

Matting Anything

Jiachen Li¹, Jitesh Jain¹, Humphrey Shi^{1,2}

¹SHI Labs @ Georgia Tech & Oregon & UIUC, ²Picsart AI Research (PAIR)

<https://github.com/SHI-Labs/Matting-Anything>

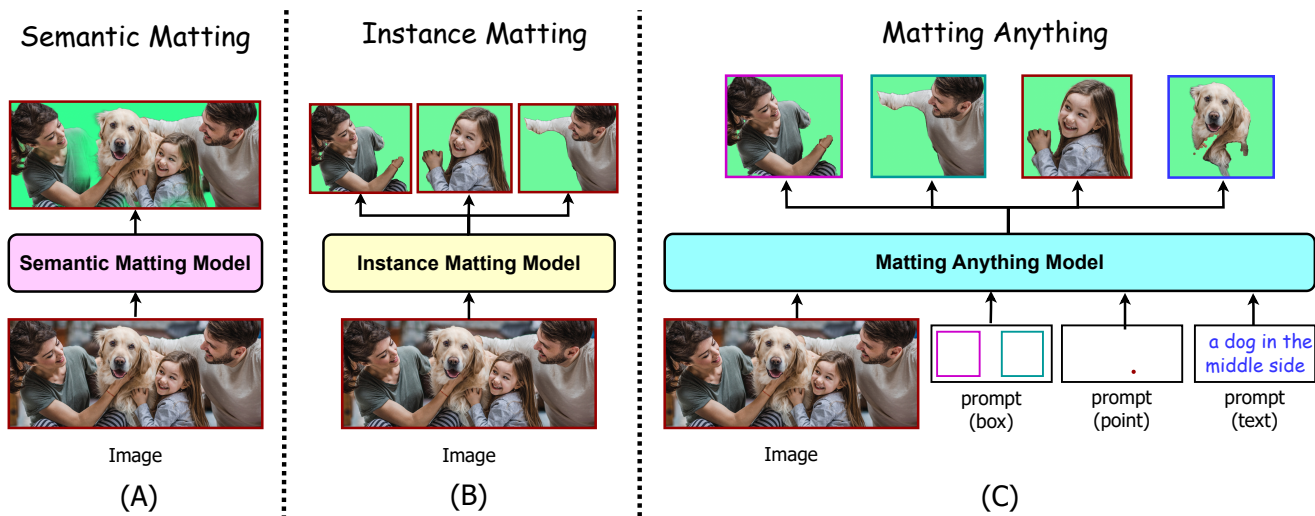


Figure 1. **Matting Anything Model (MAM)** offers a versatile framework capable of addressing various types of image matting scenarios with a single model. Compared to previous specialized models for (A) Semantic Matting, which outputs a single alpha matte of all instances in the foreground; (B) Instance Matting, which returns alpha mattes of all human instances; (C) Matting Anything Model can estimate the alpha matte of any target instance with user prompts as boxes, points, or text descriptions for interactive use by incorporating SAM [22]. It further reaches comparable performance to the specialized matting models on multiple benchmarks, and shows superior generalization ability with fewer parameters as a unified image matting model.

Abstract

In this paper, we propose the *Matting Anything Model (MAM)*, an efficient and versatile framework for estimating the alpha matte of any instance in an image with flexible and interactive visual or linguistic user prompt guidance. MAM offers several significant advantages over previous specialized image matting networks: (i) MAM is capable of dealing with various types of image matting, including semantic, instance, and referring image matting with only a single model; (ii) MAM leverages the feature maps from the *Segment Anything Model (SAM)* [22] and adopts a lightweight *Mask-to-Matte (M2M)* module to predict the alpha matte through iterative refinement, which has only 2.7 million trainable parameters. (iii) By incorporating SAM, MAM simplifies the user intervention required for the interactive use of image matting from the trimap to the

box, point, or text prompt. We evaluate the performance of MAM on various image matting benchmarks, and the experimental results demonstrate that MAM achieves comparable performance to the state-of-the-art specialized image matting models under different metrics on each benchmark. Overall, MAM shows superior generalization ability and can effectively handle various image matting tasks with fewer parameters, making it a practical solution for unified image matting. Our code and models are open-sourced at <https://github.com/SHI-Labs/Matting-Anything>.

1. Introduction

Image Matting, as a long-standing computer vision task, aims to estimate the alpha matte α given an input image I [46]. The matting target is mainly around human beings or other objects at the semantic level [23, 42, 50]. Recent

works have extended the scope of image matting to more complex scenarios like image instance matting [44], which requires instance-aware alpha matte predictions and referring image matting [29], which extracts the alpha matte given natural language description.

Previous deep learning-based image matting methods [29, 30, 37, 38, 44, 48, 52, 55, 60] have been proposed to address specific image matting tasks on corresponding benchmarks. These methods are tailored to individual datasets and lack the flexibility to handle various image matting tasks due to their fixed model designs. This limitation has hindered the development of more generalized and versatile image matting models. As a result, there is a growing interest in developing more adaptive and efficient image matting frameworks that can handle different types of image matting tasks with a single model.

Furthermore, previous image matting methods have relied on user-guided trimaps as auxiliary inputs to achieve accurate alpha matte predictions. Although some trimap-free methods have been proposed that use mask guidance or background images instead [40, 56], they are unable to estimate the alpha matte of the target instance based on the user request for interactive use. Therefore, it is crucial to develop a model that can achieve accurate alpha matte estimation without relying on user-guided trimaps, while also being capable of handling simple user requests in a flexible and efficient manner for interactive use. Such a model would significantly enhance the user experience by reducing the extra need for manual intervention.

Motivated by these limitations of image matting, we propose the Matting Anything Model (MAM), a versatile network that can estimate the alpha matte of any target instance with prompt-based user guidance in an image as shown in Figure 1. MAM leverages the recent Segment Anything Model (SAM) framework [22], which supports flexible prompting and outputs segmentation masks of any target instance for interactive use. Specifically, MAM takes the feature maps and mask outputs from SAM as inputs and adds a lightweight Mask-to-Matte (M2M) module to predict the alpha matte of the target instance. We trained MAM on a combination of five image matting datasets that cover different classes of instances, allowing the M2M module to learn generalizable features for image matting. During training, we randomly place target instances onto background images and use a pre-trained SAM to output mask predictions of the corresponding instances. The trainable M2M module then refines the mask by predicting multi-scale alpha mattes. Through an iterative refinement process based on the mask or the alpha matte, the multi-scale predictions are merged to obtain the final meticulous alpha matte.

