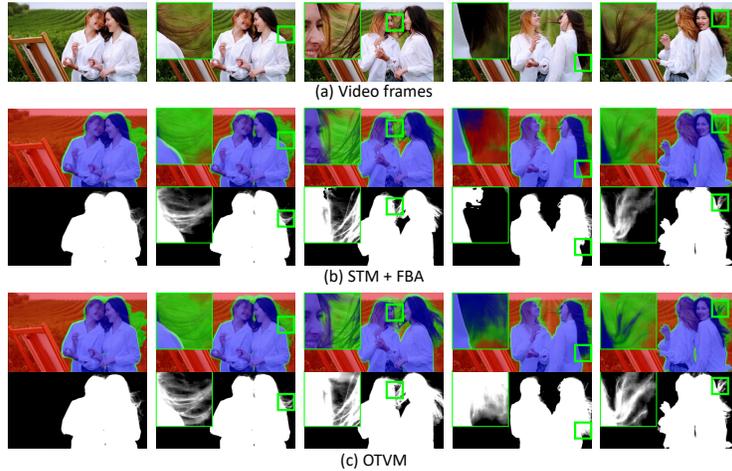
As shown in the table, OTVM consistently leads to performance improvements over the corresponding baselines.

**Accuracy of trimap propagation.** We study the effect of OTVM on trimap propagation. In Table 1(f), we compare OTVM with recent video matting methods [45,56] that naively propagate trimap using a single STM [34]. Ideally, the unknown area in a trimap needs to cover the entire soft matte area in the GT matte while being tight enough not to be trivial. To measure the quality of trimaps, we present two metrics: (1) *Precision-T*, precision of the estimated unknown area compared with the widely dilated GT unknown (dilation kernel size of  $41 \times 41$ ) and (2) *Recall-T*, recall of the estimated unknown area compared with the minimum GT unknown (*i.e.*, no dilation of the GT unknown regions). As shown in Table 1(f), OTVM significantly outperforms the state-of-the-art approach and achieves high precision and high recall.

**More analysis in the supplementary material.** We present additional results on the input of the trimap encoder, a visual analysis of the hidden feature, the effect of the refinement module, an analysis of runtime and GPU memory consumption, an analysis of temporal stability, and quantitative results with image matting metrics (*i.e.* without “-V”) in the supplementary material.

### 4.3 Comparison with State-of-the-Art Methods

**VideoMatting108 [56].** Table 2(a) shows quantitative evaluation results on the VideoMatting108 validation set with medium trimap setting. Results with other trimap settings (*i.e.*, narrow and wide) are provided in the supplementary material. In the experiment, we measured errors only on the unknown regions



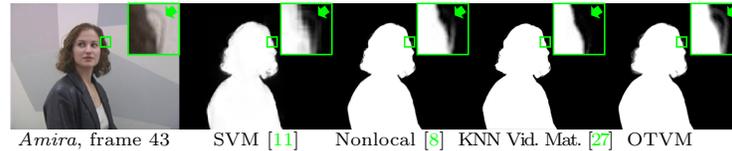
**Fig. 4.** Qualitative comparison on a real-world video sequence. As shown in zoomed trimap, OTVM generates more accurate and fine trimaps, *e.g.*, sharply separates hairs in the foreground region, but STM roughly predicts the regions as unknown. Therefore, we can extract high-quality alpha matte results.

for a fair comparison with [56]. We report two results depending on whether GT trimap is used for all frames or for the first frame. For testing other methods on 1-trimap setting, we used the baseline trimap propagation model we trained to generate trimaps. For a fair comparison with TCVOM [56], which uses GCA [28] for a backbone network, we additionally show FBA backbone results. When a single trimap is used, OTVM significantly outperforms all other methods, demonstrating the superiority of our approach.

