

# Virtual Multi-Modality Self-Supervised Foreground Matting for Human-Object Interaction

Bo Xu<sup>1</sup>, Han Huang<sup>1</sup>, Cheng Lu<sup>2</sup>, Ziwen Li<sup>1,3</sup> and Yandong Guo<sup>1,\*</sup>

<sup>1</sup>OPPO Research Institute, <sup>2</sup>Xmotors, <sup>3</sup>University of California, San Diego

yandong.guo@live.com

## Abstract

Most existing human matting algorithms tried to separate pure human-only foreground from the background. In this paper, we propose a Virtual Multi-modality Foreground Matting (VMFM) method to learn human-object interactive foreground (human and objects interacted with him or her) from a raw RGB image. The VMFM method requires no additional inputs, e.g. trimap or known background. We reformulate foreground matting as a self-supervised multi-modality problem: factor each input image into estimated depth map, segmentation mask, and interaction heatmap using three auto-encoders. In order to fully utilize the characteristics of each modality, we first train a dual encoder-to-decoder network to estimate the same alpha matte. Then we introduce a self-supervised method: Complementary Learning(CL) to predict deviation probability map and exchange reliable gradients across modalities without label. We conducted extensive experiments to analyze the effectiveness of each modality and the significance of different components in complementary learning. We demonstrate that our model outperforms the state-of-the-art methods.

## 1. Introduction

Matting is known as a chroma keying process to separate the foreground from a single image or video stream and then composite it with a new background. It has long been applied in photography and special-effect film-making [33]. Due to the rapid development of deep neural network in computer vision, automatic matting becomes increasingly matured [7, 20, 23, 29, 33, 35]. However, most existing deep learning matting methods extract pure human foreground under controlled settings.

We revisit the matting problem with the following three beliefs. First, we argue that foreground matting algorithm should be able to handle situations when humans are interacting with objects. Currently, most of human matting algorithms [6, 7, 20, 22, 23, 29, 35] only focus on the body region while ignoring the object interacting with the per-




Figure 1: Given challenging human-interactive images, one of the recent state-of-the-art foreground matting approach Background Matting (BGM) [29], fails to produce accurate alpha mattes. However our proposed multi-modality model VMFM is able to separate accurate human-object interactive foreground mattes, outperforming SOTA works.

son of interest, as illustrated in Figure 1. Second, the ideal matting model should be trained on a reliable and representative dataset, which requires the minimum amount of supervision and labeling effort. That is because the definition of good human-object interactive matting is very subjective and such labeling is typically expensive and difficult, if not impossible. Third, we believe that this human-object interactive foreground matting should be done under unconstrained conditions. Until now, most accurate matting techniques still rely on “blue-screen” or “green-screen” to remove background in a recording for traveling mattes. Some other methods require trimap (trimap is a draft marking foreground, background, and unknown areas) as an even strong prior [15, 23, 33, 35]. Learning without predefined background [29] or prior knowledge is critical.

To address the above issues, we propose the solution with following two components accordingly. First, we reformulate this problem as a multi-modality task to better capture both human of interest and the interacted objects. We believe feature extracted in the RGB space is not enough to produce high quality mattes. In our case, we introduce depth estimation as the extra virtual modality on top of RGB image and estimated segmentation mask. We argue

that while RGB and segmentation map is good at predicting human contours and regions with similar color (skin *et al.*) [29], however it struggles to recognize the association between human and the interacted object. Compared with a human body, the interacted objects are typically more diverse in terms of appearance *e.g.* on shape, color, and type. That makes detection of the human-object association based on similar appearance in RGB space almost impossible. However, the depth modality can easily group human and object into one entity because of the high degree of continuity in its feature space. However, most current segmentation networks are not good at finding depth-related feature since they are trained intentionally to separate, rather than to group, different objects. To avoid ambiguity, we follow the definition of interaction as in [12], which specifies human-object interaction into 26 classes. To make our model focus on the features of humans and interactive objects, we introduce interaction heatmaps which are also generated from RGB image as the semantic prior.

Secondly, to minimize the labeling cost we propose the Complementary Learning (CL) self-supervised module in our VMFM network. The output of each modality can supervise or be supervised by its counterpart. Collectively, they manage to find the best combination. That is inspired by the fact each modality has its unique weakness and strength. For example in Figure 2, depth modality can help us predict the association between a man and the book in his hand but is weak on distinguishing the outline between ground and feet. On the contrary, color feature can easily tell the feet from the ground but struggle to know that the book is part of the foreground. However, we can learn a collection of more reliable elements from both modalities through pixel-wise complementary learning; see Figure 6 for examples of VMFM’s outputs.

The VMFM network consists of two stages: Foreground Prediction (FP) and Complementary Learning (CL). The FP stage includes two modules: segmentation-based foreground prediction network (SFPNet) and depth-based foreground prediction network (DFPNet), both in encoder-to-decoder fashion. The depth map is estimated by a depth estimation network (D-Net) from a single RGB image and the network is based on the backbone of [2]. The I-Net, in which Hourglass [26] backbone embeds the human-object interactive feature, is used to guide the matting by generating attentional heatmaps. In the CL stage, each output pixel of one FP network can supervise, or be supervised by its counterpart (output pixel at the same location in the other network), depending on which one has a lower deviation probability. Jointly, we minimize the unreliable pixels (probability higher than the threshold we set) to better highlight the foreground pixels that belong to either human or interacted objects. To justify our solutions, we compare our algorithm, objectively and subjectively with other meth-

ods. Also, we demonstrate by ablation study that different modalities focus on different, if not unique, details. Complementary learning can combine their reliable pixels while suppressing unreliable ones. Overall, the contributions of this paper are as follows:

- This is the first end-to-end foreground/alpha matting algorithm that focuses on human-object interactive scenarios under unconstrained conditions.
- We reformulate foreground alpha matting as a cross-modal self-supervised task, which minimizes the labeling cost and produces fine matting by leveraging both modalities.
- Extensive experiments demonstrate the effectiveness of VMFM, outperforming the state-of-the-art (SOTA) methods in real-world scenarios.