We conducted extensive evaluations of MAM on six image matting benchmarks, including semantic image matting benchmark PPM-100 [41], AM2K [25] PM-10K [25],

the instance image matting benchmark RWP636 [56], HIM2K [44], and the referring image matting benchmark RefMatte-RW100 [29]. Our results demonstrate that MAM achieves performance comparable to that of state-of-the-art image matting models across all benchmarks under different evaluation metrics. The experimental results highlight the versatility and effectiveness of our proposed approach for handling various image matting tasks in an interactive and efficient manner.

2. Related Works

2.1. Image Matting

Given an image I , which can be view as a combination of foreground image F and background image B with coefficient alpha matte α ,

$$I = \alpha F + (1 - \alpha)B \quad (1)$$

Image Matting is to estimate α given only I as inputs. Traditional methods rely on a user-guided trimap, which explicitly annotates the absolute foreground area, absolute background area, and transition area. Then, sampling-based image matting solutions use low-level features to distinguish the transition areas by measuring the similarities between foreground and background neighbors [2, 3, 5, 9, 10, 12]. Recently, deep learning-based methods [28, 30, 37, 38, 48, 52, 55, 60] adopt neural networks to estimate the alpha matte in an end-to-end manner with trimap as auxiliary inputs. Some trimap-free methods use background image [40], mask guidance [33, 56], or segmentation data [6, 31] to make up the absence of trimap. When the image I contains multiple instances, the composition turns to

$$I = \sum_i^N \alpha_i F_i + (1 - \sum_i^N \alpha_i)B \quad (2)$$

α_i represents the alpha matte of instance i and InstMatt [44] adopt the target and reference mask as guidance to prediction instance-aware alpha matte prediction. Interactive matting methods [29, 49, 53] develop specialized models that use point, boxes, or text input to estimate the alpha matte of the target instance. MatAny [54] is a concurrent work that also adopts SAM for semantic image matting. SegAny [1] is an open-sourced project to improve the quality of image matting. In terms of video matting, trimap-free [24, 26, 27, 31, 41] methods are explored for real-time inference while the per-frame prediction quality is not comparable to image matting methods. However, these methods are designed for a certain scenario with corresponding benchmarks, which limits their potential to handle various image matting tasks and benchmarks.

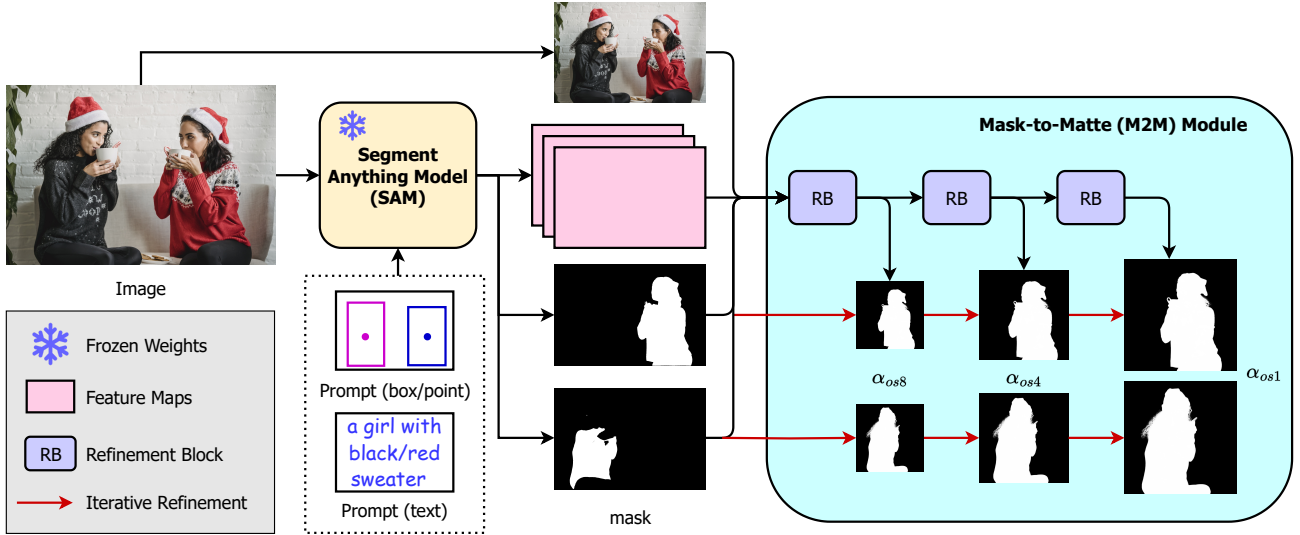


Figure 2. **Matting Anything Model Architecture.** The MAM architecture consists of a pre-trained SAM and an M2M module. Given an input image I , SAM generates the mask prediction for the target instance based on the box or point user prompt. The M2M module takes the concatenated inputs, including the image, mask, and feature maps, and produces multi-scale predictions α_{0s8} , α_{0s4} , and α_{0s1} . The iterative refinement process, detailed in Section 3, progressively improves the precision of the final meticulous alpha matte α , incorporating information from the multi-scale outputs.

2.2. Image Segmentation

Image segmentation is a close research area to image matting, while it predicts the binary mask of different instances in the image. Similar to image matting, many image segmentation methods are tailored for a specific image segmentation task, like semantic segmentation [4, 16], instance segmentation [14, 47], and panoptic segmentation [21, 45]. Recent works started to explore transformer-based frameworks [7, 13, 17, 18] for unified image segmentation. Language-guided segmentation frameworks [51, 59] look for text supervision to segment instance-aware masks. OneFormer [18] adopts a single transformer model to learn with a joint training strategy and performs universal segmentation across semantic, instance and panoptic segmentation and outperforms specialized models. SAM [22] takes a further step recently, which supports flexible prompting from users to segment any instance in an image for interactive use. Grounded-SAM [34] incorporates DINO with SAM to add text prompt support. Foundation models like SAM offer opportunities for other areas to develop versatile frameworks to support a range of applications.