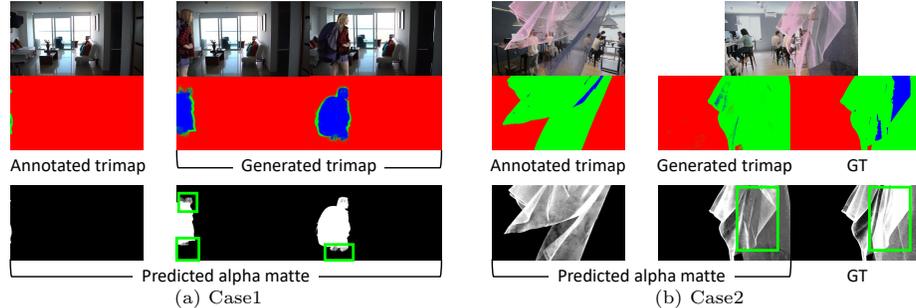
**DVM [45].** We conduct quantitative evaluation on the DVM benchmark and Table 2(b) shows the results. We computed errors only on the unknown regions according to the official metric. Table 2(b) shows that OTVM largely surpasses DVM in 1-trimap setting. Furthermore, OTVM using only 1-trimap (SAD 41.02) achieves comparable performance with DVM using full-trimap (SAD 40.91).

**Qualitative results on real-world videos.** Fig. 4 shows qualitative results on a real-world video. We compare OTVM with the cascaded STM and FBA model, denoted by STM+FBA. As shown in the figure, the STM+FBA model cannot sharply separate foreground and unknown regions on the object boundary. The cascade baseline model fails to predict accurate trimaps in the challenging scenes, resulting in poor alpha mattes. In contrast, OTVM predicts trimaps reliably, resulting in accurate alpha mattes. More qualitative results are provided in the supplementary material. In addition, we provide full-frame results online: <https://youtu.be/qkda4fHSyQE>.

**Comparison with non-deep learning methods.** Fig. 5 shows a comparison with [11,8,27] on the *Amira* benchmark [9]. The benchmark [9] does not provide the GT alpha matte and we took the results of the previous methods from [27,11]. As shown in the figure, SVM [11] result is noisy. Both Nonlocal [8] and KNN



**Fig. 5.** Comparison with non-deep learning methods.



**Fig. 6.** Limitations. We indicate error areas with green boxes. (a) The user-annotated trimap contains almost no object information. (b) There is no strong signal for the foreground in the given trimap.

Video Matting [27] fail to predict hair strand details. In contrast, OTVM predicts the precise alpha matte.

#### 4.4 Limitations

Since our framework takes only a single user-annotated trimap, not only the input trimap quality but also the rich cues of the object in the frame are important. In Fig. 6(a), OTVM struggles to generate accurate trimaps if the user-annotated frame contains almost no object information, resulting in failure to predict the alpha matte. In Fig. 6(b), although the object is presented in the annotated frame, we may struggle to predict precise alpha mattes if there is no strong signal for the foreground object in the given trimap.

## 5 Conclusion

In this paper, we present a new video matting framework that only needs a single user-annotated trimap. In contrast to the recent decoupled methods that focus on alpha prediction given the trimaps, we propose a coupled framework, OTVM, that performs trimap propagation and alpha prediction jointly. OTVM with one user-annotated trimap significantly outperforms the previous works in the same setting and even achieves comparable performance with the previous works using full-trimaps as input. OTVM is simple yet effective and works robustly in the practical one-trimap scenario. We hope that our research motivates follow-up studies and leads to practical video matting solutions.

**Acknowledgements.** This research was supported in part by the Yonsei Signature Research Cluster Program of 2022 (2022-22-0002). This research was also supported in part by the KIST Institutional Program (Project No. 2E31051-21-204).