## 2. Related works

Currently, matting is generally formulated as an image composite problem, which solves the 7 unknown variables per pixel from only 3 known values:

$$I_i = \alpha_i F_i + (1 - \alpha_i) B_i \quad (1)$$

where 3 dimensional RGB color  $I_i$  of pixel  $i$ , while foreground RGB color  $F_i$ , background RGB color  $B_i$  and matte estimation  $\alpha_i$  are unknown. In this section, we discuss the SOTA works dealing with this under-determined equation.

### 2.1. Classic methods

Classic foreground matting methods can be generally categorized into two approaches: sampling-based and propagation-based. Sampling-based methods [1, 8, 10, 16, 18, 34] sample the known foreground and background color information, and then leverage them to achieve matting in the unknown region. Various sampling-based algorithms are proposed, *e.g.* Bayesian matting [34], optimized color sampling [16], global sampling method [18] and comprehensive sampling [10]. Propagation-based methods [3, 7, 14, 19, 20, 31] reformulate the composite Eq. 1 to propagate the alpha values from the known foreground and background into the unknown region, achieving more reliable matting results. [17] provides a very comprehensive review of different matting algorithms.

### 2.2. Deep learning based methods

Classic matting methods are carefully designed to solve the composite equation and its variant versions. However, these methods heavily rely on chromatic information, which leads to bad quality when the color of foreground and background shows small or no noticeable difference.

Automatic and intelligent matting algorithms are emerging, due to the rapid development of deep neural network

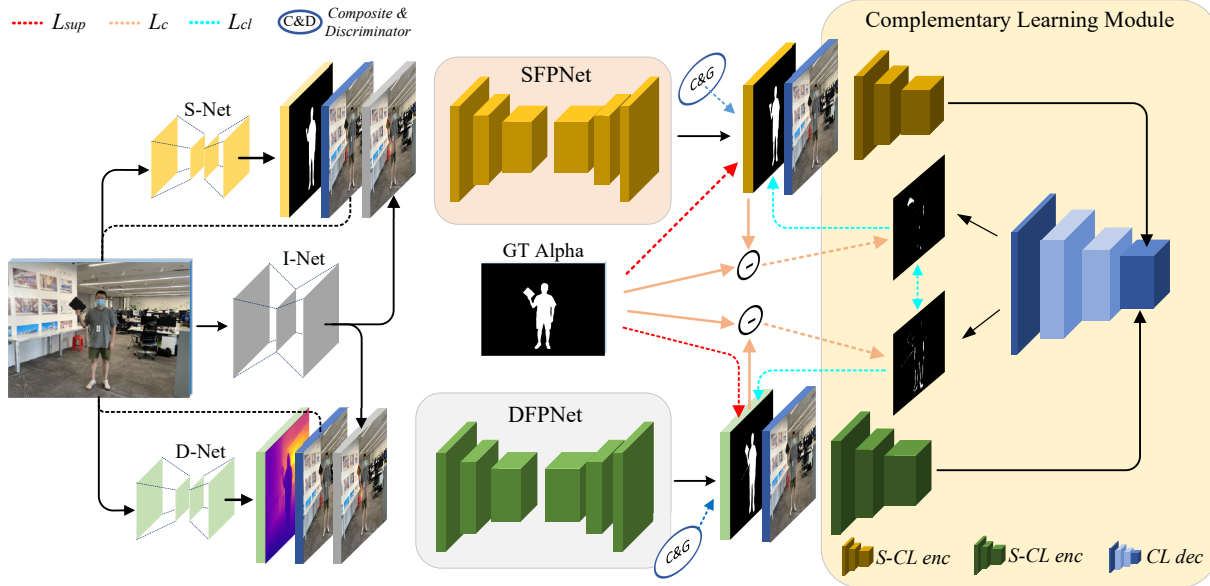


Figure 2: Architecture of the Virtual Multi-modality Foreground Matting (VMFM) network in the training stage. The Foreground Prediction (FP) module consists of dual networks **SFPnet** and **DFPnet**, which receives different modalities as inputs and estimate the same alpha matte. The Complementary Learning (CL) module consists of two encoders **S-CL enc**, **D-CL enc** and one decoder **CL dec**, outputting deviation probability maps of the predicted alpha mattes.

in computer vision. Initially, some attempts were made to combine deep learning networks with classic matting techniques, *e.g.* closed-form matting [20] and KNN matting [7]. Cho *et al.* [9] employ a deep neural network to improve results of the closed-form matting and KNN matting. These attempts are not end-to-end, so not surprisingly the matting performance is limited by the convolution back-ends. Subsequently, full DL image matting algorithms appear [6, 11, 33]. Xu *et al.* [33] propose a two-stage deep neural network (Deep Image Matting) based on SegNet [4] for alpha matte estimation and create a large-scale image matting dataset (Adobe dataset) with ground truth foreground (alpha) matte, which can be composited over a variety of backgrounds to produce training data. And we also use this data for the first-step pre-training of our network. Lutz *et al.* [24] introduce a generative adversarial network (GAN) for natural image matting and improve the results of Deep Image Matting [33]. Cai *et al.* [5] investigated the bottleneck of the previous methods that directly estimate the alpha matte from a coarse trimap, and propose to divide matting problem into trimap adaptation and alpha estimation tasks. Hou *et al.* [15] employs two encoder networks to extract essential information for matting, however it is not robust to faulty trimaps.