3. Matting Anything

In this section, we provide an overview of the Matting Anything Model (MAM) architecture, which consists of two main components: the frozen Segment Anything Model (SAM) and the trainable Mask-to-Matte (M2M) module. We first provide a brief review of the SAM, which is designed to produce high-quality instance segmentation given

user-guided prompts. We then introduce the M2M module, which enables the transformation of the binary masks into high-quality alpha mattes. Finally, we describe how we connect the M2M module with the SAM to gradually build the end-to-end MAM.

3.1. Segment Anything Model

Segment Anything is a recently proposed foundation model for segmentation. Given an image $I \in \mathbb{R}^{3 \times H \times W}$, SAM uses a ViT-based image encoder to obtain deep feature maps $F \in \mathbb{R}^{C \times \frac{H}{16} \times \frac{W}{16}}$. Then, a variety of N input prompts are encoded by the prompt encoder and sent to the mask decoder with the feature maps. The mask decoder returns a set of mask candidates $m_i \in \mathbb{R}^{1 \times H \times W}$, $i \in N$ indicated by the input prompts. With its flexible prompting mechanism, SAM allows for interactive use and is easily adaptable for downstream tasks.

3.2. Mask-to-Matte

The Mask-to-Matte (M2M) module is an integral component of our Matting Anything Model (MAM) and is designed to convert instance-aware mask predictions from SAM into instance-aware alpha matte predictions efficiently and smoothly. To achieve this, we utilize the feature maps and mask predictions generated by SAM as auxiliary inputs to M2M. To improve the accuracy of our predictions, we adopt multi-scale branches for predicting the alpha matte and merge these predictions through an iterative refinement schedule.

Multi-Scale Prediction: Given an input image $I \in$

Task	Semantic Matting			Instance Matting				Referring Matting	
Benchmark	AM2K	PM-10K	PPM-100	HIM2K		RWP636		RefMatte-RW100	
Metric	SAD _{all} ↓	MAD _{all} ↓	MSE _{all} ↓	IMQ _{mad} ^{nat} ↑	IMQ _{mse} ^{nat} ↑	IMQ _{mad} ↑	IMQ _{mse} ↑	SAD _{all} ↓	MSE _{all} ↓
<i>Specialized Models</i>									
GFM-R [25]	10.89	6.7	-	-	-	-	-	-	-
GFM-D [25]	10.26	6.9	-	-	-	-	-	-	-
MODNet [20]	-	-	4.4	-	-	-	-	-	-
MGMatting [56]	-	-	-	57.98	71.12	30.64	53.16	-	-
InstMatt [44]	-	-	-	70.26	81.34	51.10	73.09	-	-
CLIPMat-B [29]	-	-	-	-	-	-	-	107.81	59.5
CLIPMat-L [29]	-	-	-	-	-	-	-	85.83	47.4
<i>Generalized Models</i>									
SAM [22]	25.00	25.7	10.8	61.15	74.01	49.87	56.92	33.51	17.9
MAM	17.30	15.4	4.6	68.78	81.67	54.40	76.45	29.24	15.1

Table 1. Comparisons between specialized matting models and MAM on various benchmarks. ↑ / ↓ means higher / lower values indicate better performance for the corresponding metric. Gray text refers to models specifically designed for these benchmarks. MAM shows clear improvements over SAM and superior generalization ability as a unified image matting model.

$\mathbb{R}^{3 \times H \times W}$, the pre-trained SAM model produces feature maps $F \in \mathbb{R}^{C \times \frac{H}{16} \times \frac{W}{16}}$ and mask prediction $m \in \mathbb{R}^{1 \times H \times W}$ on the target instance with prompt guidance. We concatenate the rescaled image, mask, and feature maps to form the input $F_{m2m} \in \mathbb{R}^{(C+4) \times \frac{H}{8} \times \frac{W}{8}}$ to the M2M module. M2M employs several refinement blocks [8, 56], which contain connected self-attention layer [57], batchnorm layer, and activation layer, to generate alpha matte predictions at 1/8 resolution, denoted as $\alpha_{os8} \in \mathbb{R}^{1 \times \frac{H}{8} \times \frac{W}{8}}$. The feature maps are then upsampled to higher resolutions to make alpha matte predictions at 1/4 and full resolution, denoted as $\alpha_{os4} \in \mathbb{R}^{1 \times \frac{H}{4} \times \frac{W}{4}}$ and $\alpha_{os1} \in \mathbb{R}^{1 \times H \times W}$, respectively. The multi-scale predictions enable MAM to handle objects of varying scales and provide finer-grained alpha mattes for detailed object extraction.

Iterative Refinement To improve the accuracy of global and local predictions, we use an iterative refinement process. We first compute weight maps w_{os8} , w_{os4} , and w_{os1} that highlight different areas of the image during training like trimaps. These weight maps are used to compute losses for each scale of prediction, with w_{os8} emphasizing the entire image for α_{os8} predictions, w_{os4} filtering out the background for α_{os4} predictions, and w_{os1} focusing only on the transition areas. During inference, we gradually merge the predictions of α_{os8} , α_{os4} , and α_{os1} with the mask predictions m from SAM to obtain the final alpha matte prediction $\alpha \in \mathbb{R}^{1 \times H \times W}$.

3.3. Matting Anything Model

After the development of the Mask-to-Matte (M2M) module, we integrate it with the Segment Anything Model (SAM) to enable end-to-end training and inference for the Matting Anything Model (MAM). This integration allows for a comprehensive and unified framework that handles

the entire matting process, from feature extraction to alpha matte prediction.