## References

1. Apostoloff, N., Fitzgibbon, A.: Bayesian video matting using learnt image priors. In: CVPR (2004) [4](#)
2. Bai, X., Wang, J., Simons, D.: Towards temporally-coherent video matting. In: Computer Vision/Computer Graphics Collaboration Techniques (2011) [4](#)
3. Chen, Q., Li, D., Tang, C.K.: Knn matting. IEEE Transactions on Pattern Analysis and Machine Intelligence **35**(9), 2175–2188 (2013) [4](#)
4. Chen, X., Zou, D., Zhiying Zhou, S., Zhao, Q., Tan, P.: Image matting with local and nonlocal smooth priors. In: CVPR. pp. 1902–1907 (2013) [4](#)
5. Cheng, H.K., Tai, Y.W., Tang, C.K.: Rethinking space-time networks with improved memory coverage for efficient video object segmentation. NeurIPS (2021) [2](#)
6. Cheng, M.M., Mitra, N.J., Huang, X., Torr, P.H., Hu, S.M.: Global contrast based salient region detection. IEEE Transactions on Pattern Analysis and Machine Intelligence **37**(3), 569–582 (2014) [7](#)
7. Cho, D., Tai, Y.W., Kweon, I.: Natural image matting using deep convolutional neural networks. In: ECCV. pp. 626–643. Springer (2016) [4](#)
8. Choi, I., Lee, M., Tai, Y.W.: Video matting using multi-frame nonlocal matting laplacian. In: ECCV (2012) [4](#), [13](#), [14](#)
9. Chuang, Y.Y., Agarwala, A., Curless, B., Salesin, D., Szeliski, R.: Video matting of complex scenes. In: SIGGRAPH (2002) [4](#), [13](#)
10. Chuang, Y.Y., Curless, B., Salesin, D.H., Szeliski, R.: A bayesian approach to digital matting. In: CVPR. IEEE (2001) [4](#)
11. Eisemann, M., Wolf, J., Magnor, M.A.: Spectral video matting. In: VMV. pp. 121–126. Citeseer (2009) [13](#), [14](#)
12. Erofeev, M., Gitman, Y., Vatolin, D.S., Fedorov, A., Wang, J.: Perceptually motivated benchmark for video matting. In: BMVC. pp. 99–1 (2015) [10](#)
13. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International Journal of Computer Vision **88**(2), 303–338 (2010) [7](#)
14. Forte, M., Pitié, F.:  $f$ ,  $b$ , alpha matting. arXiv preprint arXiv:2003.07711 (2020) [1](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#), [10](#), [12](#)
15. Gastal, E.S.L., Oliveira, M.M.: Shared sampling for real-time alpha matting. Computer Graphics Forum **29**(2), 575–584 (2010) [4](#)
16. Gong, M., Wang, L., Yang, R., Yang, Y.H.: Real-time video matting using multi-channel poisson equations. In: Graphics Interface (2010) [4](#)
17. Grady, L., Schiwietz, T., Aharon, S., Westermann, R.: Random walks for interactive alpha-matting. In: Proceedings of VIIP. vol. 2005, pp. 423–429 (2005) [4](#)
18. Hariharan, B., Arbeláez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: ICCV. pp. 991–998. IEEE (2011) [7](#)
19. He, K., Rhemann, C., Rother, C., Tang, X., Sun, J.: A global sampling method for alpha matting. In: CVPR. pp. 2049–2056. IEEE (2011) [4](#)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016) [6](#)
21. Hou, Q., Liu, F.: Context-aware image matting for simultaneous foreground and alpha estimation. In: ICCV. pp. 4130–4139 (2019) [4](#), [12](#)
22. Jian, S., Jia, J., Tang, C.K., Shum, H.Y.: Poisson matting. ACM Transactions on Graphics **23**, 315–321 (2004) [4](#)
23. Ke, Z., Sun, J., Li, K., Yan, Q., Lau, R.W.: Modnet: Real-time trimap-free portrait matting via objective decomposition. In: AAAI (2022) [4](#)