**Trimap-free methods.** Currently, a majority of deep image matting algorithms [5, 15, 24, 33] try to estimate a boundary that divides the foreground and background, with the aid of a user-generated trimap. Several trimap-free mat-

ting methods [6, 35] predict the trimap first, followed by alpha matting.

**Additional natural background.** Qian *et al.* [27] compute a probability map to classify each pixel into the foreground or background by simple background subtraction. This algorithm is sensitive to the threshold and fails when the colors of foreground and background are close. Sen Gupta *et al.* [29] introduce a self-supervised adversarial approach - Background Matting (BGM), achieving state-of-the-art results. However, a photographer needs to take a shoot of natural background first, which is not friendly to the intensive multi-scene shooting application.

Besides, current DL foreground matting methods focus only on human matting and ignore the fact that objects interacting or attached with people can always, if not every time, be part of the foreground. This is the main reason why we propose the method of VMFM. In this paper, we quantitatively evaluate the performance of our model for alpha matting in human-object interactive scenes.

### 3. Architectures

The network architecture of the virtual multi-modality foreground matting (VMFM) is designed to fully automatically extract accurate human-object interaction foreground instead of target selection through human-computer interaction. It first learns to estimate the same foreground alpha matte through the dual network (SFPnet and DFPnet) of the Foreground Prediction (FP) module. Since predictions of




Figure 3: Illusion of interaction pair heatmaps on an example image. Interaction pair means human and one of the objects interacted with him or her. Classes of such interactions are predefined and cover most scenarios for matting purpose.

SFPnet and DFPnet are naturally biased to different regions, *e.g.* segmentation map is sensitive to the human body, while depth map focuses more on the human-object interactive area. we introduce Complementary Learning (CL) to learn more dependable elements between dual FP networks. The overall architecture of VMFM training network is shown in detail in Figure 2.

### 3.1. Foreground Prediction

**Modality preprocessing.** Given an input RGB image, we use a depth estimation backbone network (D-Net [2]) to estimate its depth map and S-Net (Mask-RCNN [13]) to automatically segment the mask of human. Moreover, we follow [32] to employ Hourglass [26] as the backbone (I-Net) for predicting human-object paired heatmaps. As shown in Figure 3, interaction heatmaps can provide semantic prior of human and his or her interactive objects. For human without any associated object, the peak of interaction heatmap generalizes to the center point of him or her.

Due to the difference between segmentation and depth features, we build a dual network for foreground prediction (FP) - segmentation-based foreground prediction network (SFPnet) and depth-based foreground prediction network (DFPnet). SFPnet receives the raw RGB image  $I$ , segmentation mask  $S$ , and interaction heatmap  $H$  as inputs, while DFPnet replaces  $S$  with the depth map  $D$ . In the following,  $SFPnet(I, S, H)$  and  $DFPnet(I, D, H)$  are referred to as  $FP^m$ ,  $m \in \{1, 2\}$ . Both of the dual networks are built in encoder-to-decoder fashion, and the architecture of each is well-modularized to make sure easy substitution to other networks that can serve the same purpose. We encourage other fellow researchers in the community to replace modules in VMFM to achieve better performance. The dual network of FP module is pre-trained on the subset of human subjects in the Adobe Matting Dataset [33]. The pre-training is supervised by minimizing the  $L_a^m$  loss,

$$L_a^m = \|a^m - a^*\|_1 + \|\nabla(a^m) - \nabla(a^*)\|_1 \quad (2)$$

where  $a^m$  is the alpha matte output of  $FP^m$ ,  $m \in \{1, 2\}$ ,  $a^*$  is the ground truth alpha matte, and the gradient term is

beneficial to remove the over-blurred alpha matte [35].

We adopt a two-stage training process for foreground prediction in our method which are regular supervision and self-supervision respectively. In the first stage, we train the FP network under supervision (with labeled data). The proposed self-supervision is achieved through complementary learning on unlabeled data in the second stage. In addition, we introduce a discriminator based on LS-GAN [25] to distinguish between fake composites and real images to improve the foreground matting network. For the generator  $\mathcal{G}$  update in the first stage, the adversarial loss term is:

$$L_{\mathcal{G}1}^m = \mathbb{E}_{X, \bar{B} \sim p_{X, \bar{B}}} [(\mathcal{D}(a^m F^* + (1 - a^m) \bar{B}) - 1)^2 + \lambda_a L_a^m + \lambda_{com} L_{com}^m] \quad (3)$$

$$L_{com}^m = \|a^m F^* + (1 - a^m) B^* - I\|_1 \quad (4)$$

where  $a^m = \mathcal{G}(X)$ ,  $X$  comprises RGB image  $I$  and its corresponding virtual modalities,  $L_{com}^m$  denotes the composite loss, which is the absolute difference between the input RGB image  $I$  and the predicted RGB image generated from the ground truth foreground  $F^*$ , the true background  $B^*$  and the predicted foreground alpha matte  $a^m$ .  $L_{com}^m$  can regularize the network to follow the compositional operation, which further reduces the error of foreground alpha prediction.  $\bar{B}$  is a given background for generating a composite image seen by the discriminator  $\mathcal{D}$ .