Multi-Dataset Training To ensure the robustness and versatility of our Matting Anything Model (MAM), we adopt a multi-dataset training approach that encompasses diverse foreground instances and background images from various image matting datasets. This selection allows us to cover a wide range of instance classes and background scenarios, enhancing the model’s ability to handle different types of instances and backgrounds effectively. During the training process, we create composite images by combining a foreground instance $F \in \mathbb{R}^{3 \times H \times W}$ with its corresponding ground truth alpha matte $\alpha_{gt} \in \mathbb{R}^{1 \times H \times W}$ and a background image $B \in \mathbb{R}^{3 \times H \times W}$. The composition is performed using the equation $I = \alpha_{gt}F + (1 - \alpha_{gt})B$. We then extract the bounding box (x_0, y_0, x_1, y_1) that encapsulates the instance of interest within the composite image. Then, we send the image I and the bounding box as a prompt to the pre-trained SAM, which returns the mask prediction of the instance. Then, we concatenate the image, mask and feature maps, and send them to the M2M module, which further returns the multi-scale alpha matte predictions $\alpha_{os8}, \alpha_{os4}, \alpha_{os1}$. The loss \mathcal{L} is computed between the multi-scale predictions and ground truth α_{gt} as

$$\mathcal{L}(\alpha_{gt}, \alpha_{os1}, \alpha_{os4}, \alpha_{os8}) = \lambda_{\mathcal{L}_1} \mathcal{L}_1 + \lambda_{\mathcal{L}_{Lap}} \mathcal{L}_{Lap} \quad (3)$$

\mathcal{L}_1 is L1 loss and \mathcal{L}_{Lap} is Laplacian loss used in [15, 31, 43]. The coefficients $\lambda_{\mathcal{L}_1}$ and $\lambda_{\mathcal{L}_{Lap}}$ control the contribution of each loss term, respectively. Both loss terms are computed on multi-scale predictions as

$$\mathcal{L}_1 = \sum_{i=1,4,8} \mathcal{L}_1(\alpha_{gt}, \alpha_{osi}) \quad (4)$$

$$\mathcal{L}_{Lap} = \sum_{i=1,4,8} \mathcal{L}_{Lap}(\alpha_{gt}, \alpha_{osi}) \quad (5)$$

Method	AM2K / PM-10K				
	SAD _{all} ↓	MSE _{all} ↓	MAD _{all} ↓	Grad _{all} ↓	SAD _{tri} ↓
SHM [6]	17.81 / 16.64	6.8 / 6.9	10.2 / 9.7	12.54 / 14.54	10.26 / 8.53
LFM [58]	36.12 / 37.51	11.6 / 15.2	21.0 / 15.2	21.06 / 21.82	19.68 / 16.36
HATT [38]	28.01 / 22.66	5.5 / 3.8	16.1 / 13.1	18.29 / 15.16	13.36 / 9.32
SHMC [33]	61.50 / 57.85	27.0 / 29.1	35.6 / 34.0	37.00 / 37.28	35.23 / 23.04
GFM-R [25]	10.89 / 11.52	2.9 / 3.8	6.4 / 6.7	10.00 / 13.07	9.15 / 8.00
GFM-D [25]	10.26 / 11.89	2.9 / 4.1	5.9 / 6.9	8.82 / 12.90	8.24 / 7.80
SAM [22]	25.00 / 44.11	10.8 / 28.8	14.8 / 25.7	60.01 / 24.56	20.72 / 31.96
MAM	17.30 / 25.82	3.5 / 9.2	10.1 / 15.4	10.65 / 14.22	15.67 / 23.99

Table 2. Results on the semantic image matting benchmark AM2K and PM-10K. Metrics with *all* and *tri* as subscript indicates the evaluation of the whole image and the transition area, separately. ↓ means lower values indicate better performance for the metric.

Method	MSE _{all} ↓	MAD _{all} ↓
DIM [52]	11.5	17.8
FDMPA [60]	10.1	16.0
LFM [58]	9.4	15.8
SHM [6]	7.2	15.2
HATT [38]	6.7	13.7
BSHM [33]	6.3	11.4
MODNet [20]	4.4	8.6
SAM [22]	10.8	13.8
MAM	4.6	9.9

Table 3. Results on the semantic image matting benchmark PPM-100.

Multi-Benchmark Inference During the inference phase, we conducted extensive evaluations of the Matting Anything Model (MAM) on multiple image matting benchmarks to assess its generality and adaptability. Given an input image I , SAM produced the initial mask prediction $m \in \mathbb{R}^{1 \times H \times W}$, which captured the rough delineation of the instance. Subsequently, M2M contributed to the refinement of the alpha matte prediction by providing multi-scale predictions α_{os8} , α_{os4} , and α_{os1} . Then, following the iterative refinements, we progressively updated the predictions by replacing the corresponding regions in the mask prediction m with the respective multi-scale predictions that demonstrated positive weight maps, while in some simple cases the replacement is directly done upon α_{os8} instead of m . This iterative refinement allowed us to refine the alpha matte estimation iteratively and enhance the precision of the final prediction $\alpha \in \mathbb{R}^{1 \times H \times W}$.

4. Experiments

4.1. Implementation Details

Training Datasets During the training process, we randomly select foreground instances from several image matting datasets, including Adobe Image Matting dataset [52], Distinctions-646 [55], AM2K [25], Human-2K [35], and RefMatte [29], to ensure a diverse range of instance classes. For background images, we select them from two datasets:

COCO [32] and BG20K [25] to provide a mix of both real-world and synthetic backgrounds.

Evaluation Benchmarks To evaluate the adaptive ability of MAM, we test it on a variety of image matting benchmarks including the semantic image matting benchmarks PPM-100 [41], AM2K [25], PM-10K [25], the instance image matting benchmark RWP636 [56], HIM2K [44], and the referring image matting benchmark RefMatte-RW100 [29]. The box prompt is used for all benchmarks and the point prompt is only used in RefMatte-RW100. This comprehensive evaluation allows us to assess the generalization capability of MAM across various image matting tasks and benchmarks.

Evaluation Metrics We evaluate the accuracy of predicted alpha matte for MAM with commonly adopted evaluation metrics. Specifically, we employ Mean Absolute Difference (MAD), Sum of Absolute Difference (SAD), Mean Squared Error (MSE), Gradient (Grad), and Connectivity (Conn) [39] as corresponding evaluation metrics. We scale MAD, MSE, Grad, and Conn by 10^3 , 10^3 , 10^{-3} , and 10^{-3} , respectively. Lower values indicate better performance for these metrics. Additionally, for instance-aware matting, we utilize Instance Matting Quality (IMQ) [44], which takes recognition and matting accuracy into consideration simultaneously. Higher values indicate better performance for the IMQ metric.