24. Lee, P., Wu, Y.: Nonlocal matting. In: CVPR. pp. 2193–2200. IEEE (2011) 4
25. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**(2), 228–242 (2007) 4
26. Levin, A., Rav-Acha, A., Lischinski, D.: Spectral matting. *IEEE transactions on pattern analysis and machine intelligence* **30**(10), 1699–1712 (2008) 4
27. Li, D., Chen, Q., Tang, C.K.: Motion-aware knn laplacian for video matting. In: ICCV (2013) 4, 13, 14
28. Li, Y., Lu, H.: Natural image matting via guided contextual attention. In: AAAI (2020) 1, 4, 11, 12, 13
29. Lin, S., Ryabtsev, A., Sengupta, S., Curless, B.L., Seitz, S.M., Kemelmacher-Shlizerman, I.: Real-time high-resolution background matting. In: CVPR. pp. 8762–8771 (2021) 4
30. Lin, S., Yang, L., Saleemi, I., Sengupta, S.: Robust high-resolution video matting with temporal guidance. In: WACV (2022) 4
31. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755. Springer (2014) 7
32. Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., Han, J.: On the variance of the adaptive learning rate and beyond. In: ICLR (April 2020) 9
33. Lu, H., Dai, Y., Shen, C., Xu, S.: Indices matter: Learning to index for deep image matting. In: ICCV. pp. 3266–3275 (2019) 4, 12
34. Oh, S.W., Lee, J.Y., Xu, N., Kim, S.J.: Video object segmentation using space-time memory networks. In: ICCV (October 2019) 1, 2, 3, 5, 6, 7, 9, 10, 12
35. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675 (2017) 7
36. Porter, T., Duff, T.: Compositing digital images. In: SIGGRAPH. pp. 253–259 (1984) 3
37. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015) 7
38. Sengupta, S., Jayaram, V., Curless, B., Seitz, S.M., Kemelmacher-Shlizerman, I.: Background matting: The world is your green screen. In: CVPR. pp. 2291–2300 (2020) 4
39. Seong, H., Oh, S.W., Lee, J.Y., Lee, S., Lee, S., Kim, E.: Hierarchical memory matching network for video object segmentation. In: ICCV. pp. 12889–12898 (2021) 1
40. Shahrian, E., Price, B., Cohen, S., Rajan, D.: Temporally consistent and spatially accurate video matting. In: Eurographics (2014) 4
41. Shahrian, E., Rajan, D., Price, B., Cohen, S.: Improving image matting using comprehensive sampling sets. In: CVPR. pp. 636–643 (2013) 4
42. Shen, X., Tao, X., Gao, H., Zhou, C., Jia, J.: Deep automatic portrait matting. In: ECCV. pp. 92–107. Springer (2016) 4
43. Shi, J., Yan, Q., Xu, L., Jia, J.: Hierarchical image saliency detection on extended cssd. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(4), 717–729 (2015) 7
44. Sun, J., Ke, Z., Zhang, L., Lu, H., Lau, R.W.: Modnet-v: Improving portrait video matting via background restoration. arXiv preprint arXiv:2109.11818 (2021) 4

45. Sun, Y., Wang, G., Gu, Q., Tang, C.K., Tai, Y.W.: Deep video matting via spatio-temporal alignment and aggregation. In: CVPR. pp. 6975–6984 (2021) [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#)
46. Sun-Young Lee, Jong-Chul Yoon, I.K.L.: Temporally coherent video matting. In: Graphic Models (2010) [4](#)
47. Tang, Z., Miao, Z., Wan, Y.: Temporally consistent video matting based on bilayer segmentation. In: ICME (2010) [4](#)
48. Tang, Z., Miao, Z., Wan, Y., Zhang, D.: Video matting via opacity propagation. In: Visual Computer (2012) [4](#)
49. Wang, J., Cohen, M.F.: Optimized color sampling for robust matting. In: CVPR. IEEE (2007) [4](#)
50. Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: Edvr: Video restoration with enhanced deformable convolutional networks. In: CVPRW (2019) [4](#)
51. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: ECCV. pp. 3–19 (2018) [4](#)
52. Xu, N., Price, B., Cohen, S., Huang, T.: Deep image matting. In: CVPR. pp. 2970–2979 (2017) [1](#), [4](#), [7](#), [8](#), [9](#), [11](#), [12](#)
53. Xu, N., Yang, L., Fan, Y., Yang, J., Yue, D., Liang, Y., Price, B., Cohen, S., Huang, T.: Youtube-vos: Sequence-to-sequence video object segmentation. In: ECCV. pp. 585–601 (2018) [7](#)
54. Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: Youtube-vos: A large-scale video object segmentation benchmark. arXiv preprint arXiv:1809.03327 (2018) [1](#)
55. Yu, Q., Zhang, J., Zhang, H., Wang, Y., Lin, Z., Xu, N., Bai, Y., Yuille, A.: Mask guided matting via progressive refinement network. In: CVPR. pp. 1154–1163 (2021) [4](#)
56. Zhang, Y., Wang, C., Cui, M., Ren, P., Xie, X., Hua, X.s., Bao, H., Huang, Q., Xu, W.: Attention-guided temporal coherent video object matting. In: ACM MM (2021) [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [9](#), [10](#), [11](#), [12](#), [13](#)
57. Zhu, B., Chen, Y., Wang, J., Liu, S., Zhang, B., Tang, M.: Fast deep matting for portrait animation on mobile phone. In: ACM MM. pp. 297–305 (2017) [4](#)