The objective for  $\mathcal{D}$  is:

$$L_{\mathcal{D}}^m = \mathbb{E}_{X, \bar{B} \sim p_{X, \bar{B}}} [(\mathcal{D}(a^m F^* + (1 - a^m) \bar{B}))^2] + \mathbb{E}_{I \in p_{data}} [(\mathcal{D}(I) - 1)^2] \quad (5)$$

For the generator's update in the second stage, we minimize:

$$L_{\mathcal{G}2}^m = \mathbb{E}_{X, \bar{B} \sim p_{X, \bar{B}}} [(D(a^m F^* + (1 - a^m) \bar{B}) - 1)^2 + \lambda_{cl} L_{cl}^m] \quad (6)$$

where  $L_{cl}^m$  denotes a complementary learning constraint for self-supervision between the dual network of FP module. For the discriminator in Stage 2, we also minimize Eq. 5. The details of the Complementary Learning (CL) module are described in Section 3.2.

### 3.2. Complementary Learning

We propose a dual network with different modalities to predict the same alpha matte. The underlying assumption is that each modality is good at predicting certain regions while weak in others. Therefore, we propose the Complementary Learning (CL) module to estimate the deviation probability map for predicted alpha mattes. The CL forces each of the dual network (SFPNet and DFPNet) to learn more reliable pixels between estimated matte  $a^1$  and  $a^2$  while rejecting the relatively unlikely one at the same

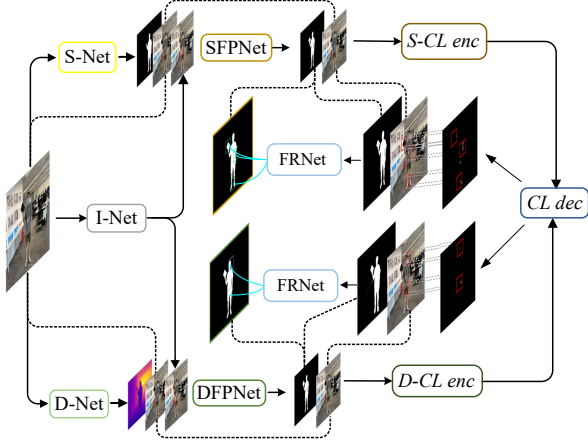


Figure 4: Architecture of VMFM in the inference stage. **FRNet**: foreground refinement network.

location. We use encoder-to-decoder (enc2dec) mechanism which includes dual encoder  $CL_{enc}^m(I, a^m)$ ,  $m \in \{1, 2\}$  and one decoder  $CL_{dec}(CL_{enc}^m)$  in the CL module.

**CL module training.** We train the CL module simultaneously with the FP network training at the first stage, so that it is directly supervised by the true deviation probability map between each FP network result  $a^m$  and its ground truth  $a^*$ . The loss is defined as

$$L_c^m = \|CL(I^m, a^m) - Q(|a^m - a^*|)\|_1 \quad (7)$$

where  $CL(I^m, P^m)$  denotes the CL module which is noted as  $CL^m$  in the following,  $Q$  denotes a dilation-normalization operation to broaden the area of complementary learning and produce a deviation probability map. We normalize the values in  $CL^m$  to  $[0, 1]$ . Interchangeably, we can also regard deviation probability as an expression of confidence. Intuitively one pixel with a higher probability value in the deviation probability map indicates lower confidence. We set a probability threshold  $\tau \in (0, 1)$  that only pixel with higher deviation probability than  $\tau$  will be supervised by its counterpart (the pixel estimated by the other FP network at the same position) whose probability is lower than  $\tau$ . We update the higher value  $CL_{i,j}^m$  (with lower confidence) at pixel  $(i, j)$  to 1:

$$CL_{i,j}^m = \begin{cases} 1, & \text{if } CL_{i,j}^m > \tau \\ CL_{i,j}^m, & \text{else} \end{cases} \quad (8)$$

**Operation of complementary learning.** In the second stage (self-supervision) of foreground prediction (FP) training, we freeze the CL module and conduct complementary learning bounded by the constraint  $L_{cl}^m$  in Eq. 9 to update the dual FP network.

$$L_{cl}^m = \lambda_{cs} L_{cs}^m + \lambda_{dc} L_{dc}^m \quad (9)$$

where we introduce a complementary supervision constraint  $L_{cs}^m$  to switch between SFPnet and DFPnet training:

$$L_{cs}^m = \beta_m \|a^m - a^{3-m}\|_1$$

$$\beta_m^{(i,j)} = \begin{cases} 1, & \text{if } CL_{i,j}^m > CL_{i,j}^{3-m} \\ 0, & \text{else} \end{cases} \quad (10)$$

$\beta_m$  is a complementary learning area map, where  $\beta_m^{(i,j)} = 1$  means that the current network  $FP^m$  needs to learn from the result  $(a_{i,j}^{3-m})$  of the other FP network ( $FP^{3-m}$ ), at the pixel  $(i, j)$  and vice versa,  $m \in \{1, 2\}$ .