Experimental Settings We trained MAM on a combination of training datasets using 8 RTX A6000 GPUs, with a batch size of 10 images per GPU. Each image was a combination of a randomly selected foreground instance and a background image. Images were cropped to a size of 1024×1024 and sent to a pre-trained ViT-B based SAM [22] with a bounding box prompt of the target instance. The feature maps and masks output by SAM were then fed into the M2M module for alpha matte prediction. We employed the Adam optimizer with $\beta_1 = 0.5$ and $\beta_2 = 0.99$, trained for 20,000 iterations with warm-up for the first 4,000 iterations. The weight map w_{os8} is always 1 at all pixels during training, while w_{os4} changes to

Method	Model Size	Synthetic Subset \uparrow				Natural Subset \uparrow			
		IMQ _{mad}	IMQ _{mse}	IMQ _{grad}	IMQ _{conn}	IMQ _{mad}	IMQ _{mse}	IMQ _{grad}	IMQ _{conn}
Mask RCNN [14]	44.3 M	18.37	25.65	0.45	19.07	24.22	33.74	2.27	26.65
CascadePSP [8]	67.7 M	40.85	51.64	29.59	43.37	64.58	74.66	60.02	67.20
GCA [30]	25.2 M	37.76	51.56	38.33	39.90	45.72	61.40	44.77	48.81
SIM [42]	46.5 M	43.02	52.90	40.63	44.29	54.43	66.67	49.56	58.12
FBA [11]	34.7 M	36.01	51.44	37.86	38.81	34.81	48.32	36.29	37.23
MGMating [56]	+ 29.6 M	51.67	67.08	53.03	55.38	57.98	71.12	66.53	60.86
InstMatt [44]	+ 29.7 M	63.59	78.14	64.50	67.71	70.26	81.34	74.90	72.60
SAM [22]	93.7 M	49.69	61.44	4.34	51.84	61.15	74.01	13.64	65.85
MAM	+ 2.7 M	54.15	68.01	30.47	55.40	68.78	81.67	51.79	72.62

Table 4. Results on the instance image matting benchmark HIM2K. Metrics with *mad*, *mse*, *grad*, and *conn* as subscript indicates the similarity metrics for IMQ are MAD, MSE, Gradient, and Connectivity, separately. \uparrow means higher values indicate better performance for the IMQ metric. MAM shows clear improvements over SAM under different metrics with only 2.7M extra trainable parameters, much lighter compared to other mask-guided methods like MGMating and InstMatt.

Method	IMQ _{mad} \uparrow	IMQ _{mse} \uparrow
Mask RCNN [14]	20.26	25.36
CascadePSP [8]	42.20	52.91
GCA [30]	33.87	46.47
SIM [42]	34.66	46.60
FBA [11]	35.00	47.54
MGMating [56]	30.64	53.16
InstMatt [44]	51.10	73.09
SAM [22]	49.87	56.92
MAM	54.40	76.45

Table 5. Results on the instance matting benchmark RWP636.

the mask guidance from SAM after the 4,000 iterations and w_{os1} changes to the boundary of α_{os4} after the 4,000 iterations as well. We set 3 refinement blocks for the prediction of α_{os8} , 3 refinement blocks for the prediction of α_{os4} , and 2 refinement blocks for the prediction of α_{os1} . As a result, the total trainable parameters of MAM is 2.7 million parameters. We applied cosine learning rate decay with an initial learning rate of 0.001 during training. During inference, we used a single GPU with a batch size of 1. Each image was resized to have its longer side at 1024 pixels and its shorter side was padded to 1024 pixels before being sent to MAM for alpha matte prediction of the target instance.

4.2. Main Results

Specialized vs Unified Model We present a high-level comparison between specialized image matting models and MAM on the semantic, instance, and referring image matting benchmarks in Table 1. It shows that MAM has clear improvements over SAM on all benchmarks. Furthermore, MAM shows comparable performance to each specialized image matting model and even reaches better performance on the HIM2K, RWP635, and RefMatte-RW100, which makes it a practical and feasible solution to unified image matting.

Method	Prompt	SAD _{all} \downarrow	MSE _{all} \downarrow	MAD _{all} \downarrow
MDETR [19]	text	131.58	67.5	75.1
CLIPSeg [36]	text	211.86	117.8	122.2
CLIPMat [29]	text	107.81	59.5	62.0
SAM [22]	text	122.76	67.9	69.0
MAM	text	120.10	65.9	67.5
SAM [22]	point	214.19	123.8	124.9
MAM	point	168.82	89.6	97.7
SAM [22]	box	33.51	17.9	19.0
MAM	box	29.24	15.1	16.6

Table 6. Results on the referring image matting benchmark RefMatte-RW100. MAM with box prompt can reach significantly better performance than with the text prompt.

Method	Model Size	Natural Subset	
		IMQ _{mad}	IMQ _{mse}
SAM [22]	93.7 M	50.47	61.66
+ Mask-Select	93.7 M	61.15	74.01
MAM Baseline	1.0 M	52.82	71.82
+ Multi-Scale Prediction	2.7 M	60.11	74.74
+ Iterative Refinement	2.7 M	65.44	78.93
+ Multi-Dataset Training	2.7 M	68.37	81.56

Table 7. Ablation study of MAM on the HIM2K benchmark. The MAM Baseline is built upon the SAM model with the box prompt.

Semantic Image Matting We evaluate the performance of MAM on three semantic image matting benchmarks: PPM-100 [41], AM2K [25], and PM-10K [25], as presented in Table 3 and Table 2. The iterative refinement process is based on the α_{os8} prediction for all three benchmarks. On the PPM-100 benchmark, MAM achieves improvements of 6.2 MSE_{all} and 3.9 MAD_{all} over SAM. Similarly, on the AM2K benchmark, MAM outperforms SAM with enhancements of 7.64 SAD_{all}, 4.4 MSE_{all}, 4.5 MAD_{all}, 41.54 Grad_{all}, and 7.59 SAD_{tri}.