$L_{dc}^m$  is a deviation correction constraint to eliminate unreliably predicted pixels of  $FP^m$ :

$$L_{dc}^m = \sigma_m \|CL^m - 0\|_2$$

$$\sigma_m^{(i,j)} = (CL_{i,j}^1 \text{ is } 1 \text{ and } CL_{i,j}^2 \text{ is } 1) \quad (11)$$

$\sigma_m$  is a deviation correction area map, where  $\sigma_m^{(i,j)} = 1$  means both pixels  $(CL_{i,j}^1 \text{ and } CL_{i,j}^2)$  show low confidence. We set constraint  $L_{dc}^m$  for updating dual FP network to limit the number of such unreliable pixel pairs. Obviously,  $L_{cs}^m$  and  $L_{dc}^m$  are non-interfering.

### 3.3. Foreground Refinement

Although complementary learning improves the performance of foreground matting in the unlabeled condition, the encoder-to-decoder structure in FP module may over-regularize the results. Therefore, during the inference stage, we extend the above proposed pipeline by adding the matting refinement network (**RN**), as illustrated in Figure 4. To reduce computational complexity, we selectively extract top- $K$  pixels with the highest estimated errors in the deviation probability map as centers to define  $16 \times 16$  patches in the predicted alpha matte for further refinement. Each selected patch is concatenated with its RGB region before being fed into the extended network. We apply loss  $L_a^m$  (Eq. 2) to supervise the matting refinement network with the parameters of FP module fixed. We illustrate the effect of our refinement network in Section 4. More network architecture details are given in the supplementary material.

## 4. Experiments

We first describe the datasets used for training and testing. Subsequently, we compare our results with existing state-of-the-art (SOTA) foreground matting algorithms. Finally, We conduct ablation experiments to show the effectiveness of each module. More implementation details are provided in the supplementary material.

### 4.1. Datasets

**Pre-train.** We follow the method of [29] to produce 26.9k composited images for pre-training our dual Foreground Prediction (FP) network: 269 Adobe [33] samples




Figure 5: Comparison of benchmarks. (a) and (b) are examples of the Adobe [33] benchmark. (c) and (d) are examples of the Distinctions-646 [28] benchmark. Prior benchmarks mainly contain scenes of portrait and human-only, extremely rare human-object interaction. However, our benchmarks provide diverse classes of interactive objects, interactive modes and scenarios to cover different matting application needs.

Dataset	Human-only	Interaction	$B^*$	Annotation
DAPM [30]	✓	-	-	✓
Adobe [33]	✓	25,3	-	✓
SHM [6]	✓	-	-	✓
BSTMCA [22]	✓	-	-	✓
Dist-646 [22]	✓	-	-	✓
BGM <sub>real</sub> [29]	✓	-	✓	-
Ours(LFM40K)	✓	✓	✓	✓
Ours(UFM75K)	✓	✓	✓	-

Table 1: The configurations of foreground matting datasets.  $B^*$ : true background. (25,3) means that only about 25 foregrounds from the training set and 3 from the testing set in Adobe benchmark are interactive scenes by definition.

of human subjects composited over 100 random COCO [21] images as backgrounds.

Most of the existing matting datasets *e.g.* [6, 22, 28, 33] consist of HD SLR images, which focus narrowly on portrait or human body alone. However, these samples are still far different from mobile captures. More importantly, the diversity of human-object interaction is severely limited in these images.

To facilitate the research in this area (especially the foreground matting in the human-object interactive settings), we construct two human-object interactive matting datasets which cover rich diversity of backgrounds. Our proposed model is trained accordingly to make sure good performance in real-world scenarios. The datasets and training strategy together significantly extend the use cases of the existing algorithms. We claim that our proposed model is bounded by none of the known constraints in the prior arts, such as trimap or predefined background.

**LFM40K.** We propose our first dataset Labeled Foreground Matting 40K (**LFM40K**), which contains more than 40000 (31950 in the training set and 8050 in the test set) labeled frames (images) from 85 videos. LFM40K has 20

interaction classes, *e.g.* carry, hold, kick, ski, work on computer and so on. We set 25% of the dataset with human-only foreground scenes to better cover the rich diversity of different application scenarios.

**UFM75K.** Our unlabeled dataset - Unlabeled Foreground Matting 75K (**UFM75K**), consists of more than 75000 unlabeled images (62750 for self-supervised training and 12250 for testing) from 172 videos indoors and outdoors. Similarly, 25% of the dataset is human-only foreground scene. UFM75K shares the same interactive classes with LFM40K but is more diverse in terms of interactive objects. We take full advantage of the inexpensive but large quantity of UFM75K under self-supervision to ensure the robustness of our network on different application scenarios. Interaction category and dataset details are introduced in supplementary material.

Table 1 and Figure 5 show the comparison between some existing foreground matting datasets with ours. Our datasets consist of both human-only and human-object interacted data. Our datasets could serve as a new challenging benchmark in the foreground matting area. LFM40K dataset is the first high-quality annotated human-object interactive dataset under unconstrained conditions while UFM75K covers a great diversity of real-world scenarios.

## 4.2. Comparative study on composition datasets

We conduct comparative study on two composition benchmarks: Adobe [33] and Dist-646 [28] datasets. To construct test benchmarks, we separately composite 11 held-out samples of human subjects from Adobe test set and 10 held-out ones of human subjects from Dist-646 with 20 random backgrounds per sample.

We report mean square error (MSE), sum of the absolute difference (SAD), spatial-gradient (Grad) and connectivity (Conn) between predicted and ground truth alpha mattes. Lower values of these metrics indicate better estimated alpha matte. The quantitative results are shown in Table 2.

Datasets	Methods	SAD	MSE	Grad	Conn
Adobe	CAM	13.92	4.67	70.32	10.02
	IM	14.61	4.92	75.68	10.61
	LFM	40.35	16.80	125.77	33.65
	BGM	14.12	4.79	72.45	10.18
	S-CL-RN*	<b>12.68</b>	<b>4.39</b>	<b>62.91</b>	<b>8.64</b>
	D-CL-RN*	13.01	4.68	71.23	9.02
Dist-646	CAM	17.40	5.11	82.29	15.33
	IM	18.69	5.82	88.61	17.26
	LFM	38.31	15.63	112.75	36.12
	BGM	12.82	4.95	70.81	11.97
	S-CL-RN*	<b>10.62</b>	<b>3.97</b>	<b>65.39</b>	<b>10.65</b>
	D-CL-RN*	12.35	4.24	68.23	11.04

Table 2: The quantitative results on composition benchmarks. We scale MSE, Grad and Conn by  $10^{-2}$ ,  $10^2$  and  $10^3$ . **CAM**: Context-Aware Matting [15], **IM**: Index Matting [23], **BGM**: Background Matting [29], **LFM**: a Late Fusion Matting [35]. \* means our proposed models. All competitive methods need additional inputs of trimaps (CAM and IM) or the true backgrounds (BGM), except LFM and ours.

Although almost all test samples in both two test composition benchmarks are human-only scenes, our VMFM method (S-CL-RN) shows significant superiority over competing trimap-free method (LFM [35]), and also outperforms trimap-based (CAM [15], IM [23]) and background-based (BGM [29]) ones. Moreover, VMFM can leverage more semantic prior from virtual modalities instead of additional inputs. Row 1 and 2 of Figure 6 visualize comparisons on Adobe testing sample and more representative visualizations are provided in supplementary material. Effectiveness of each component in VMFM is detailed in Section 4.3 and 4.5.

### 4.3. Comparative study on LFM40K

We re-scale all images to  $512 \times 512$  pixel patches in training and testing of LFM40K dataset. Table 3 compares our dual foreground matting (FP) network (SFPnet and DFPnet) with previous state-of-the-art (SOTA) foreground matting algorithms on diverse human-object interactive scenes. Because both **IM** and **CAM** require trimaps as additional inputs, we dilate-erode the ground truth alpha mattes for each image in LFM40K to form such trimaps.

The quantitative results on LFM40K are shown in Table 3. We observe that before applying CL, our SFPNet not surprisingly underperforms by small margin compared to CAM, IM, and BGM due to removal of trimap and background image, both of which require expensive manual creation and captures. Compared with LFM, both SFPNet and DFPNet exhibit significant superiority. That is because we

Methods	SAD	MSE( $10^{-2}$ )	Grad( $10^2$ )	Conn( $10^3$ )
CAM	14.48	5.01	81.62	18.34
IM	16.36	6.73	96.49	20.96
LFM	43.29	15.83	126.31	40.37
BGM	10.06	4.28	70.26	14.49
SFPNet	18.28	8.05	105.91	21.23
S-CL <sub>cs</sub>	11.85	4.29	78.63	15.70
S-CL <sub>dc</sub>	16.43	6.81	102.35	21.65
S-CL	7.02	3.65	62.72	10.91
S-CL-RN	5.23	3.04	56.42	6.35
DFPNet	16.05	6.36	90.12	20.51
D-CL <sub>cs</sub>	9.39	4.13	69.81	12.76
D-CL <sub>dc</sub>	15.24	5.07	87.09	18.22
D-CL	5.20	2.94	56.07	6.29
D-CL-RN	<b>4.39</b>	<b>2.08</b>	<b>51.62</b>	<b>5.26</b>

Table 3: The quantitative results on the test set of LFM40K. **S-CL<sub>cs</sub>** (**D-CL<sub>cs</sub>**): SFPNet (DFPNet) adds CL with only  $L_{cs}$  constraint, **S-CL<sub>dc</sub>** (**D-CL<sub>dc</sub>**): SFPNet (DFPNet) adds CL with only  $L_{dc}$  constraint.

introduce a virtual multi-modality mechanism to extend semantic space of the RGB image so that the proposed models can better capture the interactive information between human and objects. Our DFPNet is slightly inferior to CAM but superior to IM, which demonstrates that depth modality can better associate human body and the interacted object. After applying complementary learning, our methods (D-CL and S-CL) significantly outperform all previous SOTA foreground matting algorithms.

### 4.4. Comparative study on UFM75K

To evaluate and benchmark our algorithm on unlabeled data (UFM75K), we follow [29] to create the pseudo trimaps for trimap-based algorithms. The trimap creation implement details are described in the supplementary material. We compare our algorithm with other competing algorithms on UFM75K dataset for evaluation. Due to lack of ground truth alpha mattes, we composite the estimated mattes over a green or blue background and perform a psychophysical user study that aligns with [29]. Since IM [23] and our VMFM only estimate alpha matte (without generating foreground image  $F$ ), we set  $F = I$  for these algorithms.

**User study.** We compare our composited videos head-to-head with each of the competing methods on 30 test videos from UFM75K. Similar to [29], we present each user with a web page showing the raw video, our composite, and a competing composite; we randomly shuffle the order of the last two. Users were then asked to rate composite A relative to B on a scale of 1-5 (1 being ‘much worse’, 5 ‘much better’). Each video pair was rated by 20 users or more.