Instance Image Matting In Table 4 and Table 5, We eval-

uate MAM on two instance image matting benchmarks: HIM2K [44] and RWP636 [56]. For HIM2K, the iterative refinement is based on prediction mask m since it contains multiple instances per image and starting from m removes false positive predictions. Compared to other state-of-the-art methods on HIM2K, MAM reaches comparable performance with only 2.7 M extra trainable parameters, which is only 10% of the specialized models like MGMatting and InstMatt, which use the mask guidance from Mask RCNN. On the RWP636 benchmark, we apply the iterative refinement from α_{os8} and MAM reaches the new state-of-the-art with 54.40 IMQ_{mad} and 76.45 IMQ_{mse} .

Referring Image Matting In Table 6, we present the evaluation of MAM on the RefMatte-RW100 benchmark [29], a recently introduced referring image matting benchmark. While previous methods rely on text prompts for referring image matting, we leverage the bounding boxes and text descriptions as the prompts for SAM. Considering the text prompt for SAM has not been released yet, we use Grounded-SAM [34] to support text prompt guidance. Remarkably, MAM achieves superior performance when utilizing the bounding box as the prompt for SAM, surpassing the text-based methods CLIPSeg and CLIPMat by a significant margin.

4.3. Ablation Study

We conduct comprehensive ablation studies on the M2M module of MAM, considering that SAM remains frozen during the training process. To assess the performance of MAM, we select the real-world subset of the HIM2K benchmark.

SAM on HIM2K We begin by evaluating the pre-trained ViT-B-based SAM using bounding boxes and points as prompts for the target instance. SAM with box-based prompts significantly outperforms the point-based prompts and the final mask output is selected based on the mask with the highest Intersection over Union (IoU) score with the bounding box. SAM demonstrates strong performance on the HIM2K benchmark, achieving 61.15 IMQ_{mad} and 74.01 IMQ_{mse} on the natural subset.

Building MAM We then construct the M2M baseline by integrating the M2M module, which takes SAM’s mask and feature maps, as well as the image, as inputs. This baseline, comprising 3 connected refinement blocks and predicting at 1/16 resolution, yields inferior performance compared to SAM, as the low-resolution predictions lack fine details of the alpha matte. However, by gradually incorporating multi-scale predictions and iterative refinement, as described in Section 3.2, the performance of MAM improves. Additionally, the adoption of multi-dataset training, as outlined in Section 3.3, further enhances performance, resulting in 68.37 IMQ_{mad} and 81.56 IMQ_{mse} on the natural subset. Subsequently, we assess MAM’s performance




Figure 3. Visualizations of alpha matte predictions from MGMatting and MAM. Improvements are highlighted in the red boxes.

on other benchmarks without retraining to validate its generality and adaptability.

4.4. Visualization

In Figure 3, we compare matting performance between MGMatting and MAM of images that contain multiple instances. They both leverage mask guidance from SAM. It shows that MAM is able to give more accurate alpha matte predictions with only 10% parameters compared to MGMatting under the same mask guidance. It also has fewer false positive predictions in other instances. In Figure 4, we further provide visualizations of the mask and alpha matte predictions from SAM and MAM. These images are selected from the semantic image matting benchmarks and contain a single instance that can be a person, animal, or transparent object. The visualizations demonstrate that MAM achieves significantly improved predictions in the transition areas without the trimap guidance, which highlights the superior performance of MAM in refining and enhancing the quality of alpha matte predictions.

4.5. Limitation

One limitation of MAM is its dependence on the feature maps and mask predictions generated by SAM. If SAM produces incorrect mask predictions, such as returning the mask of an irrelevant instance like the bottom left case in