The subjective metrics of user study shown in Table 4

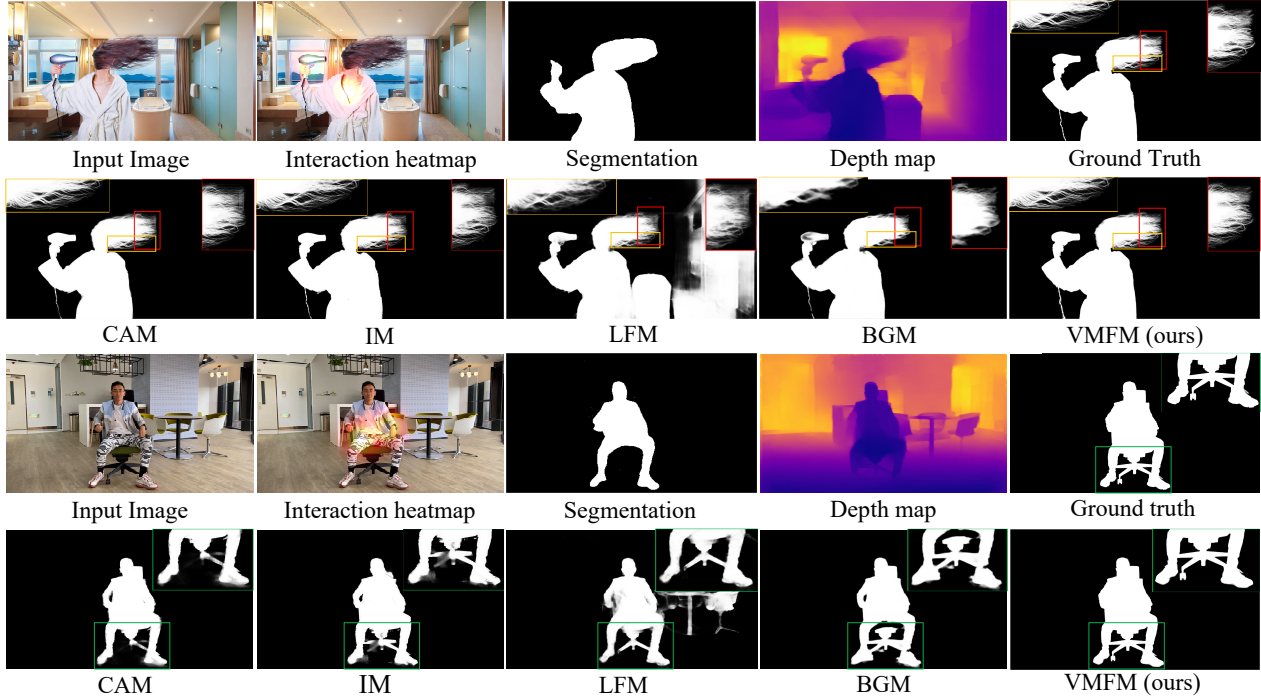


Figure 6: Comparison of foreground matting methods.

D-CL-RN vs.	much better	better	similar	worse	much worse
S-CL-RN	11.8%	22.4%	57.0%	6.2%	2.6%
LFM	60.7%	26.4%	12.9%	0%	0%
BGM	51.5%	39.2%	9.3%	0%	0%
CAM	32.8%	43.0%	20.1%	4.1%	0%
IM	35.6%	46.5%	14.0%	3.9%	0%

Table 4: User Study: compare D-CL-RN with our S-CL-RN and other competing algorithms on UFM70K.

show clearly that our foreground matting model D-CL-RN outperforms our S-CL-RN and all other competing methods by significant margin. Benefiting from the unique strength of each modality and complementary improvement across modalities, our VMFM method can capture more accurate semantic information of real images. Some representative visualizations are provided in Row 3 to 4 in Figure 6.

#### 4.5. Ablation Experiments

**Complementary Learning.** We evaluate complementary learning (CL) through dual constraint -  $L_{cs}^m$  and  $L_{dc}^m$ . Since two constraints are independent, we prune Virtual Multi-modality Foreground Matting (VMFM) network with one single CL constraint ( $L_{cs}^m$  or  $L_{dc}^m$ ). Table 3 shows the individual performance of  $L_{cs}^m$ ,  $L_{dc}^m$  and synergy of both when evaluating  $FP^m$  network (SFPNet or DFPNet) on the LFM40K test set. The CL constraint  $L_{cs}^m$  which aims

to exchange confidence score between cross-modality outputs, significantly benefits the alpha matting performance. Also,  $L_{dc}^m$  plays as an indispensable role in optimizing the collective performance of two modalities. Additionally, we demonstrate the performance gain after combining dual CL constraints.

**Refinement.** The quantitative results of matting refinement are shown in Table 3. Improvement in all quantitative metrics demonstrates that the refinement network (RN) is the last critical process that further enhances and sharpens the estimated alpha mattes. Some visualizations are provided in supplementary material.

## 5. Conclusion

In this paper, we present a matting technique that can extract high quality human-object interactive alpha mattes from a single RGB image. The training is self-supervised using our proposed complementary learning strategy which removes the bottleneck of gathering expensive alpha labels. Our method avoids using additional inputs, *e.g.* a green screen background, extra captured real background, or a manual trimap. Extensive experiments demonstrate that our model outperforms current state-of-the-art algorithms, not with more but fewer inputs. For future works, it may be possible to extend the human-object interactive method to 3D body reconstruction, which still only focuses on body voxels.