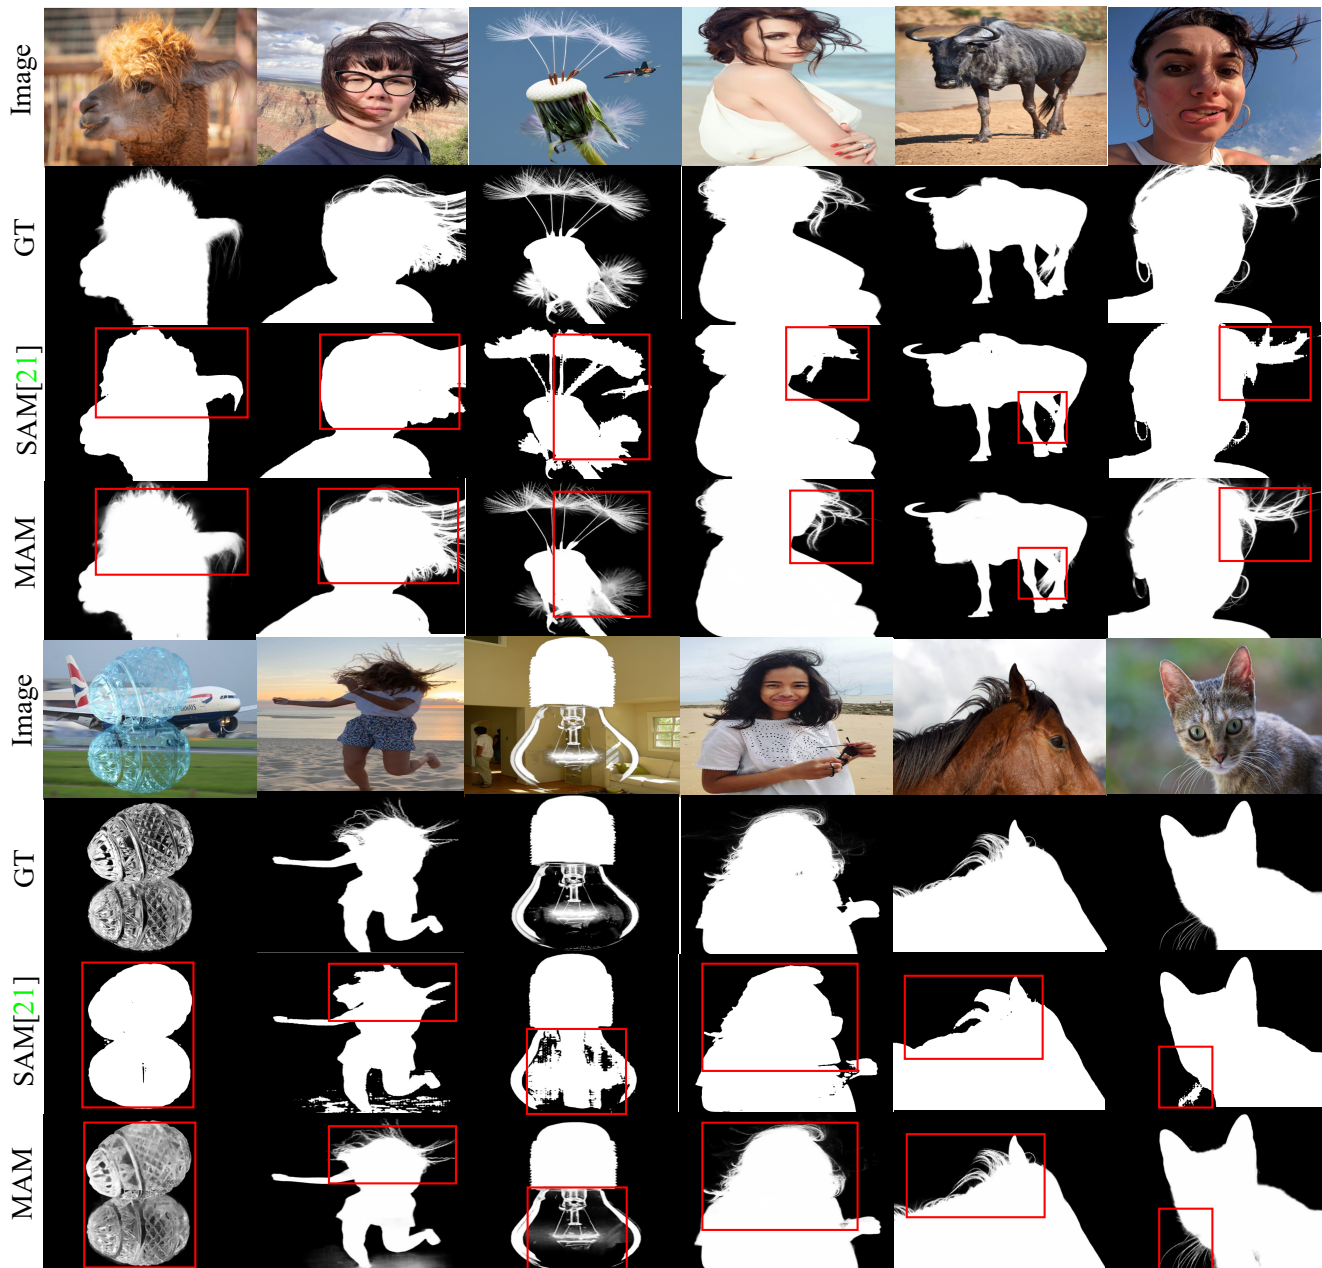


Figure 4. Visualizations of mask and alpha matte predictions from SAM and MAM. Improvements are highlighted in the red boxes.

Figure 4, MAM faces challenges in rectifying these errors. The iterative refinement process may inadvertently propagate such errors, leading to inaccurate final predictions. Therefore, addressing the issue of enabling MAM to effectively correct instance-level prediction errors remains an open question that requires further investigation.

5. Conclusion

In this paper, we introduce Matting Anything Model (MAM), which uses the Segment Anything Model (SAM)

as a guidance module with a lightweight Mask-to-Matte (M2M) module to refine the mask output into the alpha matte of the target instance. M2M is designed to handle various image matting tasks, including semantic, instance, and referring image matting, using a single model based on user prompts including points, boxes, and text. We evaluate MAM on six image matting benchmarks and demonstrate that it achieves comparable performance to the specialized state-of-the-art methods under various evaluation metrics. Our proposed model offers a more versatile and efficient solution for interactive and unified image matting.

References

- [1] segment-matting. <https://github.com/haibingtong/segment-matting>. 2
- [2] 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*, 2017. 2
- [3] Xue Bai and Guillermo Sapiro. A geodesic framework for fast interactive image and video segmentation and matting. In *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 2007. 2
- [4] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 3
- [5] Qifeng Chen, Dingzeyu Li, and Chi-Keung Tang. Knn matting. *IEEE transactions on pattern analysis and machine intelligence*, 2013. 2
- [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*, 2018. 2, 5
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *arXiv preprint arXiv:2112.01527*, 2021. 3
- [8] Ho Kei Cheng, Jihoon Chung, Yu-Wing Tai, and Chi-Keung Tang. Cascadepsp: Toward class-agnostic and very high-resolution segmentation via global and local refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 4, 6
- [9] Yung-Yu Chuang, Brian Curless, David H Salesin, and Richard Szeliski. A bayesian approach to digital matting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*. IEEE, 2001. 2
- [10] Xiaoxue Feng, Xiaohui Liang, and Zili Zhang. A cluster sampling method for image matting via sparse coding. In *European Conference on Computer Vision*. Springer, 2016. 2
- [11] Marco Forte and François Pitié. f, b, α matting. *arXiv preprint arXiv:2003.07711*, 2020. 6
- [12] Leo Grady, Thomas Schiwietz, Shmuel Aharon, and Rüdiger Westermann. Random walks for interactive alpha-matting. In *Proceedings of VIIP*, 2005. 2
- [13] Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 3
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 3, 6
- [15] Qiqi Hou and Feng Liu. Context-aware image matting for simultaneous foreground and alpha estimation. In *ICCV*, 2019. 4
- [16] Zilong Huang, Xinggang Wang, Yunchao Wei, Lichao Huang, Humphrey Shi, Wenyu Liu, and Thomas S. Huang. Ccnet: Criss-cross attention for semantic segmentation. In *TPAMI*, 2020. 3
- [17] Jitesh Jain, Anukriti Singh, Nikita Orlov, Zilong Huang, Jiachen Li, Steven Walton, and Humphrey Shi. Semask: Semantically masked transformers for semantic segmentation. *arXiv preprint arXiv:2112.12782*, 2021. 3
- [18] Jitesh Jain, Jiachen Li, MangTik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. Oneformer: One transformer to rule universal image segmentation. *CVPR*, 2023. 3
- [19] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdet: Modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. 6
- [20] Zhanghan Ke, Jiayu Sun, Kaican Li, Qiong Yan, and Rynson WH Lau. Modnet: Real-time trimap-free portrait matting via objective decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1140–1147, 2022. 4, 5
- [21] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 3
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 1, 2, 3, 4, 5, 6
- [23] Jizhizi Li, Jing Zhang, and Dacheng Tao. Deep automatic natural image matting. *arXiv preprint arXiv:2107.07235*, 2021. 1
- [24] Jiachen Li, Vidit Goel, Marianna Ohanyan, Shant Navasardyan, Yunchao Wei, and Humphrey Shi. Vmformer: End-to-end video matting with transformer. *arXiv preprint arXiv:2208.12801*, 2022. 2
- [25] Jizhizi Li, Jing Zhang, Stephen J Maybank, and Dacheng Tao. Bridging composite and real: towards end-to-end deep image matting. *International Journal of Computer Vision*, 2022. 2, 4, 5, 6
- [26] Jiachen Li, Roberto Henschel, Vidit Goel, Marianna Ohanyan, Shant Navasardyan, and Humphrey Shi. Video instance matting. *arXiv preprint arXiv:2311.04212*, 2023. 2
- [27] Jiachen Li, Marianna Ohanyan, Vidit Goel, Shant Navasardyan, Yunchao Wei, and Humphrey Shi. VideoMatt: A simple baseline for accessible real-time video matting. In *CVPR Workshops*, 2023. 2
- [28] Jizhizi Li, Jing Zhang, and Dacheng Tao. Deep image matting: A comprehensive survey. *ArXiv*, 2023. 2
- [29] Jizhizi Li, Jing Zhang, and Dacheng Tao. Referring image matting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22448–22457, 2023. 2, 4, 5, 6, 7
- [30] Yaoyi Li and Hongtao Lu. Natural image matting via guided contextual attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2, 6
- [31] Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. Robust high-resolution video matting with temporal guidance. *arXiv preprint arXiv:2108.11515*, 2021. 2, 4

- [32] 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 *European conference on computer vision*. Springer, 2014. 5
- [33] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2, 5
- [34] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 3, 7
- [35] Yuhao Liu, Jiake Xie, Xiao Shi, Yu Qiao, Yujie Huang, Yong Tang, and Xin Yang. Tripartite information mining and integration for image matting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 5
- [36] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022. 6
- [37] GyuTae Park, SungJoon Son, JaeYoung Yoo, SeHo Kim, and Nojun Kwak. Matteformer: Transformer-based image matting via prior-tokens. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [38] Yu Qiao, Yuhao Liu, Xin Yang, Dongsheng Zhou, Mingliang Xu, Qiang Zhang, and Xiaopeng Wei. Attention-guided hierarchical structure aggregation for image matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2, 5
- [39] Christoph Rhemann, Carsten Rother, Jue Wang, Margrit Gelautz, Pushmeet Kohli, and Pamela Rott. A perceptually motivated online benchmark for image matting. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009. 5
- [40] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [41] Jiayu Sun, Zhanghan Ke, Lihe Zhang, Huchuan Lu, and Rynson WH Lau. Modnet-v: Improving portrait video matting via background restoration. *arXiv preprint arXiv:2109.11818*, 2021. 2, 5, 6
- [42] Yanan Sun, Chi-Keung Tang, and Yu-Wing Tai. Semantic image matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 6
- [43] Yanan Sun, Guanzhi Wang, Qiao Gu, Chi-Keung Tang, and Yu-Wing Tai. Deep video matting via spatio-temporal alignment and aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 4
- [44] Yanan Sun, Chi-Keung Tang, and Yu-Wing Tai. Human instance matting via mutual guidance and multi-instance refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 4, 5, 6, 7
- [45] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 3
- [46] Jue Wang and Michael F Cohen. Image and video matting: a survey. 2008. 1
- [47] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. *Advances in Neural information processing systems*, 2020. 3
- [48] Yu Wang, Yi Niu, Peiyong Duan, Jianwei Lin, and Yuanjie Zheng. Deep propagation based image matting. In *IJCAI*, 2018. 2
- [49] Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Hanqing Zhao, Weiming Zhang, and Nenghai Yu. Improved image matting via real-time user clicks and uncertainty estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15374–15383, 2021. 2
- [50] Bo Xu, Jiake Xie, Han Huang, Ziwen Li, Cheng Lu, Yong Tang, and Yandong Guo. Situational perception guided image matting. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. 1
- [51] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3
- [52] Ning Xu, Brian Price, Scott Cohen, and Thomas Huang. Deep image matting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 5
- [53] Stephen DH Yang, Bin Wang, Weijia Li, YiQi Lin, and Conghui He. Unified interactive image matting. *arXiv preprint arXiv:2205.08324*, 2022. 2
- [54] Jingfeng Yao, Xinggang Wang, Lang Ye, and Wenyu Liu. Matte anything: Interactive natural image matting with segment anything models. *arXiv preprint arXiv:2306.04121*, 2023. 2
- [55] Haichao Yu, Ning Xu, Zilong Huang, Yuqian Zhou, and Humphrey Shi. High-resolution deep image matting. *AAAI*, 2021. 2, 5
- [56] Qihang Yu, Jianming Zhang, He Zhang, Yilin Wang, Zhe Lin, Ning Xu, Yutong Bai, and Alan Yuille. Mask guided matting via progressive refinement network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2, 4, 5, 6, 7
- [57] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International conference on machine learning*. PMLR, 2019. 4
- [58] 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/CVF conference on computer vision and pattern recognition*, pages 7469–7478, 2019. 5

- [59] Qiang Zhou, Yuang Liu, Chaohui Yu, Jingliang Li, Zhibin Wang, and Fan Wang. Lmseg: Language-guided multi-dataset segmentation. *arXiv preprint arXiv:2302.13495*, 2023. 3
- [60] Bingke Zhu, Yingying Chen, Jinqiao Wang, Si Liu, Bo Zhang, and Ming Tang. Fast deep matting for portrait animation on mobile phone. In *Proceedings of the 25th ACM international conference on Multimedia*, 2017. 2, 5