## References

- [1] 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 Pattern Recognition (CVPR)*, pages 29–37, 2017.
- [2] Ibraheem Alhashim and Peter Wonka. High quality monocular depth estimation via transfer learning. *arXiv e-prints*, abs/1812.11941, 2018.
- [3] Levin Anat, Rav-Acha Alex, and Lischinski Dani. Spectral matting. *IEEE transactions on pattern analysis and machine intelligence*, 30(10):1699–1712, 2008.
- [4] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.
- [5] Shaofan Cai, Xiaoshuai Zhang, Haoqiang Fan, Haibin Huang, Jiangyu Liu, Jiaming Liu, Jiaying Liu, Jue Wang, and Jian Sun. Disentangled image matting. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 8819–8828, 2019.
- [6] Quan Chen, Tiezheng Ge, Yanyu Xu, Zhiqiang Zhang, Xinxin Yang, and Kun Gai. Semantic human matting. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 618–626, 2018.
- [7] Qifeng Chen, Dingzeyu Li, and Chi-Keung Tang. Knn matting. *Proceedings of the IEEE transactions on pattern analysis and machine intelligence*, 35(9):2175–2188, 2013.
- [8] Xiaowu Chen, Dongqing Zou, Steven Zhiying Zhou, Qinpeng Zhao, and Ping Tan. Image matting with local and non-local smooth priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1902–1907, 2013.
- [9] Donghyeon Cho, Yu-Wing Tai, and Inso Kweon. Natural image matting using deep convolutional neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 626–643. Springer, 2016.
- [10] Shahrian Ehsan, Rajan Deepu, Price Brian, and Cohen Scott. Improving image matting using comprehensive sampling sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 636–643, 2013.
- [11] Chen Guanying, Han Kai, and Wong Kwan-Yee K. Tomnet: Learning transparent object matting from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9233–9241, 2018.
- [12] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, pages 2961–2969, 2017.
- [14] Kaiming He, Jian Sun, and Xiaoou Tang. Fast matting using large kernel matting laplacian matrices. In *Proceedings of the 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2165–2172. IEEE, 2010.
- [15] 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 (ICCV)*, pages 4130–4139, 2019.
- [16] Wang Jue and Cohen Michael F. Optimized color sampling for robust matting. In *2007 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8. IEEE, 2007.
- [17] Wang Jue and Cohen Michael F. *Image and video matting: a survey*. Now Publishers Inc, 2008.
- [18] He Kaiming, Rhemann Christoph, Rother Carsten, Tang Xiaoou, and Sun Jian. A global sampling method for alpha matting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2049–2056. IEEE, 2011.
- [19] Philip Lee and Ying Wu. Nonlocal matting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2193–2200. IEEE, 2011.
- [20] 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, 2007.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European conference on computer vision (ECCV)*, pages 740–755. Springer, 2014.
- [22] Jinlin Liu, Yuan Yao, Wendi Hou, Miaomiao Cui, Xuansong Xie, Changshui Zhang, and Xian-sheng Hua. Boosting semantic human matting with coarse annotations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8563–8572, 2020.
- [23] Hao Lu, Yutong Dai, Chunhua Shen, and Songcen Xu. Indices matter: Learning to index for deep image matting. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3266–3275, 2019.
- [24] Sebastian Lutz, Konstantinos Amliantis, and Aljosa Smolic. Alphagan: Generative adversarial networks for natural image matting. In *British Machine Vision Conference (BMVC)*, page 259. BMVA Press, 2018.
- [25] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, pages 2794–2802, 2017.
- [26] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016.
- [27] Richard J Qian and M Ibrahim Sezan. Video background replacement without a blue screen. In *Proceedings 1999 International Conference on Image Processing (Cat. 99CH36348)*, volume 4, pages 143–146. IEEE, 1999.
- [28] Yu Qiao, Yuhao Liu, Xin Yang, Dongsheng Zhou, Mingliang Xu, Qiang Zhang, and Xiaopeng Wei. Attention-guided hierarchical structure aggregation for image matting. In *The*

*IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

- [29] Soumyadip Sengupta, Vivek Jayaram, Brian Curless, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Background matting: The world is your green screen. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2291–2300, 2020.
- [30] Xiaoyong Shen, Xin Tao, Hongyun Gao, Chao Zhou, and Jiaya Jia. Deep automatic portrait matting. In *Proceedings of the European conference on computer vision (ECCV)*, pages 92–107. Springer, 2016.
- [31] Jian Sun, Jiaya Jia, Chi-Keung Tang, and Heung-Yeung Shum. Poisson matting. In *Proceedings of the ACM Special Interest Group on Computer Graphics (SIGGRAPH)*, pages 315–321, 2004.
- [32] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. Learning human-object interaction detection using interaction points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4116–4125, 2020.
- [33] Ning Xu, Brian Price, Scott Cohen, and Thomas Huang. Deep image matting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2970–2979, 2017.
- [34] Chuang Yung-Yu, Curless Brian, Salesin David H, and Szeliski Richard. A bayesian approach to digital matting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 264–271. IEEE, 2001.
- [35] Yunke Zhang, Lixue Gong, Lubin Fan, Peiran Ren, Qixing Huang, Hujun Bao, and Weiwei Xu. A late fusion cnn for digital matting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7469–7478, 2019.